The European Best Bid and Offer (EBBO): From Fragmented Feeds to a Consolidated Tape*

Lisa Kaminski[†‡], Ralf Laschinger[‡], Ryan Riordan[‡]

March 13, 2026

Abstract

We construct the first nanosecond-level European Best Bid and Offer (EBBO) by consolidating direct-feed quotes from ten major venues for Eurostoxx50 constituents between 2020 and 2024. The EBBO reveals that consolidation halves quoted spreads and nearly triples displayed depth, yet only 79–86% of marketable orders execute at the best available price. Geographic latency, routing frictions, and incomplete connectivity generate welfare losses of more than €50 million per year. These results quantify some of the costs of fragmentation and demonstrate how a consolidated European benchmark could enhance execution quality, transparency, and competition—directly informing the design of the Markets in Financial Instruments Directive III's (MiFID III) Consolidated Tape.

Keywords: European Best Bid and Offer (EBBO) · market fragmentation · consolidated tape · competition · high-frequency

JEL Codes: G14 · G18 · L10

*This draft is preliminary. We are very grateful for Phil Mackintosh's remarks and thank the Nasdaq's data team for valuable comments: Eugenio Piazza, Satchit Sagade, and Nicole Torskiy. We additionally thank BMLL, in particular Arnaud Bossard and Stacey Ellul, as well as Björn Hagströmer, Albert Menkveld, Andreas Park and Janak Patal. The comments we received have been invaluable in improving our work. There are no conflicts of interest to disclose.

†Corresponding author: Lisa Kaminski, Schackstraße 4, 80539 Munich, Germany; E-Mail: lisa.kaminski@lmu.de.

‡Institute for Financial Innovation and Technology, LMU Munich School of Management, Germany.

Trading in modern equity markets occurs across multiple venues designed to foster competition and innovation (Madhavan, 1995; O'Hara and Ye, 2011). This structure can reduce explicit trading costs and stimulate technological progress, yet it also fragments liquidity and price discovery. When the same security trades simultaneously at different prices across venues, traders and regulators face a fundamental question. Should markets be left to compete independently, or should regulators enforce a unified view of prices and execution? The answer determines how effectively markets allocate capital, how well competition protects investors, and how transparently market outcomes can be monitored. Without a consolidated benchmark and transparent reference prices, competition fragments rather than integrates liquidity, reducing overall execution quality.

The US addressed these issues in 2005 with Regulation National Market System (Reg NMS). Reg NMS mandates continuous consolidation of exchange quotes into the National Best Bid and Offer (NBBO) and requires routing to the best available price across venues (Hasbrouck and Saar, 2013). This regime effectively transformed a collection of individual exchanges into a single national market with a common pre-trade benchmark. Europe adopted the opposite approach. Since the Markets in Financial Instruments Directive I (MiFID I, 2007) under the supervision of the European Securities and Markets Authority (ESMA), trading has been distributed across Regulated Markets (RMs), Multilateral Trading Facilities (MTFs), and Systematic Internalisers (SIs), but without a unified and binding consolidated quote (ESMA, 2007). Instead, brokers are subject to a "best execution" obligation to take "reasonable steps" toward favourable outcomes. Europe therefore retains the benefits of venue competition but lacks the transparency and accountability afforded by a consolidated price tape and broker execution-quality reporting (Dyhrberg et al., 2025).

The absence of a European consolidated price tape stands as both, a policy and a measurement gap. With roughly €15.6 trillion in annual equity trading (2024) dispersed across more than forty venues, Europe is among the most fragmented markets globally. Without a unified best-bid-and-offer standard comparable to the US-NBBO, regulators are unable to directly evaluate execution quality, researchers cannot quantify systematic welfare losses from fragmented liquidity or assess how market design affects market efficiency. To close this measurement gap, we construct the first nanosecond-level *European Best Bid and Offer* (EBBO) by consolidating direct-feed quotes from ten major venues for all Eurostoxx50 constituents between 2020 and 2024. The EBBO integrates 1.17 billion trades and 1.31 billion quote updates from venues spanning Madrid to Stockholm. Given Europe's geographically dispersed trading landscape, we adjust local timestamps for venue-pair transmission latencies to align quotes in physical time. The resulting relative EBBOs (REBBOs) represent each venue's latency-adjusted view of the consolidated market, capturing the best available prices as observed through the lens of the respective venue.[1] This latency-adjusted consolidation provides a European pre-trade price benchmark analogous to the US-NBBO which thereby enables the systematic assessment of execution quality, quote coherence, and market efficiency in Europe. Using this benchmark, we quantify the economic costs of fragmentation and identify its structural drivers—geographic latency, routing frictions, and incomplete venue connectivity—that jointly determine how effectively liquidity is shared across markets.

Market microstructure theory highlights the trade-offs between fragmentation and consolidation. In fragmented markets, liquidity providers face higher search and adverse-selection costs because information is dispersed across venues and order flow is segmented (Chowdhry and Nanda, 1991; Madhavan, 1995). Consolidation centralizes information and reduces dispersion in quotes, improving price discovery and execution certainty, but can compress spreads and reduce market-making rents. The equilibrium depends on trading mechanisms, latency, and tick-size regimes. These insights provide the conceptual foundation for evaluating whether a consolidated price benchmark—such as the US-NBBO or our EBBO—enhances market efficiency by aligning competition while preserving incentives to supply liquidity. Our paper sheds light on the predictions of classic models of market integration and venue competition (Gehrig, 1993; Pagano and Röell, 1996; Parlour and Seppi, 2003), the economics of speed and technology choice by trading venues (Budish et al., 2015; Menkveld and Zoican, 2017; Pagnotta and Philippon, 2018), and fragmentation across lit and dark markets (Gresse, 2017). Foundational microstructure frameworks on information and limit-order dynamics provide primitives for these results (Glosten and Milgrom, 1985; Kyle, 1985; Glosten, 1994; Roşu, 2009), while coordination frictions and endogenous liquidity provision amplify the consolidation–fragmentation trade-off (Admati and Pfleiderer, 1988).

A large empirical literature shows that greater quote transparency and the consolidation of fragmented prices reduce trading costs and improve execution quality (Boehmer et al., 2005; Hendershott and Jones, 2005; O'Hara and Ye, 2011). Consolidated quoting aligns competition across venues and anchors pre-trade transparency, supporting deeper displayed liquidity and more efficient price discovery (Hasbrouck, 1995). Accurate measurement is non-trivial at high frequency: results are sensitive to how consolidated price

[1] Although the latency-neutral EBBO and the latency-adjusted REBBOs differ in construction, we use the term *EBBO* generically to denote the consolidated price tape.

tapes are constructed and timed (Hasbrouck, 2009; Holden and Jacobsen, 2014). Quote updates convey significant information, while reporting delays between the consolidated tape and venue's direct feeds generate arbitrageable price differentials (Brogaard et al., 2014; Budish et al., 2015; Easley et al., 2016; Foucault et al., 2017). These findings motivate our construction of the EBBO, which allows us to assess execution efficiency in absence of a binding consolidated tape and explicit order-protection rules.

We document European markets are surprisingly liquid with consolidation materially improving displayed direct-feed quotes. Relative to individual venues, the EBBO halves average quoted spreads—from 14.8 Euro-Cent (8.9 basis points) to 6.4 Euro-Cent (3.6 basis points)—and nearly triples displayed depth from 809 to 2,238 shares. Measured against midpoints, effective spreads are also tighter on the consolidated market (3.9 Euro-Cent, 3.4 basis points) when benchmarked against local midpoints (4.3 Euro-Cent, 3.9 basis points). In panel regressions, spreads are tighter with more venue competition and activity: more concurrently active venues and higher trading volume are associated with significantly narrower quoted and effective spreads. Higher market capitalization, a wider tick size, and lower price levels are associated with wider spreads. These results demonstrate that consolidation systematically reduces both quoted and realized trading costs, whereas Europe's tick regime and heterogeneous venue structure impose hard lower bounds on attainable spreads. For context, combining wholesalers and exchanges and taking odd-lots into consideration, S&P 500 stocks exhibit quoted and effective spreads of about 6.4 and 5.0 basis points, respectively. In our sample, EBBO effective spreads of about 3.4 basis points and quoted spreads of 3.6 basis points on average indicate trading costs below comparable US benchmarks (Dyhrberg et al., 2025).

Geographically, European equity trading is dispersed across three primary datacenter hubs—Slough/London, Frankfurt, and Bergamo—with additional venues in Madrid and Stockholm. Microwave latency corridors within the London–Frankfurt–Bergamo triangle constrain information propagation to roughly 2–4 milliseconds, while peripheral routes remain several times slower. By contrast, US venues are concentrated within a single co-located trading hub, reducing inter-datacenter propagation times to approximately 0.18 milliseconds. European quote coherence suffers accordingly: markets are locked 5% and crossed 2% of continuous trading time. Crossed states yield average arbitrageable gains of €10.8 per event, which cumulates to €4 million in annual latency-arbitrage potential across our sample. These structural transmission delays weaken real-time price formation and create predictable opportunities for speed-advantaged intermediaries.

While posted liquidity in Europe appears strong, realized execution quality is materially weaker. Only 79–86% of marketable orders execute at the EBBO, compared with over 96% in US markets (Dyhrberg et al., 2025). Roughly 18% of European trades experience price disimprovement, generating investor losses of cumulative €760 million over 2020–2024. One-third of these losses (€283 million) stem from inter-datacenter latency, and another third (€222 million) reflects mechanical price concessions due to limited available depth. Even after accounting for latency and liquidity constraints, residual welfare losses of €255 million—equivalent to €50 million per year, €198 thousand per trading day, and 0.52 basis points per disimproved trade—persist. These residual losses reflect routing frictions and incomplete venue connectivity: an implicit tax on slower or less-connected investors, with rents accruing to speed-advantaged intermediaries. Consistent with this interpretation, pan-European MTFs such as Cboe and Turquoise achieve higher best-execution rates (91%) than domestic exchanges like Xetra (84%), indicating static routing preferences even when better prices are available elsewhere.

Our results directly inform the ongoing regulatory debate on the design and implementation of a European consolidated tape under MiFID III.[2] By quantifying slippage, spreads, depth, and price (dis)improvements in nanosecond resolution, we document both the magnitude and the sources of inefficiency arising from Europe's fragmented market design. In the absence of a public pre-trade benchmark or transparency that fosters broker competition, these frictions disproportionately affect passive institutions and non-colocated participants, compounding over time into material costs for end investors and potentially increasing the cost of capital. The EBBO demonstrates that meaningful liquidity and execution gains are attainable through consolidation. Yet quote coherence and fairness depend critically on regulation, competition, technical infrastructure, and market access.

At a fundamental level, markets that are not explicitly integrated enable technologically advanced traders to earn rents by performing the de facto task of market integration. In the absence of a consolidated benchmark or synchronous price discovery, fast intermediaries arbitrage fleeting price differentials across venues, effectively supplying a coordination service that links fragmented markets in real time. While such activity can be viewed as compensation for immediacy provision and information aggregation (Menkveld, 2013; Budish et al., 2015; Pagnotta and Philippon, 2018), the persistence and magnitude of cross-venue dislocations, and lack of order-protection enforcing executions at the best available

[2] In June 2025, ESMA launched the formal selection procedure for the Consolidated Tape Provider (CTP) for equities and ETFs under the new technical standards adopted earlier that month. See ESMA Press Release, June 2025.

price, suggest that much of this income represents rent extraction rather than welfare-enhancing intermediation. Our conservative estimates imply that at least several hundred million Euro are lost annually to such inefficiencies in Europe. By comparison, the direct cost of explicit integration through order-protection and a consolidated quote is modest: the US-NBBO generates roughly \$400 million in annual revenues, with technology and operating expenses of only about 6% of cash flows (Budish et al., 2015; Aquilina et al., 2022). When regulatory and technological integration costs are this small relative to speed-based rents, the latter are best interpreted as economic rents rather than fair compensation for productive intermediation, consistent with models of speed competition and market design (Budish et al., 2015; Menkveld, 2017; Pagnotta and Philippon, 2018).

We identify three policy levers that can enhance execution quality. First, regulators could pursue latency harmonization by reducing the number of geographically dispersed datacenters or promoting colocation, thereby minimizing quote asynchrony and executions at stale prices. Second, they could require brokers to maintain connectivity to all venues that meet trading and quoting thresholds, broadening liquidity access and reducing cross-venue disparities in execution outcomes. Third, ESMA could introduce order-protection obligations that requires brokers to route orders to the venue with the best price. Whether these improvements in execution quality and trading activity outweigh the associated infrastructure, reporting, and regulatory costs remains an open empirical question and a key area for future policy evaluation.

Taken together, our evidence shows that European markets are deep, consolidation uncovers substantial cross-venue liquidity, but weak interconnections and heterogeneous venue connectivity prevent investors, especially less-informed and non-colocated participants, from realizing the market's full benefits. The results speak directly to recent MiFID III debates on consolidated tapes, latency harmonization, venue-access obligations, and order-protection mechanisms dependent on an EBBO benchmark for best-execution enforcement. Beyond the regulatory implications, the EBBO establishes a unified, high-frequency benchmark for analyzing execution quality, price discovery, and market design in fragmented European markets. By consolidating direct-feed quotes, it enables identification strategies that exploit variation in latency, venue connectivity, and order-routing behavior to isolate causal effects on execution outcomes. The EBBO transforms fragmented feeds into a single consolidated tape, allowing researchers to quantify how nanosecond-level frictions aggregate into macro-level measures of market efficiency.

1 Data & (R)EBBO reconstruction

1.1 Sample

Data are provided by BMLL Technologies. The database contains high-frequency trades and quotes obtained via direct feeds from European venues, timestamped to the nanosecond. We obtain data for all Eurostoxx50 constituents for our main analysis.[3] We restrict the sample to venue-ISIN combinations trading in Euro and extract data during continuous trading for all European venues which actively quote prices for the constituents between 2020 up to and including 2024. The final sample contains data of three MTFs—Aquis, Cboe, and Turquoise—and seven RMs—Amsterdam, Brussels, Helsinki, Madrid, Milan, Paris, and Xetra.

1.2 EBBO reconstruction

For each equity-day, we reconstruct a synthetic EBBO by aggregating local prices across all actively quoting venues. We start with identifying all changes to the local Best Bid and Best Offer (BBO) per venue. A BBO change can occur for the local Best Bid, Best Offer or the Best Bid and Best Offer simultaneously. To alleviate computational cost, we explicitly discard local BBO changes which solely involve changes in BBO size without changing the BBO price. We additionally discard one-sided BBOs and those which report a BBO price without a BBO size.

For our analysis, the accuracy of timestamps is vital. Our timestamps reflect the nanosecond in which a message has successfully been processed and is now publicly visible on the venue's direct feed. Our sample includes exchanges operating in time zones ranging from UTC+0 (e.g., London) to UTC+1 (e.g., Stockholm). All local venue timestamps are converted to a unified POSIX format, which allows for an instantaneous comparison of prices across geographically dispersed markets. In a hypothetical, latency-neutral market, we consolidate all local BBOs as if all quoting venues were situated in the same datacenter. In the sense of Budish et al. (2024), we then derive the EBBO for each timestamp by determining the highest available bid and the lowest available offer across all venues. Formally, for each equity i, venue j, date t, and nanosecond k, we define the synthetic European Best Bid (EBB) and European Best Offer (EBO) as follows:

[3] Constituents are as of 24.07.2025.

$$EBB_{itk} = \max_j BB_{ijtk}, \tag{1}$$

$$EBO_{itk} = \min_j BO_{ijtk}. \tag{2}$$

The resulting EBBO consolidates all local BBOs across venues in a latency-neutral fashion, thereby providing a hypothetical benchmark of best available prices in Europe in absence of infrastructural frictions. This synthetic benchmark will later be used to infer price slippage which arises from fragmented trading on a geographically dispersed market.

1.3 REBBO reconstruction

In practice, European equities trade on multiple venues which are geographically distant, see Figure 1. Such geographical dispersion necessitates to account for transmission latencies between venues when reconstructing a more realistic benchmark of a consolidated price tape. Transmission latency measures the propagation delay of a signal traveling between two physically distant exchanges. The exact transmission latency between two exchanges depends on the exchanges' kilometer distance and the signal's travel speed. In the exact same nanosecond, the EBBO observed on Xetra might therefore differ (slightly) from the EBBO observed on Paris. In analogy to Schwenk-Nebbe and Thimsen (2024), we thus reconstruct latency-adjusted REBBOs for each venue, which represent the EBBO as observed through the lens of the respective venue after accounting for physical transmission delays between datacenters.

