Navigating Stormier Seas? Liquidity Resilience Across Asset Classes and Time

Nihad Aliyev[†], Matteo Aquilina[¶†], Khaladdin Rzayev[‡ * ♭], and Sonya Zhu[¶]

[¶]Bank for International Settlements

[†]Macquarie University

[‡]The University of Edinburgh

[*]Koç University

[♭] Systemic Risk Centre, London School of Economics

Abstract

Paradoxically, as financial markets have grown more liquid, they appear to have become increasingly fragile. Our study of stocks, foreign exchange (FX), and government bonds across the US, Europe, and Japan, using a 25-year sample of high-frequency data, shows a significant decline in both the average and standard deviation of bid-ask spreads across all asset classes. However, there is an increase in skewness and kurtosis in equity and bond markets, indicating more frequent episodes of illiquidity, making them more fragile. In contrast, FX markets do not show significant increases in higher moments of illiquidity. We also demonstrate that structural breaks in spreads are correlated with macroeconomic shocks and shifts in market conditions. Furthermore, the rise of algorithmic and high-frequency trading (AT) and market fragmentation is associated with narrower spreads and increased skewness in equity markets, with similar effects for AT observed in FX markets but not in bond spreads.

JEL Classification: G10, G12, G14

Keywords: liquidity; market resiliency; high frequency trading; fragmentation

[0]Email address: Nihad Aliyev(nihad.aliyev@uts.edu.au), Matteo Aquilina (matteo.aquilina@bis.org), Khaladdin Rzayev (khaladdin.rzayev@ed.ac.uk), and Sonya Zhu (sonya.zhu@bis.org).

We thank Denis Gorea, Marco Lombardi, Benoit Mojon, Andreas Schrimpf, Vlad Sushko and Nicholas Zarra for helpful comments and suggestions. We also thank Alberto Americo and Sjur Nielsen for outstanding research assistance. The views expressed here are those of the authors, and not necessarily those of the Bank for International Settlements.

1 Introduction

In Aesop's famous fable, a boy's repeated false cries for help bring villagers running to fend off a wolf that never appears—until one day, a real wolf arrives with devastating consequences. Similarly, in financial markets, there are episodes where market liquidity vanishes abruptly, leaving traders unable to execute orders quickly or at prevailing prices, with equally severe outcomes. These episodes can be especially problematic if market participants have become accustomed to consistently high levels of liquidity, thereby amplifying the shock of its sudden absence.

Practitioners and regulators have shown significant interest in liquidity in recent years. A notable example is the response to a single, substantial episode of illiquidity—the *Flash Crash* in U.S. stocks in May 2010—after which the newly established CFTC-SEC Advisory Committee on Emerging Regulatory Issues issued a report with extensive recommendations for changes to the underlying market microstructure, driven by concerns that a loss of confidence could undermine market integrity and stability. Menkveld and Yueshen (2018) highlight that the Investment Company Institute reported five months of equity outflows following this episode. Importantly, this focus on liquidity was not limited to the U.S. stock market alone. For example, the Markets Committee of the Bank for International Settlements (BIS) published reports examining liquidity issues in global foreign exchange and fixed income markets (BIS Markets Committee, 2011, 2016). Additionally, the European Commission undertook substantial work on corporate bond markets, producing studies on the drivers of liquidity (European Commission, 2017).

However, the focus of much analysis has primarily been on the average level of liquidity rather than its sudden disappearance. As new regulations introduced after the 2008 financial crisis began to take effect, some market participants raised concerns that these rules might negatively impact intermediaries' ability to facilitate transactions, thereby potentially reducing the average level of liquidity. While we acknowledge that the average level of liquidity is important, we argue that its resilience and availability at all times are equally, if not more, crucial. Many asset classes, such as real estate, operate effectively even with low average liquidity due to their specific characteristics: participants know that it may take months or even years to sell a building and take that into account in their decisions. However, if market participants become accustomed to easy trading, the sudden disappearance of liquidity could compel them to take actions that may have destabilizing effects. In his influential book, Persaud (2003) highlights that such phenomena, which he terms liquidity *black holes*, can "destroy companies, cause significant economic contraction, bring down governments, rip the social fabric and steer capital away from certain markets more permanently.[1]"