The adjustment of local BBO prices for transmission latencies requires a venue-pair specific lagging of quote timestamps. We obtain the venue-pair transmission latency $\varphi(e_i, e_j)$ between exchange i and exchange j by calculating the great circle haversine-kilometer distance between the two exchanges and multiplying the distance with the signal transmission speed in optic fiber cable (204,190 kilometers per second).[4] For example, the haversine-kilometer distance between the Frankfurt and Bergamo datacenter measures 497 kilometers (309 miles). In optic fiber cable, a signal needs at least 2.44 ($497 \text{ km} \times \frac{1}{204{,}190 \text{ km/s}} \times 1,000$) milliseconds to travel from Frankfurt to Bergamo. The best prices quoted on Frankfurt are therefore observable in Bergamo only 2.44 milliseconds after they become available in the direct feed of Frankfurt. For the computa-

[4]Our results remain robust when lagging the timestamps with empirically observed venue-pair specific transmission latencies. See Table 12.

tion of the local REBBO in Bergamo, the Frankfurt timestamps thus need to be lagged by 2.44 milliseconds. Equivalently, London's prices arrive in Bergamo with a lag of 4.67 ($954 \text{ km} \times \frac{1}{204{,}190 \text{ km/s}} \times 1,000$) milliseconds after becoming available in London's direct feed. Table 3 details theoretical and empirically observable venue-pair dissemination times. Similar to Schwenk-Nebbe and Thimsen (2024), we add a constant transmission overhead of 30 percent to the transmission latency in order to account for subsequent information dissemination at the consolidating exchange.

We repeat the venue-pair specific lagging of local best prices for all venues actively quoting an instrument[5] before consolidating the latency-adjusted BBOs and deriving best REBBO prices. For an equity that trades on four different venues, we thus reconstruct four individual REBBOs, each representing a venue-specific consolidated version of the EBBO. Additionally, we snapshot conventional order book metrics on the local and global level for every REBBO update. We refer to individual trading venues as *local* markets and to the consolidated EBBO as *global* European market. We compute local and global depths, spreads and venue-level liquidity contributions.[6] Figure A.1 and Appendix A.3 conceptualize the (R)EBBO consolidation mechanism and latency-correction. After accounting for transmission latencies and information propagation, the REBBOs document best European prices observable to market participants directly colocated at each respective venue, thereby approximating the structure, function, and time-lag of a European consolidated tape.

1.4 Matching trades to (R)EBBOs

To assess overall execution quality we match all lit trades during continuous trading to prevailing REBBO prices of the trade's execution venue. E.g., all trades executed on Xetra are matched against Xetra's prevailing REBBO, all trades executed on Turquoise are matched against Turquoise's prevailing REBBO. Whenever the direction of a trade is not readily available in the data, we rely on Lee and Ready (1991) to flag the corresponding trade side based on the execution venue's midpoint. We also match trades to the latency-neutral reconstructed EBBO. The final sample spans 1,284 trading days, contains 1.17 billion trades and 1.31 billion EBBO updates.

[5] We discard trades and quotes originating on venues CHIX, TRQX and AQXE during the post-Brexit period starting in 2021. Trading and quoting activity of those venues has moved with admission of regulatory Brexit changes to the respective EU-based counterparts CEUX, TQEX and AQEU.

[6] It is worth noting that unlike American data which has previously been used for NBBO reconstruction, our data reports liquidity precisely per share, thereby avoiding NBBO liquidity distortions due to odd-lot liquidity reporting thresholds.

1.5 Assessing execution quality

We then assess the execution quality per trade by comparing the trade's execution price with prevailing (R)EBBO prices at the time of execution. Trades on the bid side of a local market are compared to the (R)EBB, while trades on the ask side of a local market are compared to the (R)EBO. The difference between the trade's price and best globally available prices is quantified as *slippage cost* which is expressed in Euro-Cents or an equity's ticks. To avoid noisy estimation of slippage cost, we classify a trade as executed at best prices whenever the trade's price deviates less than a $\frac{1}{100}$th Euro-Cent (or $\frac{1}{10{,}000}$th Euro, which is equivalent to the smallest possible tick size under MiFID II) from global best prices.

We further categorize trades according to their likelihood of walking the book. Part of the observed price slippage may arise mechanically from a trade's execution size rather than informational inefficiencies. Large trades are more likely to consume best available liquidity and thereby walk the book, generating apparent price slippage. To isolate this effect, we construct a size–liquidity flag that evaluates each trade's execution size with the contemporaneous local or global depth. We hence classify trades into executions against best quotes with (i) *local sufficient depth* and (ii) *global sufficient depth*. A trade is labeled as executed against local sufficient depth whenever the trade's execution size is smaller than or equal to the liquidity available at the best price of the respective execution venue. Analogously, a trade is executed against global sufficient depth if the trade's execution size does not exceed the maximum liquidity contributed by any single venue to the prevailing global REBBO quote. This way, we can distinguish between trades which face purely mechanical slippage due to their execution size and those which potentially could have been executed at best prices without causing mechanical price impact.

2 The European capital market

MiFID I, introduced in 2007, marked the first major liberalization of European trading by ending national exchange monopolies and fostering competition among RMs, MTFs, and SIs. The directive aimed to promote efficiency and innovation but also created a fragmented market structure in which liquidity and information became dispersed across venues. MiFID II and the Markets in Financial Instruments Regulation (MiFIR), implemented in 2018, expanded pre- and post-trade transparency and harmonized investor protection through uniform reporting obligations directly applicable across member states.

Yet the fundamental architecture—competition without consolidation—remained unchanged. Unlike the US, where Reg NMS mandates a single consolidated quote, the NBBO, Europe continues to operate without a unified reference price, leaving brokers and investors to navigate dispersed markets without a binding price benchmark for execution or oversight.

Our results emerge from this institutional setting. They show how Europe's fragmented market design generates latency asymmetries, uneven access to best prices, and measurable welfare losses. These findings speak directly to current regulatory initiatives under MiFID III, the Capital Markets Union (CMU), and the Retail Investment Strategy (RIS), all of which aim to deepen market integration, improve investor outcomes, and establish an European consolidated tape as the cornerstone of pre-trade transparency and best-execution enforcement. The evidence presented here quantifies how far current fragmentation departs from the ideal of a unified market and provides a data-driven foundation for evaluating the efficiency gains and costs of the proposed EBBO framework.

The remainder of this section describes the structure and functioning of the European capital market. We begin with a descriptive analysis of order book mechanics (Section 2.1) and market infrastructure (Section 2.2). We then examine the quality of displayed quotes (Section 2.3) and the factors that determine it (Section 2.4), followed by an analysis of execution quality (Section 2.5) and its determinants (Section 2.6). Together, this analysis provides a coherent picture of how liquidity, infrastructure, and trading behavior interact within Europe's multi-venue system and how regulatory design influences the efficiency of capital allocation.

2.1 Order book statistics

To infer about global order book statistics, we rely on the REBBO virtually consolidated in Bergamo to homogenize latencies.[7] Table 1 summarizes order book descriptives for the sampled equities. Table A.3 provides the full instrument-level panel, including each instrument's country, ISIN, the number of active market centers, average price, median tick size and local spread, the number of EBBO updates, cumulative trade counts, and market capitalization. We keep this level of detail in the appendix for completeness and replication while the main text reports aggregate statistics.

[«Insert Table 1»]

[7] For equities that do not actively trade in Bergamo, we calibrate a synthetic REBBO consolidated in Bergamo. This avoids cross-venue variation in market quality arising from heterogeneous transmission latencies.

Panel A: Equity characteristics. Our sample consists of highly liquid equities with market capitalizations between €9.4 billion and €319.0 billion and average daily trading volumes ranging from €36.9 million to €350.4 million. The average stock trades at €166 with an average tick size of roughly 4.0 Euro-Cent and a median of 1.0 Euro-Cent. Under the MiFID II tick-size regime, European equities are classified into six liquidity bands based on price level and trading activity, with minimum price increments ranging from €0.0001 (0.01 Euro-Cent) to €500 (50,000 Euro-Cent).[8] A defining structural difference between Europe and the US is that all lit trading in the EU must occur on tick—prices cannot improve within the minimum increment. US markets, by contrast, permit sub-penny executions and price improvement through internalizers, wholesalers, and dark pools operating within the lit grid. The European framework, though harmonized across venues under MiFID II, enforces discrete pricing even where competitive pressure would otherwise narrow spreads further for certain orders.

Panel B: Liquidity. The average local quoted spread is 14.8 Euro-Cent (4.18 ticks) in the mean and 4.5 Euro-Cent (2.95 ticks) in the median. By contrast, the global EBBO quoted spreads are substantially tighter: means of 6.4 Euro-Cent (1.83 ticks) and medians of 2.3 Euro-Cent (1.55 ticks), i.e., roughly 40–50% narrower than local quotes on average. Consolidated markets therefore display better quoted prices than any single venue; market orders routed to venues setting the global inside face smaller immediate price concessions than orders executed against wider local quotes. The post-consolidation compression indicates that fragmentation inflates displayed trading costs by dispersing liquidity across venues. Effective spreads show the same pattern. On local markets, the average effective spread accounts to 4.3 Euro-Cent (3.86 basis points), whereas the global EBBO effective spread measures 3.9 Euro-Cent (3.35 basis points). This difference reflects measurement as well as economics: the EBBO midpoint is tighter and less stale than venue-specific midpoints that are constrained by ticks and local order flow. Hence trades benchmarked to the EBBO midpoint appear cheaper not only because routing can access better prices, but also because the reference price is more informative.

Liquidity is likewise deeper at the consolidated inside: average depth at the EBBO is 2,238 shares versus 809 shares locally, an almost threefold increase. Aggregating depth across venues narrows effective spreads and increases executable size, lowering expected execution shortfalls; conditional on cross-venue access and timely routing. Two caveats follow. First, consolidated depth aggregates distinct queues; sweep costs and queue prior-

[8] See Commission Delegated Regulation (EU) 2017/588 (RTS 11), Official Journal of the European Union, Annex I. For the complete tick-size schedule, see Table A.8.

ity still matter for realized outcomes. Second, narrower global measures do not guarantee identical gains under a single-venue design; the counterfactual without inter-venue competition is an empirical question outside this paper.

Panel C: Competition. The fact that global quoted spreads are narrower than local quoted spreads indicates that not all venues post globally competitive prices. Across our sample, the average number of quoting venues per stock ranges from 4.4 to 7.6, yet fewer than half contribute to the global inside quote at any given time. If all venues continuously posted at the EBBO, the difference between local and global spreads would vanish. In practice, only a subset of venues consistently competes for order flow, while others remain systematically away from the market—either due to slower quote updates, higher latency, less sophisticated routing connections, business models emphasizing internalization over price competition or regulatory enforced tick-size constraints on local markets. This uneven contribution to the global inside highlights that liquidity competition in Europe is selective rather than universal. Figure A.1 visualizes how the global market spread can undercut local market spreads without violating locally enforced tick-size constraints.

Overall, these descriptive statistics show that fragmentation depresses observable liquidity and widens spreads when viewed at the single-venue level. Consolidation across venues, by contrast, reveals substantially deeper liquidity and narrower spreads, illustrating that market quality improves when liquidity is aggregated. This finding aligns with evidence that fragmentation can enhance aggregate liquidity *conditional on* multi-venue access, but that market quality deteriorates once order flow becomes siloed on individual platforms (Dzieliński et al., 2024). Limited quoting participation implies that investors relying on a single direct feed or restricted venue connectivity face systematically inferior prices and higher execution risk. In effect, fragmentation creates a two-tier market: well-connected participants with access to multiple feeds capture most of the available liquidity, while others transact at less favorable terms. Section 2.5 quantifies these disparities as effective price slippage and adverse selection costs, demonstrating the economic magnitude of incomplete venue integration.

2.2 Market infrastructure

The European capital market clusters into three datacenter hubs: (i) London/Slough, United Kingdom (Aquis (AQEU), Cboe (CEUX), and Turquoise (TQEX)), (ii) Frankfurt, Germany (Xetra (XETR)), and (iii) Bergamo, Italy (Amsterdam (XAMS), Brussels (XBRU), Milan (XMIL), Paris (XPAR)). These datacenters collectively account for 96%

of our sample trading volume, by executing 44%, 16%, and 36% of trades, respectively. The most active trading venue is Cboe while the most active quoting venue is Xetra. RMs, such as Xetra or Paris, are the primary listing venues for domestic equities and subject to stricter MiFID II obligations on transparency and market surveillance. They typically centralize domestic order flow and maintain deep books with designated market makers in their home listings, reflecting their role as reference markets for price discovery. Official opening and closing auctions are held on RMs. In contrast, MTFs like Aquis or Cboe operate as secondary trading platforms that compete for order flow. As in the US, European venues differ in their fee structures, including maker rebates and taker fees that shape effective execution costs. These cum-fee price differences can influence routing decisions across venues (Battalio et al., 2016; Li et al., 2023). While we compile representative fee statistics, comprehensive data on discounts and tiered volume incentives are unavailable. MTFs generally exhibit the lowest access costs and are the only venues in our sample that offer liquidity-provider rebates. These differences produce a natural complementarity: RMs provide baseline depth and liquidity, while MTFs inject price competition that often improves the inside spread. Together, both venue types interact to shape the EBBO by balancing depth, speed, and pricing efficiency. Table 2 reports precise venue locations, fee schedules, the number of EBBO updates, and cumulative trade counts.

SI and over-the-counter (OTC) trades also play a complementary but derivative role in this ecosystem. Unlike RMs and MTFs, which contribute directly to price formation through continuous limit order books, SIs and OTC platforms typically reference the prices established on lit venues when quoting or executing trades. Under MiFID II, both are required to provide quotes at or within the BBO available on public markets, effectively anchoring their internal prices to those discovered on RMs and MTFs. As a result, SIs and OTC transactions borrow their price integrity from the lit market while internalizing order flow that would otherwise contribute to displayed liquidity. This structural dependence reinforces the central role of RMs and MTFs in determining European equity prices: while off-exchange venues expand execution channels and cater to specific order types (e.g., large-in-scale or riskless principal trades), they remain price-takers with respect to the EBBO. Consequently, the quality and coherence of lit-market prices directly shape the efficiency of the trading in European equities.

[«Insert Table 2»]

The datacenters of our sample venues are distributed across five countries, with an average geodistance of 611 kilometers (379 miles) between actively quoting venues for the average stock. Figure 1 illustrates the spatial dispersion of European trading infrastructure. Such geographic dispersion contrasts with the US market, where the NBBO infrastructure is co-located within a 20 kilometer (12 mile) radius in northern New Jersey limiting empirical transmission latencies to roughly 0.18 milliseconds between venues. In Europe, by comparison, inter-datacenter distances routinely exceed 1,600 kilometers, with the longest link—between Bolsa de Madrid in Spain and Nasdaq's Baltic datacenter in Sweden—spanning 2,600 kilometers. These physical separations translate into propagation delays that are an order of magnitude higher than in the US; the average sample stock experiences an inter-datacenter transmission latency of 2.99 milliseconds across its actively quoting venues. Such geographic fragmentation introduces persistent latency differences and jitter that directly affect price synchronization, routing decisions, and the probability of best execution.

[«Insert Figure 1»]

Since datacenter operators and network providers rarely disclose details on precise network topology and infrastructural design, we estimate European inter-datacenter transmission latencies using two complementary and relatively straight-forward approaches. First, we compute theoretical inter-datacenter transmission latencies based on the haversine-distance between datacenter pairs and the signal travel speed in vacuum and optic fiber cable. While these measures provide a first approximation of theoretical transmission limits and act as a natural benchmark to validate empirical transmission times, they do not reveal the presence of microwave linkages or potential low-latency corridors. Thus, and second, we exploit our nanosecond-timestamped data to derive empirical estimates of inter-datacenter transmission latencies. Specifically, we calculate the venue-pair *quote-reactiveness*. Quote-reactiveness measures the median time between a trade on one venue and the next subsequent quote revision on the other venue. This empirical estimate reflects the effective speed at which information propagates across datacenters, and thereby accounts for path inflation, repeaters, and routing through intermediate microwave hubs. The comparison between theoretical and empirically observable transmission latencies allows to identify fiber and microwave connected datacenters. Table 3 reports theoretical and empirical transmission latencies in milliseconds.

[«Insert Table 3»]

The fastest and most frequently used inter-datacenter routes are typically supported by microwave links, which outpace theoretical transmission latencies in fiber cable. Microwaved routes connect the London–Frankfurt–Bergamo triangle. Those microwave-supported latency corridors reduce propagation time to 2–4 milliseconds, whereas peripheral routes remain several times slower. By contrast, venues at the geographic periphery, such as Madrid and Helsinki, exhibit disproportionately longer transmission times. Although the distance between Madrid and Frankfurt measures 1,454 kilometers with a theoretical transmission time of 7.12 milliseconds in optic fiber cable, empirical transmission latency accounts to 11.59 milliseconds, almost 63% higher than the theoretical fiber-based benchmark.

The resulting latency differentials across venues illustrate how Europe's fragmented geography translates directly into measurable market inefficiencies. Geographic dispersion generates structural, multi-millisecond delays between trading hubs, creating two fundamental challenges. First, market participants must reconstruct real-time prices from heterogeneous and delayed data feeds, each reflecting a different view of the market. Second, every venue maintains its own local BBO, producing systematic asynchrony across order books. For brokers and routers seeking a consolidated view, these transmission gaps of several milliseconds make it impossible to observe a perfectly synchronized market state, increasing exposure to stale quotes and execution errors. In the absence of binding order-protection and best-execution rules, these structural latencies manifest as frequent and economically meaningful price discrepancies across venues. Such differences are consistent with theoretical models linking lower exchange latency to heightened adverse-selection risk for liquidity providers (Menkveld and Zoican, 2017). Traders operating on incomplete or delayed representations of the EBBO therefore face a systematic risk of transacting at inferior prices. The next sections quantify the economic magnitude of these latency-induced dislocations and their welfare implications for European market participants.

2.3 Quote quality

Latency asymmetries show that fragmentation is not only a technological issue but also a fundamental source of inefficiency. When quotes update at different speeds across venues, the integrity of displayed prices deteriorates because what appears to be the EBB or EBO in one venue may already be stale in another. As a result, latency differences directly affect quote quality, defined by how accurately and consistently displayed prices represent executable trading opportunities. High-quality quotes are tight, deep, and synchronized, allowing market participants to identify the best available price and trade with minimal slippage. In fragmented markets, however, quote quality depends on more than local liquidity, it also hinges on how effectively information is shared and competition across venues. Differences in access speed, routing tables, and quoting incentives can distort the global inside market, leading to stale, locked, or crossed quotes that obscure price discovery. Weak quote quality undermines both liquidity provision and execution quality: liquidity takers face higher uncertainty about attainable prices, while liquidity providers risk being adversely selected. As a result, quote quality serves as a core indicator of market efficiency and the degree to which dispersed venues operate as a single, integrated market.

Key aspect of Europe's fragmented markets is the behavior of quotes across venues. On a single market, the bid can never match or exceed the ask, since such orders would immediately execute. In a fragmented system, however, identical securities may display overlapping bids and offers on different venues without triggering execution, creating temporary states of disequilibrium. Such conditions blur the true best available price: liquidity takers face uncertainty about where to execute, while liquidity providers risk being adversely selected as stale quotes are lifted by faster traders. As a result, quote incoherence elevates trading costs and weakens both price discovery and confidence in displayed liquidity. These effects are unlikely to be evenly distributed across market participants. Sophisticated and colocated traders can monitor and exploit fleeting price discrepancies, whereas retail investors relying on slower consolidated feeds or broker routing systems are most exposed to execution at stale prices.

Table 4 summarizes the frequency and duration of market states. We distinguish between normal, locked and, crossed markets. During locked (crossed) markets, global spreads become zero (negative) and violate ordered market conditions while local venue dynamics remain intact and continue to quote positive spreads. Panel A shows that an average stock experiences more than 17 thousand price updates under normal market conditions during a continuous trading day alongside nearly 1 thousand and 2 thousand price updates which result in locked and crossed markets respectively. This translates to 91% of continuous trading under normal, ordered market conditions, with locked states accounting for roughly 5% and crossed states for 2% of continuous trading time. Finally, crossed market states generate arbitrage opportunities for traders equipped with technological infrastructure and an informational advantage to realize a seemingly riskless profit. Panel C reports that an average global market cross generates an arbitrageable profit of €10.79. In median, the average daily arbitrageable profit per stock amounts to €305 which cumulates to €4 million of yearly arbitrage potential across our sample.

[«Insert Table 4»]

Figure 2 illustrates the evolution of crossed markets over time. On average, the frequency of crossed states remains relatively stable, but we observe clear spikes during episodes of market stress. During periods of market stress, the cross-sectional dispersion also widens, with some stocks experiencing a much higher incidence of crossed quotes. This complements by showing that quote incoherence is not only structural but also cyclical, intensifying precisely when liquidity and transparency are most valuable. Most market crossings are transitory and persist only for a few microseconds.

[«Insert Figure 2»]

The majority of market crossings arises from flickering quote insertions generated by algorithmic insert-and-delete strategies. Such strategies build on low latency infrastructure and are calibrated on local markets to minimize reaction times. The local setup of algorithmic quoting strategies thus does not consider the potential generation of global mispricings and market inefficiencies. Algorithmic quote flickers usually contain elements of randomization, e.g. variations in order size or inserted price level, to conceal strategic and widely predictable algorithmic activity. While the average quote flicker does not generate a global market crossing, such randomness can cause flickering quotes to occasionally cross the global market, especially if local spreads are already tight.

Although not central to the paper but in line with our previous observations, we document that crossed markets are arbitraged based on local market reactions rather than global ones. In theory, an aggressive quote insertion which crosses the global market creates an arbitrageable mispricing visible to any sufficiently fast market participant who consolidates individual data feeds. In practice, we observe the exploitation of such crossing quotes based on local market reactions rather than a precise evaluation of the global, consolidated market. Similarly, Cartea et al. (2025) provide a more detailed perspective on the precise quoting and arbitrage mechanisms on a local level.

Analogously, Figure 3 plots the evolution of quoted and effective spreads over time. As elaborated in Section 2.1, global spreads are consistently narrower and less volatile than local spreads, confirming that cross-venue aggregation reveals deeper, more competitive liquidity. Both quoted and effective spreads widen sharply during episodes of market stress before reverting as conditions stabilize. Overall, the figure illustrates that consolidation smooths liquidity fluctuations, while local market fragmentation amplifies them, particularly during crises when transparency and coordination are most valuable.

[«Insert Figure 3»]

2.4 Determinants of quoted liquidity

Evidence on crossed markets shows that fragmentation and asynchronous quoting shape the structure of observable liquidity. Building on this link between market states and price formation, we next examine which market characteristics account for cross-sectional differences in quoted and effective spreads across European equities. Table 5 presents regression results.

[«Insert Table 5»]

Quoted spreads capture the prices liquidity providers post ex-ante, while effective spreads reflect the costs investors actually incur ex-post the trade. Both measures narrow significantly as the number of active venues increases, confirming that competition among venues compresses trading costs. Higher trading volume is also associated with tighter spreads, consistent with liquidity concentrating where activity is greatest. In contrast, larger firms exhibit wider spreads after controlling for volume, reflecting the mechanical effect of higher price levels and broader tick sizes. The tick-size coefficient is strongly positive, showing that regulatory grid constraints directly inflate spreads, while the negative price coefficient indicates that higher-priced stocks trade on finer relative increments. Overall, the regressions highlight that quoted liquidity improves with venue competition and trading activity but remains bounded by tick-size regulation and the heterogeneous market structure that defines Europe's fragmented equity landscape. Taken together, the joint behavior of quoted and effective spreads indicates that while competition enhances visible liquidity, underlying market design still limits how efficiently order flow is translated into execution quality.

2.5 Execution quality

The structural transmission delays have direct implications not only for quoted and effective spreads but also for overall execution quality and regulatory compliance. Under MiFID II, brokers and trading venues are required to take "all reasonable steps" to ensure best execution across markets. While the regulation provides only limited guidance on how best execution should be assessed in practice, we define best execution within the scope of this study as best *price* execution – that is, trades executed at global best prices. When quotes update asynchronously across datacenters, traders relying on delayed or incomplete feeds may systematically execute at stale prices, thereby incurring measurable slippage relative to the true global best quote. These frictions motivate our subsequent analysis of fragmentation's causal impact on market quality, quantified through each trade's effective price slippage.

The average sample trade executes at the EBBO, with roughly one in eight trades occurring at prices worse than the globally best quote. Table 6 decomposes execution quality across four progressively stringent filters that isolate distinct sources of slippage. Panel A shows how the share of trades executed at best, improved, or worsened prices evolves as latency and mechanical price-impact effects are sequentially controlled for. Panel B reports the fraction of shares executed at EBBO prices per venue, highlighting systematic cross-market differences in connectivity and routing efficiency.

[«Insert Table 6»]

To disentangle distinct sources of price slippage, we apply four increasingly stringent filters to our initial sample of trades. The baseline scenario (Filter i) contains the full set of trades which are matched to the latency-neutral version of the EBBO. This scenario establishes the hypothetical case of all venues being directly colocated in a single datacenter and thereby acts as a theoretical benchmark. The second scenario (Filter ii) matches the full set of trades to the venue-specific latency-adjusted REBBOs and thereby accounts for the geographic dispersion of markets. The third filter (Filter iii) subsets all trades to those small enough to execute on their respective local venue without causing mechanical price impact, i.e. the local liquidity currently available at the inside market suffices to execute the trade. A precise categorization of such trades allows us to account for trades which look like they have been executed at worse than globally available prices when in reality the trade's execution size caused a mechanical price impact by walking levels of the book. The same logic applies to the final filter. The final filter (Filter iv) subsets all trades to those small enough to execute on any venue contributing to the REBBO without causing mechanical price impact. While the first two scenarios (Filter i and Filter ii) compare execution quality for the latency-neutral and the latency-adjusted case, the two consecutively applied filters (Filter iii and iv) progressively narrow the set of trades to isolate residual slippage. Execution quality after application of the final filter (Final v) reflects *effective slippage* after accounting for mechanical and latency-induced effects.

Table 7 reports adverse selection cost measured by execution quality, Panel A, cumulative slippage in Euro, Panel B, and per-share slippage in basis points, Panel C, through successively applied filters. Figure 5 visualizes the results.

[«Insert Table 7 and Figure 5»]

In the aggregate, these execution frictions quantify the allocative inefficiency that results from Europe's fragmented market architecture. In absence of dissemination latencies (Filter i) roughly 81% of marketable orders execute at prevailing best prices, while the remainder incur cumulative slippage cost of €477 million. For trades executed at inferior prices, this corresponds to an average slippage of 0.56 basis points per share. Incorporating venue-specific transmission latencies (Filter ii) reduces the share of trades executed at best prices to about 79% and increases cumulative slippage to €760 million. This deterioration in execution quality indicates that a sizable fraction of observed slippage arises from information delays across geographically dispersed datacenters. Accounting for inter-venue transmission latencies reveals that many trades occur after superior quotes have expired—quantifying the hidden cost of geographic dispersion through systematic execution against stale prices. Further restricting the sample to trades small enough to execute on their local venue without exhausting displayed depth (Filter iii) removes price concessions mechanically driven by limited available liquidity. This adjustment raises the share of best-executed trades to 82% and reduces aggregate slippage to €594 million. Restricting the sample to trades small enough to execute globally without price impact (Filter iv) raises the best-execution ratio to 86%. A cumulative residual of effective €255 million price slippage persists. The average price-disimproved trade faces an effective slippage of 0.52 basis points per share.

The step-wise decomposition of execution frictions shows that one third of measured slippage reflects structural latency asymmetries, another third stems from mechanical price impact due to insufficient market depth. A sizable fraction of observed slippage can thus be rationalized by market mechanics while a large and economically meaningful residual worth €255 million persists, providing a direct welfare benchmark comparable to the magnitude of standard liquidity-taking fees. This residual slippage is unexplained by trade mechanics and dissemination latencies, thereby representing routing inefficiencies and incomplete venue connectivity. Traders executing aggressively at displayed quotes incur a persistent cost premium raising the cost of immediacy beyond explicit taking fees.

The overall improvement in execution quality observed with increasingly stringent filters also persists within individual venues. Cumulatively, Cboe and Turquoise provide best execution, while trades executing on Madrid or Xetra miss best prices more often. The heterogeneity across venues likely reflects differences in broker connectivity and access architecture. Pan-European MTFs such as Cboe and Turquoise are integrated into the smart-order routers of large global brokers and colocated in major datacenter hubs, enabling order flow to react quickly to changes in the EBBO. In contrast, domestic exchanges, such as Madrid and Xetra, depend more on national broker networks with limited cross-venue access, making their trades more likely to execute at stale prices even when better quotes exist elsewhere. Xetra, for example, executes the largest share of trades at best for German primary-listed stocks, reflecting its strong domestic connectivity but relatively limited reach across other European venues. We document similar patterns for non-German stocks, e.g. for French primary listed stocks, Paris shows the greatest fractions of trades executed at best prices. With price discovery fragmented across venues, best available prices remain anchored to the primary listing exchanges. The heterogeneity in execution quality across venues confirms that connectivity and routing are central determinants of investor outcomes.

Our results indicate that brokers employing dynamic smart-order-routing systems connected to a larger number of venues can achieve superior execution quality. By contrast, brokers relying on static routing tables and limited venue connectivity forgo better execution opportunities for their clients. While we cannot observe the fixed and operational costs of broader venue connectivity, and it is possible that the incremental execution gains do not fully offset these costs, evidence suggests that a common European market infrastructure, akin to Reg NMS and a consolidated tape, could improve welfare by spreading connectivity costs proportionally and eliminating redundant technological investments.

Table 8 provides a detailed breakdown of price deviations and slippage costs arising from inefficient routing and residual market frictions (Filter v). Panel A groups observable price deviations by absolute Euro-Cent magnitude, while Panel B reports relative deviations in ticks per equity.[9] The majority of trades (86%) execute at the global best, with the remaining price deviations distributed asymmetrically around the reception of best prices. The distribution is heavily skewed toward price disimprovements, as worsened executions vastly outnumber and outweigh price improvements in both frequency and cost. Such observation indicates we are not merely observing statistical timestamp noise but rather systematic market frictions.

[«Insert Table 8»]

The systematic imbalance between price improvements and disimprovements translates directly into measurable slippage costs, which escalate with the size of the price deviation. Large dislocations are especially costly: trades which occur more than €0.10 away from the EBBO, generate €95 million in losses. Additional bands of up to €0.10, €0.05, and

[9] Tables A.5, A.6, and A.7 present analogous decompositions for the three alternative filtering steps.

€0.01 account for €30 million, €60 million, and €26 million, respectively. Even deviations as small as €0.005 affect 11% of trades, costing €44 million. Price improvements are vanishingly rare: fewer than 1% of trades beat the EBBO, with total positive slippage around €10 million. Tick-based results are similarly asymmetric: 8% of trades miss best prices by one tick, while 4% of trades miss best prices by up to five ticks. One-tick deviations reflect minor price flickers due to stale quote evaluation or timestamp misalignment, whereas deviations of up to five ticks point to more systematic market frictions.

Figure 4 visualizes the evolution of execution quality over time. The figure plots the share of trades executed at the EBBO across all Eurostoxx50 constituents and highlights pronounced declines during periods of general market stress. Execution quality averages 85–90% during most of the sample, but drops sharply in early 2020, late 2020, and early 2022. During these episodes, execution at the EBBO falls for the lowest quartile of stocks, and cross-sectional dispersion widens markedly.

[«Insert Figure 4»]

The breakdown of execution quality proposes two main findings: first, slippage cost is non-random and, second, slippage cost is persistent. The concentration of slippage cost shows that traders systematically fail to access superior quotes that exist simultaneously on other venues. While transmission latencies and price differentials explain a considerable portion of price slippage, a significant residual slippage remains which is attributable to routing inefficiencies and incomplete venue connectivity. The persistence of executions at locally inferior prices indicates that routing to best global prices does not occur systematically. Order flow thus appears segmented by venue rather than integrated across the European trading landscape. This creates recurring arbitrage opportunities for faster participants and persistent adverse-selection costs for slower or less-connected traders. The absence of a low-latency, consolidated EBBO benchmark effectively taxes uninformed participants by denying them access to the true EBBO. Our findings thus identify a tangible source of allocative inefficiency: market participants pay more or receive less than the competitive price, not because of fundamental risk but because of infrastructural opacity. The frictions we document primarily burden retail and less-sophisticated investors, as technologically advanced traders can mitigate or arbitrage around these inefficiencies.