To address this gap in the literature, instead of focusing solely on the average level of liquidity (i.e., the first moment of liquidity), we analyse the entire distribution of liquidity in arguably the three most important asset classes: stocks, FX and government bonds. We construct a comprehensive dataset of the bid-ask spread that spans the US, European and Japanese markets, using high-frequency data that we aggregate into monthly and yearly values. We first document several interesting patterns in the evolution of the distribution of liquidity. We then concentrate our attention on the first and third moments, as these convey most of the information we are interested in. We statistically identify structural breaks in the time series and explore their potential causes. To do so, we integrate traditional analytical approaches with the application of large language models to uncover potential explanatory factors. We then conduct a regression analysis across all asset classes, investigating the correlations between microstructural and macroeconomic variables and both the mean and skewness of the bid-ask spread. Finally, we perform a simulation exercise to illustrate the cost implications for traders resulting from the increasing fragility of liquidity in the current landscape compared to the beginning of our study period. Our main results are as follows:

- The average bid-ask spread has declined across all asset classes over the past 25 years. While the absolute and relative magnitude of these changes varies by asset class and geographical region, the overall decline is notable.

- The standard deviation of the bid-ask spread shows patterns that closely mirror those of the mean spread, with a near-universal reduction observed across asset classes and geographies.

- The higher moments of the distribution (skewness and kurtosis) present a different pattern. In stock and government bond markets, these higher moments have increased, indicating more frequent episodes of substantial illiquidity. In FX markets, skewness generally follows an inverted-U pattern, while kurtosis fluctuates without showing a significant trend.

- Statistical tests designed to detect structural changes in time series indicate that such breaks are associated with various factors across all asset classes, including broad economic and geopolitical events as well as market-specific developments.

[1]See page xv.

Overall, breaks in the mean tend to occur in proximity to macroeconomic shocks, while skewness appears more responsive to changes in the underlying market structure and, to a lesser extent, regulatory changes.

- Regression analysis shows that two major shifts in equity markets—the rise of algorithmic trading/high-frequency trading (AT)[2] and increased market fragmentation—are associated with lower average spreads but higher skewness. These developments may thus contribute to making markets more fragile and less resilient. A similar result is observed for AT in the FX market. Consistent with the notion that AT is less prevalent in bond markets, we do not find any significant association between our AT measure and the bid-ask spread moments in this market.[3]

- Instrument-level volatility, measured as the absolute value of midprice return, is a significant determinant of bid-ask spread moments across all markets. Specifically, higher volatility is strongly associated with higher average bid-ask spreads. The relationship between volatility and skewness is also positive, although less statistically significant. At the market level, the VIX index is a significant determinant of bid-ask spread moments in equity markets, while the MOVE index plays a similar role in bond markets.

- Simulation results demonstrate that changes in the skewness of the bid-ask spread have a direct impact on trading profitability and costs. Specifically, increasing skewness by about 50%—a change similar to that observed in stock markets—while keeping the mean and standard deviation constant, reduces trading profitability by approximately 7.2%.

In a nutshell, our results reveal a fascinating trend across many financial markets worldwide: trading has become easier than ever, with spreads narrowing significantly. However, the turbulent waves of illiquidity and the disruptions they can cause to other markets have become much more frequent. It is akin to a mostly serene sea that can suddenly give way to powerful storms.

The rest of the paper is structured as follows. Section 2 introduces our liquidity measure, provides details on the data, and presents summary statistics for our empirical analysis. Section 3 describes the high-level evolution of the distribution of the bid-ask spread across asset classes and geographies. We then move to analysing statistically

[2]The acronym "AT" refers to both algorithmic trading and high-frequency trading.