2.6 Determinants of execution quality

To examine the determinants of execution quality more directly, we estimate the probability that a trade executes at the prevailing EBBO using a linear–probability specification. Our dependent variable is a binary indicator equal to 1 if a trade executes at best and equal to zero if it does not. To limit computational complexity, we randomly sample 100 trades per equity and day. Since we are interested in investigating residual slippage unexplained by latency differentials or trade mechanics, we subsample our set of trades to those executed against globally sufficient liquid REBBOs (Filter iv). This leaves us with a representative sample of 6 million trades. Table 9 presents regression results.

[«Insert Table 9»]

Specification (1) regresses best execution on the geodistance between the executing venue and the venues currently posting best prices. As the sample of trades is already matched to latency-adjusted REBBOs, the explanatory power of geodistance to best execution is, as expected, marginal. Nevertheless, for every additional 100km of physical distance between the venues, best execution probability decreases by roughly 1 percentage point. Specifications (2) and (4) offer the most comprehensive analysis, incorporating controls for liquidity, venue characteristics, and market features. Results are consistent across both models. Geographic distance and transmission latency remain strongly negative, confirming that physical separation and slower information propagation materially reduce the likelihood of best execution. Execution quality improves with the number of venues quoting at the global inside: each additional venue posting the global best price increases the probability of best execution by roughly 2.8%. This effect highlights that simultaneous quoting across venues—a sign of market integration and competition—enhances execution quality beyond what can be achieved through proximity or low latency alone.

Both global and local spread coefficients are negative and statistically significant, indicating that wider spreads—whether on local or global markets—reduce the probability of best execution. This pattern is consistent with price granularity under Europe's tick-size regime: when spreads widen, price competition weakens and routing algorithms face fewer opportunities to match or improve upon the best available price. Larger trades further reduce the probability of best execution, reflecting a fundamental trade-off between immediacy and price improvement, as large orders tend to prioritize completion speed over marginal price gains.

During locked and crossed markets, the likelihood of best execution drops sharply by 8.0% and 31.6%, respectively. These coefficients quantify the cost of asynchronous order books, where transient quote discrepancies lead routing algorithms to execute against stale local prices. Such states reflect fleeting arbitrage opportunities rather than genuine price discovery, highlighting how latency-induced desynchronization undermines execution quality even in liquid markets. We also include the time elapsed since the last EBBO update. While older EBBO updates should, in principle, allow more time for price information to propagate across venues, the coefficient is statistically insignificant. This suggests that exchanges do not systematically condition routing decisions on contemporaneous best prices, implying that execution algorithms may be driven by static rules rather than real-time quote evaluation.

Venue effects are economically significant and consistent with institutional frictions. Trades on MTFs are roughly 7.7% more likely to achieve best execution, while trades on the stock's primary listing exchange gain a smaller but positive 1.7% effect. In contrast, executions on foreign regulated markets, for example, a German stock trading in Paris or Madrid, are significantly less likely to obtain best prices. These differences suggest that familiarity, local market infrastructure (including fees, designated market makers, and order flow concentration), and network connectivity jointly enhance execution quality. Overall, the results show that fragmentation, latency, and incomplete venue connectivity, rather than liquidity supply alone, govern the probability of achieving best execution in Europe's multi-venue environment.

While market distance explain a substantial share of observable execution patterns, latency alone cannot account for the full variation in best-execution outcomes. Consistent with the results in Section 2.5, the evidence indicates that some venues are structurally less connected to the consolidated market, leading to persistently lower execution quality—especially for equities not primarily listed on those venues. These connectivity gaps show that fragmentation imposes unequal costs across markets: trades executed on peripheral or weakly integrated platforms are systematically more likely to miss the EBBO. By quantifying these shortfalls relative to a unified best-price standard, the EBBO provides a practical benchmark for assessing how proposed market-structure reforms trade off regulatory costs against tangible investor welfare gains.

3 Extensions & Mechanisms

3.1 Cross-sectional evidence

To assess the external validity and economic generality of our findings, we extend the main analysis of execution quality beyond the Eurostoxx50 constituents to a broader cross-section of 523 EurostoxxTotal constituents. We evaluate execution quality for the full trading month of July 2024, adding 31 million trades to the sample.

[«Insert Tables 10 and 11»]

General execution-quality patterns remain virtually unchanged and confirm that latency and insufficient depth explain a substantial portion of price disimprovement, while a smaller residual is attributable to routing frictions and incomplete venue connectivity. Table 10 decomposes the origin of execution frictions for the cross-section. The sequential application of the four filters raises the share of trades executed at best prices from 87% in a latency-neutral scenario (Filter i) to 92% once we restrict to trades executed against globally sufficient depth (Filter iv); the fraction of worsened executions falls from 13% to 8%. Absolute slippage shrinks from €10.6 million to €4.8 million, and per-share slippage on worsened trades declines from 1.45 basis points to 1.24 basis points. Table 11 reports the distribution of price deviations, showing that roughly 8% of trades deviate from the EBBO, with large slippage events exceeding €0.10 being rare (less than 1% of trades) and price improvement occurring in only fewer than 1% of cases. The patterns are consistent across absolute and tick-based measures. We note that July 2024 is generally characterized by a relatively high execution quality and low volatility, as evidenced in Figure 4.

The cross-sectional results confirm the main findings from Section 2.5, but they also reveal important economic heterogeneity. Execution quality for the broader cross-section is marginally higher than for the Eurostoxx50 baseline. The improvement is largely driven by lower fragmentation: while an average Eurostoxx50 constituent trades on 4.8 venues, the average cross-sectional stock trades on only 4.3. Fewer active venues reduce routing complexity, synchronization errors, and latency-induced misalignments, thereby reducing execution shortfalls. From an equilibrium perspective, less fragmented trading implies a smaller informational dispersion across venues, reducing search costs for liquidity takers, and adverse-selection risk for liquidity providers.

The economic trade-off mirrors theoretical predictions from Chowdhry and Nanda (1991) and Madhavan (1995): as trading fragments, informed and uninformed order flow becomes segmented, widening effective spreads and increasing the probability of execution at stale prices. Our evidence shows that this mechanism operates primarily through the number of active venues and their physical separation. Stocks that trade on many geographically dispersed venues exhibit greater cross-venue latency differentials and more frequent quote incoherence, while more concentrated trading enhances price alignment and execution certainty. Moreover, differences in liquidity levels matter. The average cross-sectional stock is less liquid than the average Eurostoxx50 constituent, leading to fewer global price updates per unit of time. This slower quote dynamic effectively reduces the time sensitivity of routing decisions, allowing best-execution algorithms to synchronize more easily. The result is paradoxical but consistent with economic intuition: moderate illiquidity can reduce the frequency of informational shocks, thereby reducing latency-induced welfare losses.

4 Conclusion

Fragmentation and consolidation represent the core tension in modern market design. By reconstructing a nanosecond-level consolidated view of Europe's equity markets, we show that execution frictions stem less from a lack of liquidity than from its dispersion across venues and datacenters. Our results establish the EBBO as both an empirical benchmark and a policy instrument. For regulators, it provides a foundation for MiFID III's consolidated tape, and a basis for calibrating latency harmonization, venue-access obligations, and order-protection rules. For researchers, it offers a unified, high-frequency benchmark for studying trading, information transmission, and welfare in fragmented markets.

The EBBO uncovers substantial liquidity that remains inaccessible to many investors due to latency, routing frictions, and incomplete connectivity. Consolidated liquidity is strong—quoted spreads halve and depth triples—yet realized execution quality remains materially weaker: only 86% of marketable orders achieve best prices. The inefficiency is economically meaningful: routing frictions and incomplete connectivity generate annual price slippage of roughly €50 million per year equivalent to €198 thousand per day, imposing an implicit tax of 0.52 basis points per price-disimproved share on less-informed traders. The trade-off is clear: accept modest, transparent integration costs ex-ante to identify global best prices, or tolerate a persistent implicit premium ex-post.

A pragmatic first step toward price consolidation could build on technology that already exists in European markets. The Financial Information Exchange (FIX) protocol used by brokers and trading venues already includes a drop-copy function that distributes information on trades and quotes in real-time. By weaving together these drop-copy feeds, standard FIX messaging, and the current MiFID transparency framework—which already makes post-trade data public—regulators and market participants could create a real-time shadow consolidated tape for an ex-post assessment of execution quality. This approach repurposes established infrastructure, offering timely consolidated price information without major new technological investments, and demonstrates a practical path towards a full European consolidated tape.

Taken together, our evidence suggests that the European capital market is liquid, but poorly connected. A consolidated pre-trade benchmark such as the EBBO could align competition with efficiency, transforming fragmented order flow into a coherent market. Whether the gains from consolidation ultimately outweigh the costs of harmonization and infrastructure investment remains an empirical question—but one that can now be answered.

References

Admati, A.R., Pfleiderer, P., 1988. A theory of intraday patterns: Volume and price variability. The Review of Financial Studies 1, 3–40. doi:10.1093/rfs/1.1.3.

Aquilina, M., Budish, E., O'Neill, P., 2022. Quantifying the high-frequency trading "arms race". The Quarterly Journal of Economics 137, 493–564. doi:10.1093/qje/qjab032.

Aramian, F., Comerton-Forde, C., 2024. Listing exchanges versus other lit venues: Price formation. Working Paper. URL: https://static1.squarespace.com/static/6310c0b9bb63a25599f4418c/t/66fd5552dc736c48c64f74a8/1727878483183/PriceFormationFinal.pdf.

Battalio, R., Corwin, S.A., Jennings, R., 2016. Can brokers have it all? On the relation between make-take fees and limit order execution quality. The Journal of Finance 71, 2193–2237. doi:10.1111/jofi.12422.

Boehmer, E., Saar, G., Yu, L., 2005. Lifting the veil: An analysis of pre-trade transparency at the NYSE. The Journal of Finance 60, 783–815. doi:10.1111/j.1540-6261.2005.00746.x.

Brogaard, J., Hendershott, T., Riordan, R., 2014. High-frequency trading and price discovery. The Review of Financial Studies 27, 2267–2306. doi:10.1093/rfs/hhu032.

Budish, E., Cramton, P., Shim, J., 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. The Quarterly Journal of Economics 130, 1547–1621. doi:10.1093/qje/qjv027.

Budish, E., Lee, R.S., Shim, J.J., 2024. A theory of stock exchange competition and innovation: Will the market fix the market? Journal of Political Economy 132, 1209–1246. doi:10.1086/727284.

Cartea, A., Chang, P., Graumans, R., 2025. Anonymity, signaling, and collusion in limit order books. Working Paper. doi:10.2139/ssrn.5080700. Available at SSRN.

Chowdhry, B., Nanda, V., 1991. Multimarket trading and market liquidity. The Review of Financial Studies 4, 483–511. doi:10.1093/rfs/4.3.483.

Dyhrberg, A.H., Shkilko, A., Werner, I.M., 2025. The retail execution quality landscape. Journal of Financial Economics 168, 104051. doi:10.1016/j.jfineco.2025.104051.

Dzieliński, M., Hagströmer, B., Qu, C., 2024. Exchange competition, fragmentation, and market quality. Working Paper. doi:10.2139/ssrn.5063382. Available at SSRN.

Easley, D., O'Hara, M., Yang, L., 2016. Differential access to price information in financial markets. Journal of Financial and Quantitative Analysis 51, 1071–1110. doi:10.1017/S0022109016000491.

ESMA, 2007. Best Execution under MiFID I. Technical Report. European Securities and Markets Authority. URL: https://www.esma.europa.eu/sites/default/files/library/2015/11/07_050b.pdf.

ESMA, 2018. Final Report: Amendments to RTS 11 Tick Size Regime. Technical Report. European Securities and Markets Authority. URL: https://www.esma.europa.eu/press-news/esma-news/esma-publishes-final-report-tick-size-regime. eSMA71-99-1074, 14 December 2018.

Foucault, T., Kozhan, R., Tham, W.W., 2017. Toxic arbitrage. The Review of Financial Studies 30, 1053–1094. doi:10.1093/rfs/hhw103.

Gehrig, T., 1993. Intermediation in search markets. Journal of Economics & Management Strategy 2, 97–120. doi:10.1111/j.1430-9134.1993.00097.x.

Glosten, L.R., 1994. Is the electronic open limit order book inevitable? The Journal of Finance 49, 1127–1161. doi:10.1111/j.1540-6261.1994.tb02450.x.

Glosten, L.R., Milgrom, P., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71–100. doi:10.1016/0304-405X(85)90044-3.

Gresse, C., 2017. Effects of lit and dark market fragmentation on liquidity. Journal of Financial Markets 35, 1–20. doi:10.1016/j.finmar.2017.05.003.

Hasbrouck, J., 1995. One security, many markets: Determining the contributions to price discovery. The Journal of Finance 50, 1175–1199. doi:10.1111/j.1540-6261.1995.tb04054.x.

Hasbrouck, J., 2009. Trading costs and returns for U.S. equities: Estimating effective costs from daily data. The Journal of Finance 64, 1445–1477. doi:10.1111/j.1540-6261.2009.01469.x.

Hasbrouck, J., Saar, G., 2013. Low-latency trading. Journal of Financial Markets 16, 646–679. doi:10.1016/j.finmar.2013.05.003.

Hendershott, T., Jones, C.M., 2005. Island goes dark: Transparency, fragmentation, and regulation. The Review of Financial Studies 18, 743–793. doi:10.1093/rfs/hhi013.

Holden, C.W., Jacobsen, S., 2014. Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions. The Journal of Finance 69, 1747–1785. doi:10.1111/jofi.12127.

Holden, C.W., Pierson, M., Wu, J., 2023. In the blink of an eye: Exchange-to-sip latency and trade classification accuracy. Working Paper. doi:10.2139/ssrn.4441422. Available at SSRN.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335. doi:10.2307/1913210.

Lee, C.M., Ready, M.J., 1991. Inferring trade direction from intraday data. The Journal of Finance 46, 733–746. doi:10.1111/j.1540-6261.1991.tb02683.x.

Li, S., Ye, M., Zheng, M., 2023. Refusing the best price? Journal of Financial Economics 147, 317–337. doi:10.1016/j.jfineco.2022.11.004.

Madhavan, A., 1995. Consolidation, fragmentation, and the disclosure of trading information. The Review of Financial Studies 8, 579–603. doi:10.1093/rfs/8.3.579.

Menkveld, A.J., 2013. High frequency trading and the new market makers. Journal of Financial Markets 16, 712–740. doi:10.1016/j.finmar.2013.06.006.

Menkveld, A.J., 2017. High-frequency trading as viewed through an electron microscope. Financial Analysts Journal 74.

Menkveld, A.J., Zoican, M.A., 2017. Need for speed? Exchange latency and liquidity. The Review of Financial Studies 30, 1188–1228. doi:10.1093/rfs/hhx006.

Nasdaq, 2020. How trades speed between venues. https://www.nasdaq.com/articles/how-trades-speed-between-venues. Accessed: 2025-08-14.

O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality? Journal of Financial Economics 100, 459–474. doi:10.1016/j.jfineco.2011.02.006.

Ozturk, A.R., van der Wel, M., van Dijk, D., 2017. Intraday price discovery in fragmented markets. Journal of Financial Markets 32, 28–48. doi:10.1016/j.finmar.2016.10.001.

Pagano, M., Röell, A., 1996. Transparency and liquidity: A comparison of auction and dealer markets with informed trading. The Journal of Finance 51, 579–611. doi:10.1111/j.1540-6261.1996.tb02695.x.

Pagnotta, E.S., Philippon, T., 2018. Competing on speed. Econometrica 86, 1067–1115. doi:10.3982/ECTA10762.

Parlour, C.A., Seppi, D.J., 2003. Liquidity-based competition for order flow. The Review of Financial Studies 16, 301–343. doi:10.1093/rfs/hhg008.

Ramelli, S., Wagner, A.F., 2020. Feverish stock price reactions to COVID-19. The Review of Corporate Finance Studies 9, 622–655. doi:10.1093/rcfs/cfaa012.

Roşu, I., 2009. A dynamic model of the limit order book. The Review of Financial Studies 22, 4601–4641. doi:10.1093/rfs/hhp011.

Schwenk-Nebbe, S., Thimsen, C., 2024. Accurate Inference from TAQ using RNBBO. Working Paper. doi:10.2139/ssrn.3984827. Available at SSRN.

SEC, 2025. Amendments to Rule 612 of Regulation NMS for the quoting of certain NMS stocks. https://www.sec.gov/resources-small-businesses/small-business-compliance-guides/tick-sizes.

Wu, F., Zhan, X., Zhou, J., Wang, M., 2023. Stock market volatility and Russia–Ukraine conflict. Finance Research Letters 55, 103919. doi:10.1016/j.frl.2023.103919.

Tables

Table 1:
Order book statistics

	Unit	Q25	Mean	Median	Q75
Panel A: Equity characteristics					
Price	Euro	24.00	166.27	67.00	160.00
Tick size	Euro-Cent	0.50	4.01	1.00	5.00
Daily trades	Count	9,399	16,612	14,059	20,631
Daily trade volume	Thousand Euro	46,058	90,514	72,225	112,524
Market capitalization	Million Euro	41,310	73,119	52,010	88,511
Panel B: Liquidity					
Local quoted spread	Euro-Cent	1.68	14.81	4.49	11.40
Global quoted spread	Euro-Cent	0.84	6.41	2.28	6.83
Local quoted spread	Basis points	4.16	8.94	5.80	8.80
Global quoted spread	Basis points	2.28	3.55	2.88	4.21
Local effective spread	Euro-Cent	0.69	4.30	2.04	5.45
Global effective spread	Euro-Cent	0.61	3.93	1.85	5.11
Local effective spread	Basis points	1.97	3.86	2.92	4.34
Global effective spread	Basis points	1.79	3.35	2.65	3.95
Local depth	Shares	115.61	809.24	249.01	728.89
Global depth	Shares	304.10	2,237.62	654.43	2,016.92
Local depth	Euro	12,028.76	20,092.18	16,400.64	23,299.18
Global depth	Euro	28,750.93	56,198.75	44,204.98	70,130.31
Panel C: Competition					
Active venues	Count	4.23	4.77	5.00	6.00
Active venues at EBBO	Count	1.81	2.11	2.30	2.55

Note: Table 1 reports descriptive statistics for Eurostoxx50 constituents from 2020 to 2024. Local measures refer to liquidity and spreads observed within individual trading venues, whereas global measures refer to the consolidated market as represented by the EBBO. Panel A summarizes generic instrument characteristics and trading activity. Panel B presents liquidity measures, including quoted and effective spreads and depth at the inside quote. Panel C aggregates competition metrics as the number of actively quoting venues. We report the unit of observation, 25th percentile, mean, median and 75th percentile. To infer global order book statistics, we rely on the REBBO consolidated in Bergamo to homogenize transmission latencies. Spread measures are reported in Euro-Cent ($\frac{1}{100}$th Euro) and basis points and are restricted to normal EBBO market states. Variables are defined in Table A.1.

Table 2:
Venue details

MIC	Name	Group	Datacenter location	Maker fee	Taker fee	#EBBO updates	#Trades
Panel A: Regulated Markets							
XETR	Xetra	Deutsche Börse	Frankfurt, Germany	0.00	0.36	125,619,141	184,627,295
XHEL	Nasdaq Helsinki	Nasdaq Baltic	Upplands, Sweden	0.05	0.385	13,502,065	17,847,549
XMAD	Bolsa De Madrid	BME	Madrid, Spain	0.30	0.30	13,950,134	34,538,107
Pre-Euronext relocation							
XAMS	Amsterdam Exchange	Euronext	Basildon, United Kingdom	0.45	0.95	32,236,244	56,629,477
XBRU	Euronext Brussels	Euronext	Basildon, United Kingdom	0.45	0.95	5,999,344	7,806,342
XMIL	Borsa Italiana SPA	Euronext	Basiglio, Italy	0.45	0.95	16,119,943	43,552,546
XPAR	Euronext Paris	Euronext	Basildon, United Kingdom	0.45	0.95	75,072,543	130,949,862
Post-Euronext relocation							
XAMS	Amsterdam Exchange	Euronext	Bergamo, Italy	0.45	0.95	33,951,536	44,451,060
XBRU	Euronext Brussels	Euronext	Bergamo, Italy	0.45	0.95	2,277,227	4,131,525
XMIL	Borsa Italiana SPA	Euronext	Bergamo, Italy	0.45	0.95	19,306,743	40,176,135
XPAR	Euronext Paris	Euronext	Bergamo, Italy	0.45	0.95	67,188,676	92,890,142
Panel B: Multilateral Trading Facilities							
Pre-Brexit							
AQXE	Aquis Exchange	Aquis	Slough, United Kingdom	n.a.	n.a.	45,793,372	12,451,930
CHIX	Cboe Europe CXE	CBOE	Slough, United Kingdom	0.00	0.30	11,711,615	50,702,492
TRQX	Turquoise	LSE	London, United Kingdom	-0.29	0.30	7,902,726	17,138,651
Post-Brexit							
AQEU	Aquis Exchange Europe	Aquis	Slough, United Kingdom	0.00	0.63	97,714,573	58,769,603
CEUX	Cboe Europe DXE	CBOE	Slough, United Kingdom	0.00	0.30	65,474,656	290,946,361
TQEX	Turquoise Europe	LSE	London, United Kingdom	-0.24	0.30	26,896,780	86,271,015

Note: Table 2 provides descriptive metadata on the sample trading venues for Regulated Markets (RMs), Panel A, and Multilateral Trading Facilities (MTFs), Panel B. Venue types follow MiFID II classifications: A RM is an exchange formally authorized as official trading venue under MiFID II managed by a market operator. A MTF is a trading system managed by an investment firm or market operator. Both facilitate the interaction of multiple third-party buying and selling interests in financial instruments. We report the each venue's Market Identifier Code (MIC), name, group, datacenter location, maker and taker fee in basis points, the number of active EBBO price sets as #EBBO updates and the cumulative number of #Trades on the respective venue. Euronext migrated its datacenter in June 2022 from Basildon, United Kingdom, to Bergamo, Italy.

Table 3:
Millisecond transmission latencies

From	To	Haversine kilometer distance	Theoretical transmission time vacuum	Theoretical transmission time fiber	Empirical transmission time	Physical connection type
Basildon	Frankfurt	601	2.00	2.94	2.10	Microwave
Basildon	Madrid	1,279	4.27	6.26	9.14	Fiber
Basildon	Slough	77	0.26	0.38	0.35	Microwave
Basildon	Upplands	1,399	4.67	6.85	10.27	Fiber
Bergamo	Frankfurt	497	1.66	2.44	1.82	Microwave
Bergamo	London	954	3.18	4.67	3.05	Microwave
Bergamo	Madrid	1,234	4.12	6.05	11.67	Fiber
Bergamo	Slough	990	3.30	4.85	3.38	Microwave
Bergamo	Upplands	1,631	5.44	7.99	10.99	Fiber
Frankfurt	Madrid	1,454	4.85	7.12	11.59	Fiber
Frankfurt	Slough	675	2.25	3.31	2.38	Microwave
Frankfurt	Upplands	1,193	3.98	5.84	6.28	Fiber
London	Madrid	1,260	4.20	6.17	8.87	Fiber
London	Slough	42	0.14	0.21	0.46	Fiber
Madrid	Slough	1,253	4.18	6.13	9.18	Fiber
Slough	Upplands	1,458	4.86	7.14	10.25	Fiber

Note: Table 3 tabulates cross-datacenter transmission latencies. We rely on the haversine-distance measure to gauge the physical distance between two datacenters. To compute theoretical inter-datacenter transmission latencies as a benchmark to classify empirical transmission latencies, we multiply the haversine-distance with a signal's travel speed in vacuum (299,792 kilometers per second) and optic fiber cable (204,190 kilometers per second). Since the haversine-distance between venues does not necessarily reflect the actual physical cable routing or transmission path, we empirically estimate the effective inter-venue transmission latency. We compute the venue-pair quote reactiveness for Eurostoxx50 constituents on the full set of trades. The quote-reactiveness between two venues measures the duration between the recognition of a publicly observable signal on one venue, more precisely a trade, and the next consecutive quote revision on the other venue. We examine quote-reactiveness for all venue combinations per equity and group the venues into datacenter clusters. We report the median datacenter-pair transmission latency as empirical transmission time. Empirically estimated transmission times which are slightly above the vacuum transmission threshold but below the fiber transmission threshold can be attributed to the use of microwave routes. For brevity, we only report major datacenter routes in this table. Table A.4 details the distribution of estimated transmission latencies for the full sample of datacenters.

Table 4:
Market state statistics

	Unit	Q25	Mean	Median	Q75
Panel A: State occurrence					
Daily normal price updates	Count	5,740	17,009	11,713	21,417
Daily locking price updates	Count	52	928	191	685
Daily crossing price updates	Count	944	2,156	1,739	2,909
Panel B: Continuous trading in states					
Normal global market	Percent	93.96	90.95	98.10	99.22
Locked global market	Percent	0.05	5.76	0.40	2.48
Crossed global market	Percent	0.56	1.96	1.20	2.43
Panel C: Arbitrageable profit					
Local spread during cross	Euro	0.01	0.11	0.03	0.10
Global spread during cross	Euro	-0.05	-0.13	-0.02	-0.01
Average arbitrageable profit	Euro	1.11	10.79	2.04	5.76
Cumulative arbitrageable profit	Euro	93.24	25,347.51	305.37	1,580.23

Note: Table 4 reports descriptive statistics of global quote market states for Eurostoxx50 constituents from 2020 to 2024 per stock-day. Panel A summarizes the number of market state occurrences. Panel B reports state duration in percent of continuous trading time. Panel C estimates average and cumulative arbitrageable profits during crossed market states and reports local and global spreads during a cross. We report the unit of observation, 25th percentile, mean, median and 75th percentile. To infer global order book statistics, we rely on the REBBO consolidated in Bergamo to homogenize transmission latencies. Variables are defined in Table A.1 and winsorized on the upper 5% level.

Table 5:
What determines quoted and effective spreads?

	Quoted spread		Effective spread	
	(1)	(2)	(3)	(4)
Constant	1.0085***	1.0648***	0.4738***	0.2902***
	(10.609)	(6.547)	(10.807)	(5.036)
Active venues	-0.1238***	-0.1347***	-0.018***	0.0760***
	(-7.063)	(-6.863)	(-2.187)	(5.547)
Active venues at EBBO		-0.0105***		-0.1601***
		(-0.464)		(-10.152)
Log(Trade volume)		-0.3465***		-0.0751***
		(-5.608)		(-3.364)
Log(Market capitalization)		0.4314***		0.0903***
		(5.778)		(3.692)
Tick size		1.0659***		0.3478***
		(3.380)		(6.079)
1/Price		0.8375***		2.063***
		(15.662)		(10.313)
Equity fixed	No	Yes	No	Yes
Month fixed	No	Yes	No	Yes
Observations	63,705	63,696	63,706	63,582
R2	0.022	0.167	0.001	0.121

Note: Table 5 reports OLS estimates of determinants of local quoted and effective spreads in basis points. The sample consists of daily equity-level observations for all Eurostoxx50 constituents between 2020 and 2024. T-statistics are shown in parentheses. Asterisks ***, ** and * indicate statistical significance at the 1, 5, and 10% level respectively. Fixed effects are included as indicated. The sample comprises equity-day observations under Filter (iv). Variables are defined in Table A.1.

Table 6:
Execution quality by filter

	Latency-neutral EBBO	Latency-adjusted REBBOs	Local liquidity sufficient REBBOs	Global liquidity sufficient REBBOs
Panel A: Percentage by execution category				
at best	81.42	79.34	82.46	86.49
improved	0.44	0.63	0.48	0.61
worsened	18.14	20.03	17.06	12.91
Panel B: Percentage at best by venue				
AQEU	85.67	84.79	86.80	87.89
AQXE	85.27	83.14	84.82	89.96
CEUX	88.37	87.04	88.99	90.18
CHIX	77.70	75.19	77.59	82.93
TQEX	90.53	88.75	89.94	90.82
TRQX	81.47	78.32	80.17	85.19
XAMS	82.10	80.08	82.94	88.16
XBRU	86.40	81.92	84.36	87.42
XETR	80.54	76.94	78.90	84.43
XHEL	80.20	77.02	80.45	85.99
XMAD	77.21	75.26	80.37	84.14
XMIL	79.19	77.13	80.87	85.72
XPAR	81.86	79.32	81.36	85.67
#Trades	1,173,880,092	1,173,874,190	1,108,176,417	1,062,746,222
#Shares	352,039,325,754	352,018,137,247	308,573,673,207	283,725,393,435

Note: Table 6 summarizes execution quality by filter. Panel A reports execution outcomes by execution category. *At best* denotes trades receiving best globally available prices, *worsened* captures trades executed at worse than globally available prices subject to slippage, *improved* captures trades executed at better than globally available prices. Panel B reports the percentage of shares executed at best per venue. All measures are computed as the fraction of shares executed within each respective category after full sample cumulation over all days and stocks. We report the number of #Trades and the number of #Shares after application of each filter.

Table 7:
Changes in execution quality by filter

	Latency-neutral EBBO	Latency-adjusted REBBOs	Local liquidity sufficient REBBOs	Global liquidity sufficient REBBOs
Panel A: Percentage by execution category				
at best	81.42	79.34	82.46	86.49
improved	0.44	0.63	0.48	0.61
worsened	18.14	20.03	17.06	12.91
Δ at best		-2.08	3.12	4.03
Δ improved		0.19	-0.15	0.13
Δ worsened		1.89	-2.97	-4.15
Panel B: Absolute slippage in Euro				
at best	0	0	0	0
improved	6,467,680	14,974,450	10,212,120	10,809,897
worsened	-476,969,108	-760,019,317	-593,732,536	-254,665,393
Δ at best		0	0	0
Δ improved		8,506,770	-4,762,330	597,777
Δ worsened		-283,050,209	166,286,781	339,067,143
Panel C: Per-share slippage in basis points				
at best	0	0	0	0
improved	0.25	0.43	0.44	0.40
worsened	-0.56	-0.81	-0.84	-0.52
Δ at best		0	0	0
Δ improved		-0.25	-0.03	0.32
Δ worsened		0.18	0.01	-0.04
#Trades	1,173,880,092	1,173,874,190	1,108,176,417	1,062,746,222
#Shares	352,039,325,754	352,018,137,247	308,573,673,207	283,725,393,435

Note: Table 7 summarizes the changes in adverse selection cost by filter. Panel A reports execution outcomes by execution category. Panel B reports cumulative slippage in Euro, and Panel C reports per-share slippage in basis points for each execution category. *At best* denotes trades receiving best globally available prices, *worsened* captures trades executed at worse than globally available prices subject to slippage, *improved* captures trades executed at better than globally available prices. Additionally, we report the absolute change per execution category as Δat best, Δworsened and Δimproved. All measures are computed as the fraction of shares executed within each respective category after full sample cumulation over all days and stocks. We report the number of #Trades and the number of #Shares after application of each filter.

Table 8:
Execution quality decomposition for global liquidity sufficient REBBOs

Panel A: Euro-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 10 Cent	9,242,983	244,937,614	-94,535,300	0.09
worsened up to 10 Cent	12,839,449	369,529,127	-30,266,024	0.13
worsened up to 5 Cent	35,664,742	2,386,418,268	-60,417,976	0.84
worsened up to 1 Cent	22,342,872	2,852,645,235	-25,522,454	1.01
worsened up to 0.5 Cent	45,215,927	30,765,583,991	-43,923,639	10.84
executed at best	929,996,863	245,380,062,852	0	86.49
improved up to 0.5 Cent	2,737,617	1,388,940,286	1,879,850	0.49
improved up to 1 Cent	1,459,325	169,294,291	1,225,282	0.06
improved up to 5 Cent	2,028,072	136,450,961	2,560,510	0.05
improved up to 10 Cent	472,540	17,825,633	1,188,554	0.01
improved more than 10 Cent	745,832	13,705,177	3,955,701	0.00
Total worsened	125,305,973	36,619,114,235	-254,665,393	12.91
Total improved	7,443,386	1,726,216,348	10,809,897	0.61

Panel B: Tick-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 20 ticks	1,941,156	569,486,924	-72,136,745	0.20
worsened up to 20 ticks	1,195,471	406,532,065	-11,639,784	0.14
worsened up to 10 ticks	3,023,770	1,156,323,249	-15,867,940	0.41
worsened up to 5 ticks	36,452,718	12,134,131,543	-75,666,516	4.28
worsened up to 1 ticks	82,584,969	22,327,820,824	-70,641,336	7.87
executed at best	929,922,233	245,370,812,005	0	86.49
improved up to 1 ticks	2,938,312	673,124,632	2,379,644	0.24
improved up to 5 ticks	4,195,403	973,571,025	3,958,670	0.34
improved up to 10 ticks	121,283	34,299,386	648,102	0.01
improved up to 20 ticks	81,697	19,241,167	906,697	0.01
improved more than 20 ticks	106,380	25,967,732	2,916,454	0.01
Total worsened	125,198,084	36,594,294,605	-245,952,321	12.90
Total improved	7,443,075	1,726,203,942	10,809,567	0.61

Note: Table 8 reports residual slippage cost by price deviation bucket in Cents, Panel A, and in ticks, Panel B. Filter (iv) is applied which subsets all trades to those executed against global liquidity sufficient REBBOs. We report the total number of #Trades, the total number of #Shares, the cumulative Euro slippage cost and the relative share in percent as the fraction of shares executed within each deviation bucket after full sample cumulation over all days and stocks. We acknowledge minor discrepancies in execution quality between Euro-based and tick-based execution which arise from missing tick size information for a handful of stock-days.