[3]We cannot measure fragmentation for market segments other than equities.

the breaks in the time series of the mean and skewness in Section 4 and analyse the determinants of these two moments in Section 5. Section 6 reports the results of the simulation exercise, and Section 7 concludes.

2 Liquidity measure, data and summary statistics

To compare the distribution of market liquidity across different asset classes and jurisdictions, an ideal liquidity measure should possess a number of features. First, it should capture the ease of trading. In other words, a liquidity measure that captures the cost of trading is more appropriate for our purpose than a liquidity measure that captures the volume of trading. Second, it should be widely available and computationally easy to calculate given the breadth of markets that we aim to include in the analysis. This rules out model-based liquidity measures such as price impact (Kyle (1985)) and Roll's implicit liquidity measure (Roll (1984)). Third, it should be comparable across asset classes and time.

To this end, our analysis relies on market liquidity measured by relative quoted spread. For all markets, we construct relative quoted spreads based on intraday snapshots of the market. Our relative quoted spread for instrument i and at time t is defined as the difference between the ask and bid price, divided by the mid-price:

$$s_{it} = \frac{Ask_{it} - Bid_{it}}{MidPrice_{it}}. \tag{1}$$

Our main analysis relies on the first to fourth moments of relative quoted spread at the instrument and month level. For all instruments we first calculate all liquidity moments for each instrument and on each trading day. Afterwards, we take the average of these moments for each instrument by month.

2.1 Data sources

Our source for the relative quoted spread data is the LSEG *Tick History* dataset. Tick History provides historical information at different levels of aggregation (from tick by tick to daily) for a number of asset classes. Its coverage goes as far back as 1996 for many times series and is widely used by industry practitioners. The dataset consists of recorded trades and quotes from a number of real-time feeds across more than 500 trading venues including all types of participants (i.e., not just dealers). In the literature, it has been used extensively in the analysis of the liquidity of stock markets,[4] but also across bond (Sakiyama and Kobayashi (2018)) and FX ones (Krohn and Sushko (2022)).

For the government bond and FX markets, we obtain bid and ask quotes at 1-minute snapshots. For the equity market, we start by calculating the liquidity of individual stocks traded in major economies. Because the number of stocks is very large (about 10 thousand), we rely on 5-minute snapshots for computational reasons.

On the basis of these intraday data, we calculate daily moments (mean, standard deviation, skewness and kurtosis) of the relative bid-ask spread and then average these either monthly or yearly to investigate longer-term trends in the distribution of liquidity. We have also implemented a number of steps to clean the data. First, we only use local trading hours (9:30 to 16:00) and remove weekends. We then eliminate data when the order book is crossed (i.e., ask price is lower than or equal to the bid price). We also check for significant data gaps. If the month does not contain at least 9 days of data it is removed from the sample as are years for which more than 2 months of data are missing. We then only keep the time series for which at least 5 years of data. Finally, to remove the impact of extreme values[5] on our estimated moments, we trim the high-frequency bid-ask spread at the 99% percentile.

Our main goal is to characterise the distribution of liquidity - with a specific focus on its higher moments - across global financial markets. Hence, we focus on three high-level geographical areas: America, Asia and Europe. America and Asia are represented by the United States and Japan in the equity and bond segments, given their out-sized importance in such markets. Europe is represented by Germany, France and the UK in the equity market and by Germany, Italy and the UK in the government bond market to reflect the relative importance of French equities and Italian government debt, respectively. In the FX market, there is not an equivalent concept of an instrument based in a specific jurisdiction. We nonetheless use a similar approach and focus on the exchange rate of the US dollar against the Euro, the British Pound and the Japanese Yen, respectively.