Table 9:
What determines best execution?

	Probability of receiving best prices			
	(1)	(2)	(3)	(4)
Constant	0.906***	0.771***	0.906***	0.771***
	(6,517.740)	(612.092)	(6,517.740)	(612.092)
Geodistance (100km)	-0.010***	-0.013***		
	(-234.358)	(-234.196)		
Transmission latency			-0.017***	-0.023***
			(-234.358)	(-234.196)
Venues at best		0.028***		0.028***
		(270.513)		(270.513)
Bid		0.003***		0.003***
		(11.195)		(11.195)
Log(execution size)		-0.005***		-0.005***
		(-53.175)		(-53.175)
Locked global market		-0.080***		-0.080***
		(-150.830)		(-150.830)
Crossed global market		-0.316***		-0.316***
		(-342.015)		(-342.015)
Global quoted spread		-0.001***		-0.001***
		(-11.459)		(-11.459)
Local quoted spread		-0.002***		-0.002***
		(-4.538)		(-4.538)
Time since last EBBO update		0.000		0.000
		(0.948)		(0.948)
Midpoint volatility		-0.072***		-0.072***
		(-20.505)		(-20.505)
Executed at primary		0.017***		0.017***
		(11.009)		(11.009)
Executed at MTF		0.077***		0.077***
		(81.329)		(81.329)
Equity fixed	No	Yes	No	Yes
Venue fixed	No	Yes	No	Yes
Month fixed	No	Yes	No	Yes
Trade sample	Filter (iv)	Filter (iv)	Filter (iv)	Filter (iv)
Observations	5,915,866	5,564,764	5,915,866	5,564,764
R2	0.012	0.131	0.012	0.131

Note: Table 9 shows OLS regression results for execution quality. Dependent variable is execution quality with $y = 1$ indicating best execution and $y = 0$ indicating worse execution. T-statistics are shown in parenthesis. Asterisks ***, ** and * indicate statistical significance at the 1, 5, and 10% level respectively. For computational flexibility, we sample 100 trades per day and instrument over the full period from 2020 until 2024. All sampled trades are executed against global liquidity sufficient REBBOs (Filter iv). Variables are defined in Table A.1.

Table 10:
Changes in execution quality by filter (cross-section)

	Latency-neutral EBBO	Latency-adjusted REBBOs	Local liquidity sufficient REBBOs	Global liquidity sufficient REBBOs
Panel A: Percentage by execution quality				
at best	87.10	86.25	89.17	91.86
improved	0.38	0.27	0.18	0.24
worsened	12.51	13.48	10.65	7.90
Δ at best		-0.85	2.92	2.69
Δ improved		-0.11	-0.09	0.06
Δ worsened		0.97	-2.83	-2.75
Panel B: Absolute slippage in Euro				
at best	0	0	0	0
improved	245,433	129,657	76,930	90,221
worsened	-10,596,266	-10,828,954	-7,693,632	-4,792,761
Δ at best		0	0	0
Δ improved		-115,776	-52,727	13,291
Δ worsened		-232,688	3,135,322	2,900,871
Panel C: Per-share slippage in basis points				
at best	0	0	0	0
improved	0.73	0.53	0.52	0.50
worsened	-1.45	-1.37	-1.39	-1.24
Δ at best		0	0	0
Δ improved		0.08	-0.02	0.15
Δ worsened		-0.2	-0.01	-0.02
#Trades	31,028,009	30,169,559	28,450,599	27,381,237
#Shares	9,146,366,662	8,935,993,086	7,824,575,088	7,279,635,955

Note: Table 10 summarizes the changes in adverse selection cost by filter for the cross-sectional sample of 523 EurostoxxTotal constituents. Panel A reports execution outcomes by execution category. Panel B reports cumulative slippage in Euro, and Panel C reports per-share slippage in basis points for each execution category. *At best* denotes trades receiving best globally available prices, *worsened* captures trades executed at worse than globally available prices subject to slippage, *improved* captures trades executed at better than globally available prices. Additionally, we report the absolute change per execution category as Δat best, Δworsened and Δimproved. All measures are computed as the fraction of shares executed within each respective category after cumulation over all trading days in July 2024 and all 523 stocks. We report the number of #Trades and the number of #Shares after application of each filter.

Table 11:
Execution quality decomposition for global liquidity sufficient REBBOs (cross-section)

Panel A: Euro-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 10 Cent	156,474	6,746,668	-1,461,132	0.09
worsened up to 10 Cent	200,357	7,702,368	-638,216	0.11
worsened up to 5 Cent	694,682	49,330,484	-1,274,258	0.68
worsened up to 1 Cent	401,273	64,275,619	-585,036	0.88
worsened up to 0.5 Cent	836,715	446,708,425	-834,119	6.14
executed at best	24,996,295	6,687,420,036	0	91.86
improved up to 0.5 Cent	42,488	13,794,228	25,608	0.19
improved up to 1 Cent	16,678	1,927,984	14,312	0.03
improved up to 5 Cent	24,307	1,533,245	28,648	0.02
improved up to 10 Cent	4,489	127,317	7,779	0.00
improved more than 10 Cent	7,479	69,581	13,873	0.00
Total worsened	2,289,501	574,763,564	-4,792,761	7.90
Total improved	95,441	17,452,355	90,221	0.24

Panel B: Tick-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 20 ticks	33,589	5,179,997	-902,698	0.07
worsened up to 20 ticks	21,464	3,878,418	-227,070	0.05
worsened up to 10 ticks	49,293	13,673,703	-305,534	0.19
worsened up to 5 ticks	611,986	178,775,772	-1,671,092	2.46
worsened up to 1 ticks	1,573,043	373,236,498	-1,686,030	5.13
executed at best	24,995,532	6,687,192,152	0	91.86
improved up to 1 ticks	47,712	7,614,750	42,513	0.10
improved up to 5 ticks	47,298	9,756,964	44,084	0.13
improved up to 10 ticks	231	59,284	1,114	0.00
improved up to 20 ticks	37	2,900	138	0.00
improved more than 20 ticks	157	16,871	2,360	0.00
Total worsened	2,289,375	574,744,388	-4,792,425	7.90
Total improved	95,435	17,450,769	90,210	0.24

Note: Table 11 reports residual slippage cost by price deviation bucket in Cents, Panel A, and in ticks, Panel B, for the cross-sectional sample of 523 EurostoxxTotal constituents. Filter (iv) is applied which subsets all trades to those executed against global liquidity sufficient REBBOs. We report the total number of #Trades, the total number of #Shares, the cumulative Euro slippage cost and the relative share in percent as the fraction of shares executed within each deviation bucket after full sample cumulation over all trading days in July 2024 and all 523 stocks. We acknowledge minor discrepancies in execution quality between Euro-based and tick-based execution which arise from missing tick size information for a handful of stock-days.

Table 12:
Percentage at best by latency correction mechanism

	Latency-neutral EBBO	Latency-adjusted REBBOs	Local liquidity sufficient REBBOs	Global liquidity sufficient REBBOs
Theoretical vacuum	81.42	81.01	83.95	87.54
Theoretical fiber	81.42	79.34	82.46	86.49
Empirical	81.42	80.67	83.61	87.31

Note: Table 12 summarizes execution quality by latency correction mechanism for the full Eurostoxx50 sample. We rely on theoretical and empirical transmission latencies tabulated in Table 3 to physically align quotes. Our results remain robust, whether accounting for theoretical or empirical transmission latencies.

Figures

Figure 1:
European market infrastructure

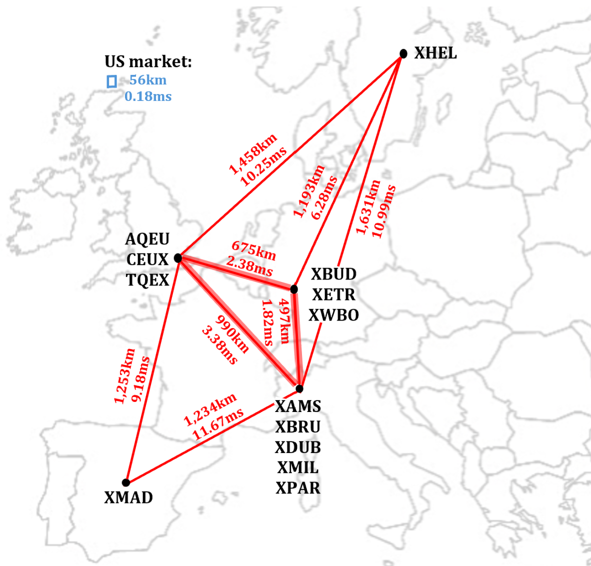

Note: Figure 1 illustrates physical distances in kilometers (km) and empirically estimated transmission latencies in milliseconds (ms) between major European datacenters. Bold connections indicate microwave links which establish latency corridors between datacenters. The map is depicted after Euronext relocation to Bergamo. Venues are grouped into four regional clusters: London/Slough, United Kingdom (AQEU, CEUX, TQEX), Frankfurt, Germany (XETR, XBUD, XWBO), Bergamo, Italy (XAMS, XBRU, XDUB, XMIL, XPAR), and Stockholm (XHEL). The US market structure is shown for reference in the top left. The distance between the two datacenters located in Mahwah (NYSE) and Carteret (NASDAQ), New Jersey, accounts to 56 kilometers (35 miles) with a transmission latency of 0.18 milliseconds. By contrast, European venue dispersion spans an average of 611 kilometers (379 miles), with a minimum transmission latency of 2.99 milliseconds in optic fiber. For additional information on the US regarding inter-datacenter distances and transmission latencies, see Nasdaq (2020) and Holden et al. (2023).

Figure 2:
Crossed markets time series

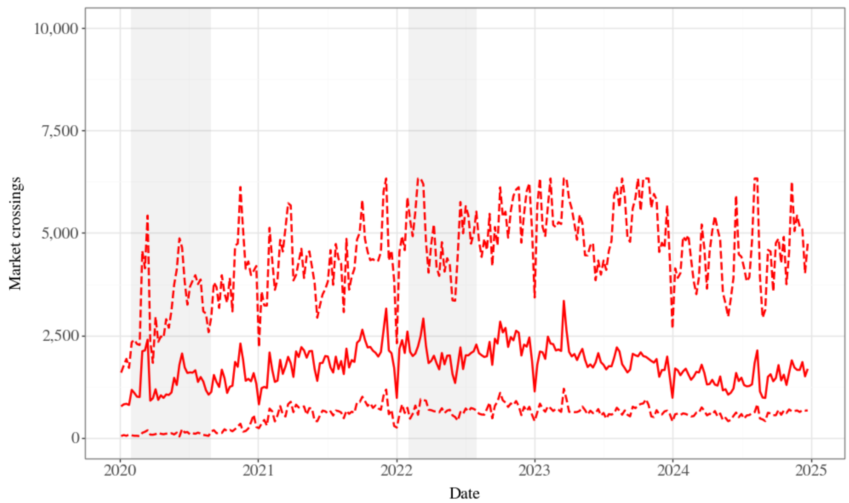

Note: Figure 2 visualizes the number of absolute daily market crossings as average per stock-day. The solid line presents the cross-sectional average, the dashed lines indicate 10th and 90th percentile. Global liquidity Filter (iv) is applied which subsets all trades to those executed against global liquidity sufficient REBBOs. Shaded areas highlight periods of global market stress: the COVID onset in January 2020 (Ramelli and Wagner, 2020), and the Russia–Ukraine invasion on 24 February 2022 (Wu et al., 2023). We report average weekly metrics for ease of visual interpretation.

Figure 3:
Spread time series

Panel A: Basis point local quoted spread

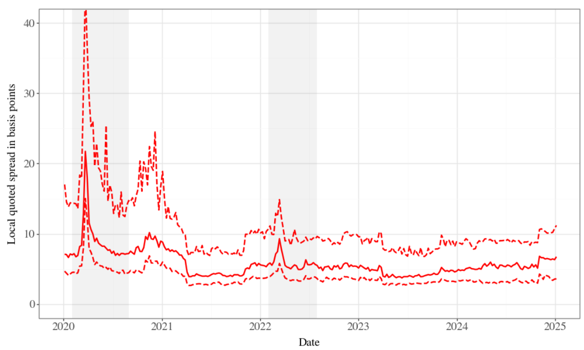

Panel B: Basis point global quoted spread

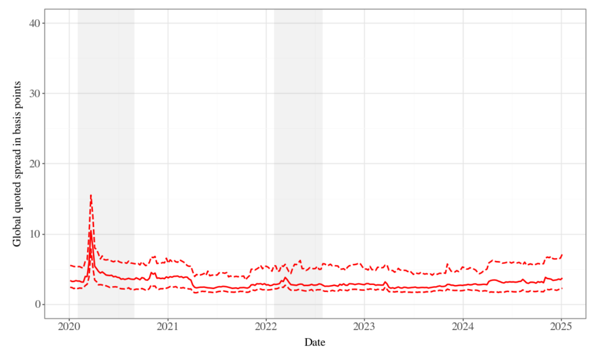

Panel C: Basis point local effective spread

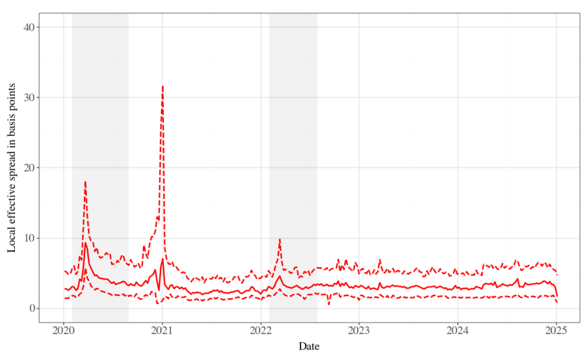

Panel D: Basis point global effective spread

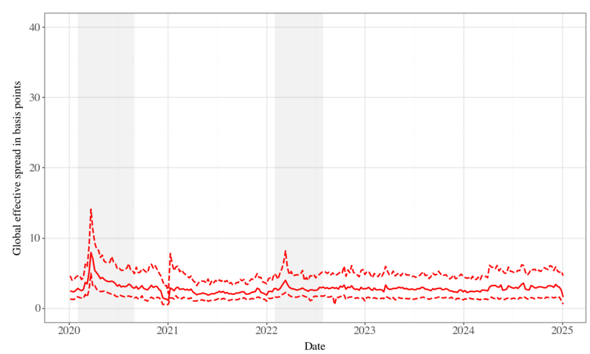

Note: Figure 3 visualizes spread time series as averages per stock-day. The solid line presents the cross-sectional average, the dashed lines indicate 10th and 90th percentile. Global liquidity filter (Filter iv) is applied. Panel A and B represent the quoted local and global market spreads in basis points respectively. Panel C and D plot the effective local and global market spreads in basis points respectively. Shaded areas highlight periods of global market stress: the COVID onset in January 2020 (Ramelli and Wagner, 2020), and the Russia–Ukraine invasion on 24 February 2022 (Wu et al., 2023). We report average weekly metrics for ease of visual interpretation and winsorize spreads on the upper 1% level.

Figure 4:
Execution quality time series

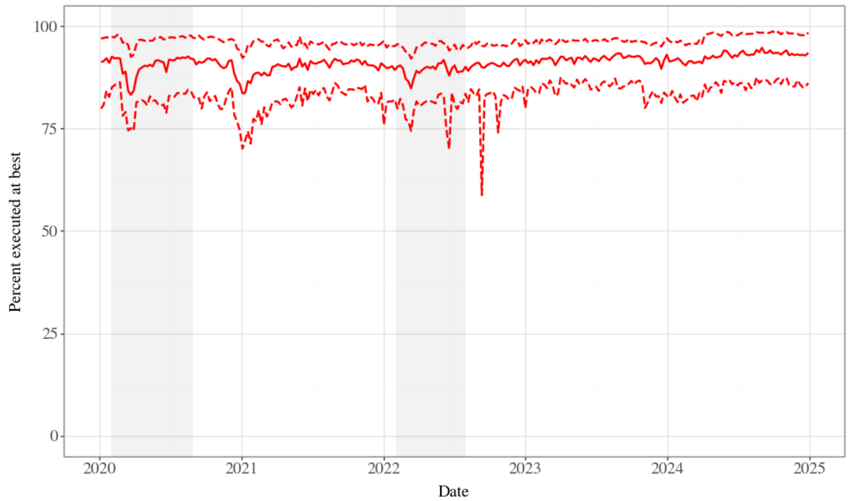

Note: Figure 4 visualizes the percentage of shares executed at best as average per stock-day. The solid line presents the cross-sectional average, the dashed lines indicate 10th and 90th percentile. Global liquidity filter (Filter iv) is applied. Shaded areas highlight periods of global market stress: the COVID onset in January 2020 (Ramelli and Wagner, 2020), and the Russia–Ukraine invasion in 24 February 2022 (Wu et al., 2023). We report average weekly metrics for ease of visual interpretation.

Figure 5:
Changes in execution quality by filter

Panel A: Percent executed at best

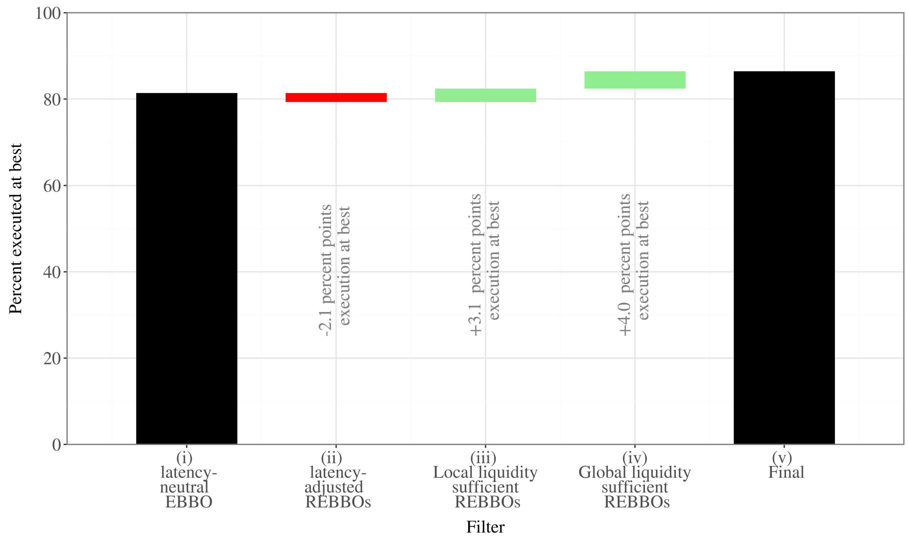

Panel B: Absolute slippage in Euro

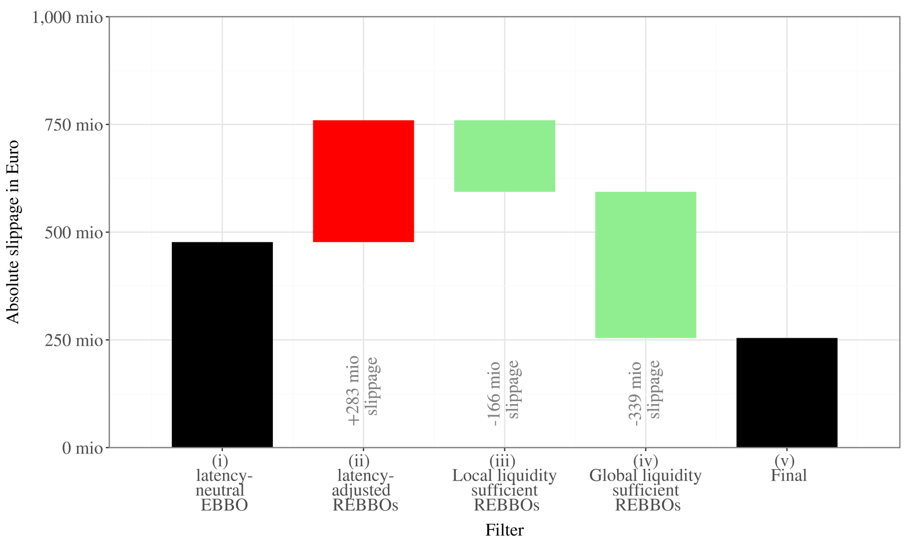

Note: Figure 5 shows the percentage of shares executed at best, Panel A, and the absolute slippage in Euro, Panel B, by filter. The baseline scenario, Filter (i), illustrates best execution for trades matched against latency-neutral EBBOs. Filter (ii) illustrates best execution for trades matched against latency-adjusted REBBOs, Filter (iii) subsets all trades to those executed against local liquidity sufficient REBBOs, Filter (iv) subsets all trades to those executed against global liquidity sufficient REBBOs. The final scenario (Filter v) quantifies effective slippage after accounting for mechanical and latency-induced effects which is attributable to inefficient routing and incomplete venue connectivity. The Figures visualize values tabulated in Table 7.

Online appendix

A.1 Supplementary tables

Table A.1:
Variable description

	Unit	Description
Panel A: Core variables		
Local		Metrics on individual venues.
Global		Metrics on the consolidated (R)EBBO market.
EBBO	Euro	Consolidated best bid and best offer across all actively quoting venues at each timestamp; latency-neutral benchmark as if all venues were co-located.
REBBO	Euro	Consolidated best bid and best offer across all actively quoting venues at each timestamp; latency-adjusted, venue-specific after venue-pair specific lagging of timestamps benchmark.
Quoted spread	Euro/bps/ticks	Difference between best bid and best ask. Measured both locally per venue and globally on the consolidated (R)EBBO.
Effective spread	Euro/bps/ticks	Difference between trade's execution price and midpoint. Measured both locally per venue and globally on the consolidated (R)EBBO.
Depth	Shares	Number of shares available at the best bid and best ask level on a venue's local order book during continuous trading. Measured both locally per venue and globally on the consolidated (R)EBBO.
Slippage	Euro/bps/ticks	Difference between a trade's execution price and the prevailing best global (R)EBBO price on the same side at the time of trade execution. Measured both locally per venue and globally on the consolidated (R)EBBO.
Normal	Binary	Global market is normal (positive global quoted spread).
Locked	Binary	Global market is locked (zero global quoted spread).
Crossed	Binary	Global market is crossed (negative global quoted spread).
At best	Binary	Trade that is executed at best globally available prices.
Worsened	Binary	Trade that is executed at worse than best globally available prices.
Improved	Binary	Trade that is executed at better than best globally available prices.
Tick size	Euro	Smallest price increment per ISIN and day measured as smallest price difference between all trades.
Panel B: Control variables		
Active venues	Count	Number of venues actively quoting prices for an instrument.
Active venues at EBBO	Count	Number of venues actively quoting best prices for an instrument.
Trade volume	Log(Notional)	Cumulative daily trade volume in logarithmic notional.
Market capitalization	Log(market capitalization)	Market capitalization in logarithmic million Euro.
1/Price	Ratio	Reciprocal of the average daily trade price.
Geodistance	Kilometers	Haversine-distance between trade's execution venue and venues currently quoting best prices.
Transmission latency	Milliseconds	Average transmission latency based on theoretical transmission latencies in optic fiber cable between trade's execution venue and venues currently quoting best prices.
Bid	Binary	Trade is a buy.
Execution size	Log(shares)	Trade's size in logarithmic shares.
Time since last EBBO update	Log(milliseconds)	Time difference between last EBBO update and trade execution.
Midpoint volatility	Deviation	Standard deviation of the global midpoint in the 10 seconds prior trade.
Executed at primary	Binary	Trade is executed at its primary listing exchange.
Executed at MTF	Binary	Trade is executed at a MTF.

Note: Table A.1 defines core and control variables used in the analysis. Central variables are calculated on local markets, referring to metrics measured on individual trading venues, and the consolidated global market, referring to metrics measured on the consolidated (R)EBBO. For readability, basis points are abbreviated as bps.

Table A.2:
Institutional differences between European and US markets

	EU	US
Panel A: Fragmentation		
Trading venues	41 trading, 35 listing.	16 trading, 3 listing.
Market geography	Dispersed. Average distance between datacenters of 611 km (ranging from 77 to 2,600 km).	Clustered in New Jersey. Average distance between datacenters of 56 km (ranging from 13 to 89 km) (Nasdaq, 2020; Holden et al., 2023).
Fragmentation	Average stock trades on 4.7 venues. Price discovery mainly via listing venue (Aramian and Comerton-Forde, 2024).	Price discovery and majority of trading volume mainly via listing venue (O'Hara and Ye, 2011; Ozturk et al., 2017).
Panel B: Data		
Availability	Only direct feeds; no central aggregator.	Direct feeds next to publicly available aggregator.
Consolidated tape	None. Expected by 2026 as EBBO generated from direct feeds.	Available. NBBO as mandatory instance (Reg NMS, Rule 603).
Panel C: Regulation		
Regulation	MiFID II (2018): mandates transparency and tick rules.	Reg NMS (2007): mandates consolidation and competition.
Tick size	Variable tick size bands under MiFID II ranging from €0.0001 to €10 conditional on an equity's price and transaction volume (ESMA, 2018).	Fixed at §0.01 for all equities (Reg NMS, Rule 612). As of November 2025, §0.005 or §0.01 conditional on liquidity band of the stock (SEC, 2025).
Best execution	Loosely defined. Market participants should take *"all reasonable steps"* (Directive 2014/65/EU, Art. 27(1); Delegated Directive (EU) 2017/593, Art. 66(1)) to ensure best execution.	Explicit. NBBO-based routing (Reg NMS, Rule 611) to enforce best price execution. Requires exchanges to implement policies which ensure that trades do not occur at worse than best available prices.

Note: Table A.2 contrasts key institutional and regulatory features of European and US equity markets, including listing structures, venue fragmentation, data availability, and regulatory frameworks.

Table A.3:
Instrument details

Name	Country code	ISIN	#Venues	Price	Tick size	Local spread	#EBBO updates	#Trades	Market capital-ization
ADIDAS	DE	DE000A1EWWW0	4.44	219	0.050	0.119	26,501,715	21,071,646	41,909
ADYEN	NL	NL0012969182	5.02	1,445	0.200	0.880	37,790,464	23,345,514	46,310
AIRBUS	NL	NL0000235190	7.04	114	0.020	0.075	61,946,562	36,572,439	89,041
ALLIANZ	DE	DE0008404005	5.03	212	0.050	0.096	24,750,009	20,937,537	86,920
ANHEUSER-BUSCH	BE	BE0974293251	6.03	54	0.010	0.035	24,768,714	22,035,121	93,180
ASML HOLDING	NL	NL0010273215	6.03	578	0.100	0.244	50,155,880	57,678,117	231,723
AXA	FR	FR0000120628	6.09	26	0.005	0.012	24,799,182	22,368,449	59,372
BANCO SANTANDER	ES	ES0113900J37	6.03	3	0.001	0.002	27,240,686	25,823,861	53,018
BASF	DE	DE000BASF111	4.49	49	0.005	0.021	35,366,749	24,396,414	47,700
BAYER	DE	DE000BAY0017	4.50	51	0.010	0.021	26,306,793	24,013,408	47,403
BBV.ARGENTARIA	ES	ES0113211835	5.43	6	0.001	0.005	17,732,217	18,397,217	37,110
BMW	DE	DE0005190003	5.05	83	0.010	0.035	33,064,623	21,164,497	49,157
BNP PARIBAS	FR	FR0000131104	6.08	54	0.010	0.032	47,666,125	35,196,580	63,627
DANONE	FR	FR0000120644	5.06	57	0.010	0.025	12,625,814	16,427,574	39,143
DEUTSCHE BOERSE	DE	DE0005810055	5.03	162	0.050	0.106	7,589,427	10,395,439	30,948
DHL Group	DE	DE0005552004	4.51	40	0.010	0.020	22,339,378	18,841,171	51,002
DEUTSCHE TELEKOM	DE	DE0005557508	5.09	18	0.005	0.009	16,570,093	19,266,465	92,845
ENEL	IT	IT0003128367	5.05	7	0.001	0.003	21,722,293	23,878,541	67,570
ENI	IT	IT0003132476	5.04	13	0.002	0.006	24,741,601	21,360,368	42,514
ESSILORLUXOTTICA	FR	FR0000121667	5.66	165	0.050	0.107	15,259,555	16,013,843	72,611
FERRARI	NL	NL0011585146	5.02	202	0.050	0.184	8,387,260	12,698,052	44,541
HERMES	FR	FR0000052292	5.02	1,338	0.500	0.923	14,028,581	13,663,950	151,407
IBERDROLA	ES	ES0144580Y14	5.45	11	0.005	0.010	5,944,309	16,726,663	69,714
INDITEX	ES	ES0148396007	6.02	29	0.010	0.023	7,358,584	13,271,175	97,864
INFINEON	DE	DE0006231004	5.08	32	0.005	0.014	36,949,655	28,556,603	39,800
ING GROEP	NL	NL0011821202	6.07	11	0.002	0.006	48,955,675	35,228,043	41,110
INTESA SANPAOLO	IT	IT0000072618	5.03	2	0.001	0.001	26,483,344	25,392,249	45,022

(continued)

Name	Country code	ISIN	#Venues	Price	Tick size	Local spread	#EBBO updates	#Trades	Market capital-ization
(continued)									
KONINKLIJKE AHOLD	NL	NL0011794037	6.05	28	0.005	0.016	10,934,103	16,123,826	28,102
AIR LIQUIDE	FR	FR0000120073	5.04	148	0.020	0.049	18,342,568	21,053,732	75,719
L'OREAL	FR	FR0000120321	5.04	365	0.050	0.124	23,675,633	24,367,759	195,199
LVMH	FR	FR0000121014	6.03	655	0.100	0.243	35,441,961	42,819,187	319,032
MERCEDES-BENZ	DE	DE0007100000	5.08	64	0.010	0.025	36,135,553	29,171,700	65,857
MUNICH RE	DE	DE0008430026	5.03	263	0.100	0.198	10,518,559	12,153,331	42,056
NOKIA	FI	FI0009000681	7.02	4	0.001	0.004	36,658,103	21,713,409	22,736
NORDEA BANK	FI	FI4000297767	4.75	10	0.001	0.006	19,488,874	10,572,068	35,133
PERNOD-RICARD	FR	FR0000120693	5.03	169	0.050	0.101	7,662,472	12,827,032	43,389
PROSUS N	NL	NL0013654783	5.04	66	0.010	0.027	34,001,989	33,780,826	118,570
RHEINMETALL	DE	DE0007030009	4.25	177	0.050	0.195	10,275,994	8,862,934	9,408
SAFRAN	FR	FR0000073272	4.66	120	0.020	0.056	22,302,538	20,171,896	56,287
SAINT GOBAIN	FR	FR0000125007	5.05	54	0.010	0.025	21,749,104	17,831,865	27,924
SANOFI	FR	FR0000120578	6.05	89	0.010	0.033	31,654,926	32,210,785	113,834
SAP	DE	DE0007164600	4.48	121	0.020	0.044	30,882,163	33,341,244	157,955
SCHNEIDER ELECTRIC	FR	FR0000121972	5.04	144	0.020	0.056	25,739,555	25,061,187	86,178
SIEMENS	DE	DE0007236101	5.04	137	0.020	0.048	33,233,636	27,876,203	112,893
STELLANTIS	NL	NL00150001Q9	5.99	16	0.002	0.007	51,757,753	29,067,341	43,688
TOTALENERGIES	FR	FR0000120271	5.11	51	0.010	0.019	41,740,843	46,551,021	124,123
UNICREDIT	IT	IT0005239360	5.03	12	0.002	0.009	40,818,484	27,350,597	32,503
VINCI	FR	FR0000125486	5.05	94	0.020	0.042	20,797,346	20,537,173	57,075
VOLKSWAGEN Vz.	DE	DE0007664039	5.04	135	0.020	0.070	37,171,179	24,987,462	29,274
WOLTERS KLUWER	NL	NL0000395903	4.66	99	0.020	0.066	6,940,337	10,603,661	26,461

Note: Table A.3 provides descriptives per sample instrument. We report the name, country code, ISIN, the number of actively quoting #Venues, the average trading price in Euro, median tick size in Euro as the smallest observable price increment, median local spread in Euro, the cumulative number of #EBBO updates and #Trades, and market capitalization in million Euro.