To carry out the analysis in Section 5, we complement Tick History data with information coming from other sources. In particular we source FX volatility indexes from JP Morgan and the TED spread and the VIX from FRED. In both the stock market and the FX market analysis, we control for AT, defined as the ratio between a number of trades and quotes obtained from LSEG. For the FX market analysis, because the trading vol-

[4] See Aquilina et al. (2024), Werner et al. (2023), Ibikunle et al. (2021), Comerton-Forde et al. (2019), and Degryse et al. (2015).

[5] These are data errors rather than statistical outliers.

ume of the spot market is not available, we use the trading volume of FX futures instead. In the stock market analysis, we also control for fragmentation, defined as the inverse Herfindahl–Hirschman (HHI) index based on volumes across different trading venues. For the US, we use the trading volume data for all trading venues provided by LSEG. For Europe, we collect the trading volume data from Xetra Germany, SIX Swiss Exchange, London Stock Exchange, Euronext Paris, CBOE (both BATS and Chi-X), and Euronext Amsterdam. For Japanese equities, we source the data from the Osaka Stock Exchange, Tokyo Stock Exchange, Nagoya Stock Exchange, Fukuoka Stock Exchange, and Sapporo Stock Exchange.

Our final sample consists of more than 2 billion observations in high-frequency bid-ask spread data, a significantly higher number compared to those used in the relevant literature. While the use of high-frequency data requires intensive computational effort, it is often considered more precise for evaluating trading costs than end-of-day liquidity measures. The literature generally acknowledges that *"low-frequency measures should be used only when high-frequency data are not available"* (Vayanos and Wang, 2013). Hence, we use high-frequency intraday data in our study. However, high-frequency data can be more susceptible to data errors, so we cross-check our high-frequency measures against their corresponding low-frequency counterparts. We also compare the time-series evolution of the bid-ask spread with findings from other studies in the literature. Generally, the spreads based on high-frequency and low-frequency data are highly correlated and capture similar variations, with time-series trends consistent with those documented in the literature.

The only exception is the U.S. Treasury bond data, where we detected a potential issue between late 2004 and early 2009. During this period, spreads exhibit an unusual pattern, first jumping up and then down within a single month (see Figure 3.3), a behavior not observed in other sources focusing on dealer-to-dealer transactions (Fleming and Ruela, 2020). We investigated this issue with the data provider, who confirmed that the data were delivered correctly and were not corrupted. They specifically affirmed that the data matches those they receive from the trading venues. We suspect, though cannot confirm, that this anomaly may be due to the migration of U.S. bond trading from voice trading to electronic platforms like BrokerTec and eSpeed by early 2005, and that the data LSEG receives may not fully capture this transition.[6] Given the assurances from LSEG, the fact that our study encompasses transactions beyond those involving only dealers, and the

[6]See Mizrach and Neely (2006) and Fleming et al. (2018) for detailed descriptions of the shift from voice to electronic trading in the U.S. Treasury market.

likelihood that market participants traded based on these quotes, we decided to retain the data in our sample. However, in the regression analysis in Section 5, we repeat the analysis excluding this period to ensure the robustness of our results.

2.2 Summary statistics

We commence our analysis by describing the distribution of the bid-ask spread across our entire sample. Table 1 reports the average moments by asset type and by country. For each stock on each trading day, we first calculate the moments of its bid-ask spread distribution. Then, we group stocks into terciles based on their average market capitalization for each year. In the U.S. market, all four levels of moments exhibit a clear correlation with market capitalization. The mean and standard deviation of the bid-ask spread decrease as market capitalization increases. During our sample period, the average bid-ask spread for small stocks is 148 basis points (bps), more than three times higher than that of medium-sized stocks (41 bps) and eight times higher than that of large stocks (17 bps). Conversely, skewness and kurtosis decrease with market capitalization. This relationship between market capitalization and the moments of the bid-ask spread is not unique to the U.S.; it holds across all four countries in our sample. Across all jurisdictions, the bid-ask spread of stocks is generally positively skewed, highlighting the prevalence of highly illiquid periods, and has excessive kurtosis.

In the government bond market, both the mean and standard deviation of the bid-ask spread increase with bond maturity across all geographies. For U.S. government bonds, the average bid-ask spread is 1.81 bps for the two-year bond, 2.28 bps for the five-year bond, and 3.03 bps for the ten-year bond. Among the five countries we study, U.S. government bonds are the most liquid, while Italian government bonds are the least liquid. Government bonds with shorter maturities not only have a narrower average bid-ask spread but also exhibit less volatility in the spread. Generally, the skewness and kurtosis of two-year government bonds are higher than those of ten-year government bonds, with the exception of German bonds.

In the FX market, the costs of trading GBP, Yen, or Euro against the U.S. dollar are quite similar, averaging around 3 bps. The average bid-ask spread in the FX market is also comparable to that of government bonds from the same jurisdictions. In terms of liquidity volatility, GBP/USD is the least volatile, while JPY/USD is the most volatile, though the difference in their standard deviations is only about half a basis point. The standard deviation of the bid-ask spread in the FX market is roughly one-third of the mean spread, similar to the government bond market, suggesting that both markets exhibit relatively low liquidity risk. Compared to other asset classes, the skewness and kurtosis of the bid-ask spread in FX markets are much lower and close to zero.

3 The evolution of the distribution of liquidity

In this section, we examine the evolution of liquidity distribution. We first calculate the four daily moments of high-frequency bid-ask spreads and then use yearly averages of these daily moments to understand the long-term trends.

3.1 Equity

We investigate liquidity moments in the equity markets of the U.S., France, Germany, the U.K., and Japan from 1996 to 2023. These markets are among the most liquid in the world and collectively represent approximately 57% of the total market capitalization of global stock markets.[7]

(Figure 1) reports the average bid-ask spread across all countries and size groups. Aggregating across regions, the bid-ask spread decreased from 60 bps in 1996 to 13 bps in 2023, from 103 bps to 34 bps for medium stocks, and from 184 bps to 111 bps for small stocks. A similar trend was observed in the standard deviation of the bid-ask spread, which averaged around 34 bps in 1996 and decreased to 19 bps by 2023, representing an 84% reduction.

However, the bid-ask spread in the equity market is positively skewed, regardless of country and firm size. Before the global financial crisis, there was a universal upward trend in skewness across the U.S., Japan, and Europe. After the crisis, skewness in the U.S. equity market increased again, while in Japan and Europe, the pattern differed. In Japan, the level of skewness remained relatively stable since 2010, whereas in Europe, it started to decline during the same period. Skewness is also positively related to firm size, a pattern that generally holds across all regions. However, the increase in skewness over time is relatively smaller (larger) in large (small)-cap stocks.

The trend in the kurtosis of bid-ask spreads mirrors that of skewness, with even more pronounced increases. Specifically, the kurtosis of the bid-ask spread increased from 0.23 in 1996 to 6.13 in 2023 for small stocks, from 0.52 to 7.65 for medium stocks, and from 0.83 to 7.72 for large stocks.

[7]See, for example, https://data.worldbank.org/indicator/CM.MKT.LCAP.CD.

3.2 Foreign exchange

Figure 2 reports the distribution of bid-ask spread for trading the Euro, Japanese Yen and British Pound against the U.S. dollar. These are the three most actively traded currencies in the FX spot market. In April 2022, the average daily turnover of trading these currencies against the U.S. dollar accounted for nearly 50% of the total turnover in the foreign exchange spot market.((McGuire et al., 2024)).

The average bid-ask spread in the FX market has approximately halved from the mid-1990s to the present day. For example, the cost of trading Japanese yen in 1996 was around 6 bps, while by the mid-2000s, it declined to about 3 bps. This implies that a transaction of 100,000 U.S. dollars incurs a spread of about 60 dollars in 1996 and about 30 dollars in 2005. In the years following the global financial crisis, bid-ask spreads increased by around one-third across all currency pairs. After the financial crisis, the long-term downward trend in bid-ask spreads stopped.