Table A.4:
Distribution of millisecond transmission latencies

From	To	#Trades	Q10	Q25	Q50	Q75	Q90
Basildon	Frankfurt	1,743,513	0.20	1.05	2.10	5.74	59.23
Basildon	London	41,910,979	0.01	0.04	0.16	0.30	3.25
Basildon	Madrid	33,861	0.34	1.56	9.14	24.43	135.66
Basildon	Slough	145,895,102	0.03	0.14	0.35	0.77	7.36
Basildon	Upplands	8,348,945	2.26	9.09	10.27	11.71	93.15
Bergamo	Frankfurt	3,474,901	0.46	1.24	1.82	3.97	34.03
Bergamo	London	42,588,309	0.27	1.36	3.05	3.54	7.28
Bergamo	Madrid	61,953	0.46	1.99	11.67	21.41	130.92
Bergamo	Slough	189,304,449	0.27	1.53	3.38	3.49	7.85
Bergamo	Upplands	12,139	0.93	3.57	10.99	13.33	77.75
Frankfurt	London	1,747,681	1.17	2.72	8.31	208.48	1,305.95
Frankfurt	Madrid	126,481	1.51	4.40	11.59	48.84	457.19
Frankfurt	Slough	263,621,235	0.38	1.47	2.38	3.77	11.81
Frankfurt	Upplands	443,473	0.98	3.50	6.28	14.08	145.83
London	Madrid	7,673,637	1.46	4.99	8.87	9.87	13.50
London	Slough	40,420,921	0.14	0.25	0.46	29.27	1,235.28
London	Upplands	1,140	18.28	37.56	56.14	1,038.09	2,630.16
Madrid	Slough	33,409,781	1.68	5.39	9.18	10.10	15.29
Slough	Upplands	138,801	3.92	8.97	10.25	10.77	447.89

Note: Table A.4 reports the distribution of inter-datacenter transmission latencies between European datacenter pairs. Latencies are measured in milliseconds and computed from the venue-pair quote-reactiveness of all trades in the Eurostoxx50 constituents during the full sample period. We report quote-reactiveness for all venue combinations per equity and group the venues into datacenter clusters. We report the 10th, 25th, 50th, 75th and 90th percentile. Table 3 contrasts the median estimated transmission latencies with theoretical transmission latencies and details the estimation procedure.

Table A.5:
Execution quality decomposition for EBBOs

Panel A: Euro-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 10 Cent	10,887,473	531,949,215	-196,356,111	0.15
worsened up to 10 Cent	14,825,736	719,450,627	-57,126,743	0.20
worsened up to 5 Cent	41,364,949	4,131,732,471	-103,941,597	1.17
worsened up to 1 Cent	25,916,402	4,929,264,915	-43,040,761	1.40
worsened up to 0.5 Cent	55,494,803	53,545,827,347	-76,503,894	15.21
executed at best	1,019,947,030	286,624,328,598	0	81.42
improved up to 0.5 Cent	2,367,857	1,339,207,221	1,684,150	0.38
improved up to 1 Cent	998,642	113,213,790	823,226	0.03
improved up to 5 Cent	1,342,175	89,196,618	1,621,695	0.03
improved up to 10 Cent	229,215	8,280,430	559,914	0.00
improved more than 10 Cent	505,810	6,874,522	1,778,696	0.00
Total worsened	148,489,363	63,858,224,575	-476,969,108	18.14
Total improved	5,443,699	1,556,772,581	6,467,680	0.44

Panel B: Tick-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 20 ticks	2,696,703	1,503,594,847	-179,632,226	0.43
worsened up to 20 ticks	1,725,905	1,076,730,517	-26,903,689	0.31
worsened up to 10 ticks	4,127,516	2,531,208,249	-32,818,895	0.72
worsened up to 5 ticks	44,814,501	21,796,625,193	-127,330,672	6.19
worsened up to 1 ticks	95,041,849	36,925,077,397	-105,730,811	10.49
executed at best	1,019,883,465	286,615,345,076	0	81.42
improved up to 1 ticks	2,185,958	624,757,767	1,656,818	0.18
improved up to 5 ticks	3,148,461	881,203,745	2,907,183	0.25
improved up to 10 ticks	52,723	27,857,760	344,804	0.01
improved up to 20 ticks	24,774	11,443,022	328,248	0.00
improved more than 20 ticks	31,760	11,506,799	1,230,612	0.00
Total worsened	148,406,474	63,833,236,203	-472,416,294	18.13
Total improved	5,443,676	1,556,769,093	6,467,665	0.44

Note: Table A.5 reports slippage cost by price deviation bucket in Cents, Panel A, and in ticks, Panel B, for Eurostoxx50 constituents. Filter (i) is applied which matches the full set of trades to latency-neutral EBBOs. We report the total number of #Trades, the total number of #Shares, the cumulative Euro slippage cost and the relative share in percent as the fraction of shares executed within each deviation bucket after full sample cumulation over all trading days and stocks. We acknowledge minor discrepancies in execution quality between Euro-based and tick-based execution which arise from missing tick size information for a handful of stock-days.

Table A.6:
Execution quality decomposition for REBBOs

Panel A: Euro-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 10 Cent	15,099,974	1,012,730,554	-413,348,982	0.29
worsened up to 10 Cent	17,503,510	1,022,085,547	-79,760,341	0.29
worsened up to 5 Cent	47,512,161	5,297,107,252	-133,124,431	1.50
worsened up to 1 Cent	29,565,770	5,720,583,405	-49,874,413	1.63
worsened up to 0.5 Cent	60,570,033	57,459,355,409	-83,911,150	16.32
executed at best	995,762,358	279,275,945,562	0	79.34
improved up to 0.5 Cent	2,937,696	1,793,288,145	2,377,655	0.51
improved up to 1 Cent	1,519,288	208,758,152	1,510,294	0.06
improved up to 5 Cent	2,115,479	183,706,973	3,589,712	0.05
improved up to 10 Cent	493,720	25,100,860	1,686,606	0.01
improved more than 10 Cent	794,201	19,475,388	5,810,183	0.01
Total worsened	170,251,448	70,511,862,167	-760,019,317	20.03
Total improved	7,860,384	2,230,329,518	14,974,450	0.63

Panel B: Tick-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 20 ticks	5,514,158	2,864,819,308	-401,126,234	0.81
worsened up to 20 ticks	2,599,620	1,433,899,427	-43,165,193	0.41
worsened up to 10 ticks	5,093,824	2,904,936,164	-41,936,245	0.83
worsened up to 5 ticks	50,608,933	23,604,661,850	-144,643,196	6.71
worsened up to 1 ticks	106,314,952	39,671,561,457	-118,609,239	11.27
executed at best	995,681,914	279,265,226,300	0	79.34
improved up to 1 ticks	3,071,489	841,581,761	2,649,136	0.24
improved up to 5 ticks	4,429,880	1,239,884,662	5,091,932	0.35
improved up to 10 ticks	140,879	67,351,380	1,167,529	0.02
improved up to 20 ticks	92,593	34,660,586	1,405,060	0.01
improved more than 20 ticks	125,167	46,827,363	4,660,093	0.01
Total worsened	170,131,487	70,479,878,206	-749,480,106	20.02
Total improved	7,860,008	2,230,305,752	14,973,749	0.63

Note: Table A.6 reports slippage cost by price deviation bucket in Cents, Panel A, and in ticks, Panel B, for Eurostoxx50 constituents. Filter (ii) is applied which matches the full set of trades to latency-adjusted REBBOs. We report the total number of #Trades, the total number of #Shares, the cumulative Euro slippage cost and the relative share in percent as the fraction of shares executed within each deviation bucket after full sample cumulation over all trading days and stocks. We acknowledge minor discrepancies in execution quality between Euro-based and tick-based execution which arise from missing tick size information for a handful of stock-days.

Table A.7:
Execution quality decomposition for local liquidity sufficient REBBOs

Panel A: Euro-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 10 Cent	13,531,126	791,822,767	-327,272,101	0.26
worsened up to 10 Cent	15,723,051	797,527,294	-62,463,024	0.26
worsened up to 5 Cent	42,626,903	4,070,273,797	-103,182,631	1.32
worsened up to 1 Cent	26,376,253	4,327,626,881	-38,236,870	1.40
worsened up to 0.5 Cent	52,611,334	42,654,534,462	-62,577,909	13.82
executed at best	950,341,651	254,445,998,987	0	82.46
improved up to 0.5 Cent	2,466,047	1,173,389,488	1,577,020	0.38
improved up to 1 Cent	1,389,958	155,394,519	1,124,864	0.05
improved up to 5 Cent	1,919,907	126,230,055	2,387,889	0.04
improved up to 10 Cent	459,291	17,223,293	1,148,730	0.01
improved more than 10 Cent	730,896	13,651,664	3,973,618	0.00
Total worsened	150,868,667	52,641,785,201	-593,732,536	17.06
Total improved	6,966,099	1,485,889,019	10,212,120	0.48

Panel B: Tick-based execution quality

	#Trades	#Shares	Euro slippage cost	Share in %
worsened more than 20 ticks	4,979,860	2,135,114,929	-316,086,982	0.69
worsened up to 20 ticks	2,305,690	989,177,311	-33,407,331	0.32
worsened up to 10 ticks	4,360,214	1,922,939,178	-30,447,414	0.62
worsened up to 5 ticks	43,743,073	16,631,807,427	-108,213,580	5.39
worsened up to 1 ticks	95,372,207	30,940,678,670	-97,138,212	10.03
executed at best	950,266,316	254,436,666,931	0	82.46
improved up to 1 ticks	2,742,575	575,002,961	2,133,161	0.19
improved up to 5 ticks	3,919,313	834,568,548	3,602,404	0.27
improved up to 10 ticks	115,889	33,156,512	627,948	0.01
improved up to 20 ticks	79,784	18,372,285	892,198	0.01
improved more than 20 ticks	108,239	24,776,983	2,956,090	0.01
Total worsened	150,761,044	52,619,717,515	-585,293,519	17.05
Total improved	6,965,800	1,485,877,289	10,211,802	0.48

Note: Table A.7 reports slippage cost by price deviation bucket in Cents, Panel A, and in ticks, Panel B, for Eurostoxx50 constituents. Filter (iii) is applied which subsets all trades to those executed against local liquidity sufficient REBBOs. We report the total number of #Trades, the total number of #Shares, the cumulative Euro slippage cost and the relative share in percent as the fraction of shares executed within each deviation bucket after full sample cumulation over all trading days and stocks. We acknowledge minor discrepancies in execution quality between Euro-based and tick-based execution which arise from missing tick size information for a handful of stock-days.

Table A.8:
Tick size table

Price level in Euro	Average daily number of transactions					
	<10	10–80	80–600	600–2,000	2,000–9,000	≥9,000
$0 \leq$ price < 0.1	0.0005	0.0002	0.0001	0.0001	0.0001	0.0001
$0.1 \leq$ price < 0.2	0.0010	0.0005	0.0002	0.0001	0.0001	0.0001
$0.2 \leq$ price < 0.5	0.0020	0.0010	0.0005	0.0002	0.0001	0.0001
$0.5 \leq$ price < 1	0.0050	0.0020	0.0010	0.0005	0.0002	0.0001
$1 \leq$ price < 2	0.0100	0.0050	0.0020	0.0010	0.0005	0.0002
$2 \leq$ price < 5	0.0200	0.0100	0.0050	0.0020	0.0010	0.0005
$5 \leq$ price < 10	0.0500	0.0200	0.0100	0.0050	0.0020	0.0010
$10 \leq$ price < 20	0.1000	0.0500	0.0200	0.0100	0.0050	0.0020
$20 \leq$ price < 50	0.2000	0.1000	0.0500	0.0200	0.0100	0.0050
$50 \leq$ price < 100	0.5000	0.2000	0.1000	0.0500	0.0200	0.0100
$100 \leq$ price < 200	1.0000	0.5000	0.2000	0.1000	0.0500	0.0200
$200 \leq$ price < 500	2.0000	1.0000	0.5000	0.2000	0.1000	0.0500
$500 \leq$ price $< 1{,}000$	5.0000	2.0000	1.0000	0.5000	0.2000	0.1000
$1{,}000 \leq$ price $< 2{,}000$	10.0000	5.0000	2.0000	1.0000	0.5000	0.2000
$2{,}000 \leq$ price $< 5{,}000$	20.0000	10.0000	5.0000	2.0000	1.0000	0.5000
$5{,}000 \leq$ price $< 10{,}000$	50.0000	20.0000	10.0000	5.0000	2.0000	1.0000
$10{,}000 \leq$ price $< 20{,}000$	100.0000	50.0000	20.0000	10.0000	5.0000	2.0000
$20{,}000 \leq$ price $< 50{,}000$	200.0000	100.0000	50.0000	20.0000	10.0000	5.0000
price $\geq 50{,}000$	500.0000	200.0000	100.0000	50.0000	20.0000	10.0000

Note: Table A.8 shows that tick sizes depend on both the price level and the average daily number of transactions. The smallest permissible tick size in the EU equity tick-size regime is €0.0001 (0.01 Euro-Cent), applying to instruments priced below €0.10 that trade in the highest liquidity band with an average daily number of transactions exceeding 9,000. Source: Commission Delegated Regulation (EU) 2017/588, Annex I. (EU) 2017/588 (RTS 11).

A.2 Supplementary figures

Figure A.1:
Consolidation from local to global order books

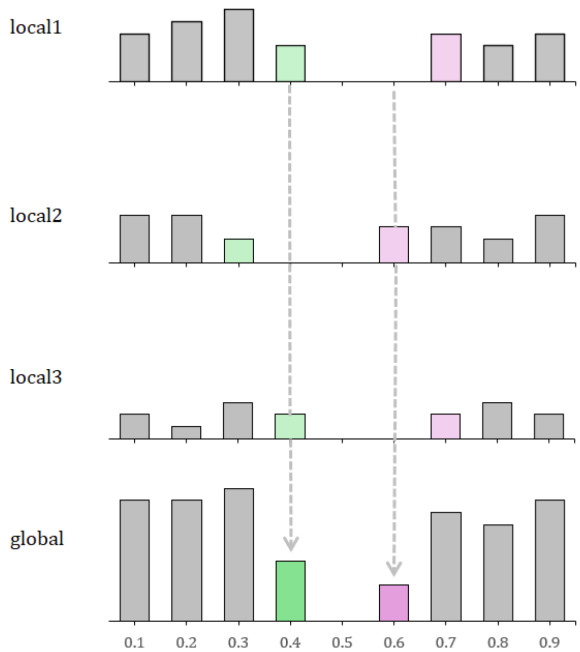

Note: Figure A.1 illustrates the consolidation of local order books to a global order book. Each local order book reports the available bid and ask depth at specific price levels. In this stylized example, local order books are tick-size constrained to a minimum quoted spread of 0.3 Euro. By aligning price levels across venues and summing the corresponding depths, the consolidated order book reflects the aggregate liquidity available in the market. The global spread accounts to 0.2 Euro as not all venues post synchronous best prices and thereby undercuts the regulatory enforced minimum tick size of 0.3 Euro on local markets. Venues quoting best global prices do so one-sided. When adjusting for inter-datacenter transmission latency, each local order book is assigned a venue-specific time lag to its quotes before aggregation, thereby accounting for geographic dispersion across venues.

A.3 Pseudo-code (R)EBBO reconstruction

```
Input: V: Set of all quoting venues
Input: Q_v: Set of BBO updates for each v ∈ V
Data: φ(e_i, e_j): Transmission latency between venue i and j

speed fiber = 204,190 ;                          // transmission speed in fiber
κ = 1.3 ;                                        // transmission overhead factor

foreach e_i ∈ V do                               // for every consolidating venue
    foreach e_j ∈ V do                           // for every quoting venue
        latency ← φ(e_i, e_j, speed fiber) · κ ; // obtain transmission latency
        foreach q ∈ Q_{e_j} do
            q.timestamp ← q.timestamp + latency ;   // lag venue quotes
        end
    end

    Q ← ⋃_{v∈V} Q_v ;                            // merge all BBO updates
    Q ← sort(Q by timestamp) ;                   // sort updates chronologically

    foreach t ∈ Q do                             // derive best global prices
        best_bid_t ← max(bids_t)
        best_ask_t ← min(asks_t)
    end
end
```

Note: Appendix A.3 illustrates the pseudo-code for (R)EBBO reconstruction per equity. We reconstruct the venue-specific REBBO after lagging all venue timestamps based on the venue-pair specific transmission latencies in optic-fiber cable. $\varphi(e_i, e_j)$ represents the transmission latency between exchange i and j. When computing the latency-neutral EBBO, $\varphi(e_i, e_j)$ is essentially set to zero.