The standard deviation of the bid-ask spread follows a similar trend as the mean spread, exhibiting a persistent decline across all currencies until the global financial crisis. Since 2015, however, the standard deviation has been gradually increasing, though this rise is less pronounced for the USD/JPY pair.

The skewness of the bid-ask spread has shown an inverted-U pattern, closely co-moving across all countries before 2014 with a generally increasing trend. After 2015, the skewness of the GBP and EUR began to decline. While the Japanese yen initially followed a similar downward trend, its skewness spiked sharply in 2019 and has since plateaued.

The excess kurtosis of the bid-ask spread for foreign exchanges has generally remained stable over time but experienced sharp spikes around the global financial crisis, gradually declining into negative territory afterward. However, following the outbreak of the COVID-19 pandemic, the kurtosis of various currencies began to increase again.

3.3 Government bonds

Figure 3 illustrates the evolution of the bid-ask spread distribution in government bonds. The moments of the bid-ask spread across different maturities closely co-move with one another. The average bid-ask spread generally declined across the U.S., Japan, and Europe.[8] The data issue associated with U.S. government bonds does not appear in

[8] We take the average across British, German, and Italian government bonds

other countries. Among European countries, our unreported results show that the bid-ask spread for Italian bonds experienced sharp spikes during the Global Financial Crisis and the European sovereign debt crisis. In contrast, the bid-ask spread for UK and German government bonds has steadily improved over the past two decades, although these changes are less pronounced than those in the FX market. Regardless of the region, two-year government bonds tend to be more liquid than five- and ten-year government bonds.

The standard deviation of the bid-ask spread for the US and European government bonds follows a long-term downward trend. However, the standard deviation for Japanese government bonds has been increasing since 2010. Starting from 2016, the bid-ask spread for the US government bonds also started to be more volatile. The standard deviation of the bid-ask spread largely reflects the liquidity risk of the underlying asset. When comparing bonds of the same maturity, the US (Japanese) government bonds are least (most) exposed to liquidity risk among the countries that we study. Additionally, Japanese government bonds have a wider spread across different maturities compared to U.S. Treasuries. In the U.S., the standard deviation of the bid-ask spread for the ten-year bond is around 0.63 bps, nearly twice that of the two-year bond (0.38 bps). However, in Japan, the liquidity risk of the ten-year government bond (2.08 bps) is approximately five times higher than that of the two-year bond (0.52 bps).

The skewness of the bid-ask spread across regions and maturities appears to follow a common time-series trend. Before the global financial crisis, skewness remained stable, fluctuating mostly around zero. However, following the crisis, there was a universal upward trend in skewness, regardless of region or maturity. This trend is particularly pronounced in the U.S. and Japan. Recently, the skewness of the bid-ask spread has started to decline following the pandemic shock, though it remains in positive territory.

Similar to skewness, the kurtosis of the bid-ask spread increased rapidly from 2015 to 2018 across regions and maturities, only starting to decline recently. While the mean and standard deviation of the bid-ask spread show a clear relationship with maturities, the levels of skewness and kurtosis remain similar across different maturities. This difference suggests that the underlying drivers of the lower and higher moments of the bid-ask spread might be different, as the lower moments demonstrate a clear term structure, whereas the higher moments do not.

4 Step changes in the distribution of liquidity and their potential drivers

In the previous section, we documented the behavior of the bid-ask spread distribution across various asset classes and geographical areas. In this section, we dig deeper by attempting to statistically identify structural changes in the moments of these distributions and hypothesize about their potential drivers. This second step is inherently speculative, as there are numerous possible reasons for breaks in the time series, including market characteristics, competition dynamics in liquidity provision, regulatory changes, macroeconomic factors, and more. Nevertheless, we believe it is valuable to assess whether these breaks coincide with other market changes, as this could illuminate potential avenues for future research.

We focus our analysis on the mean and skewness of the distribution for two main reasons: first, the correlation between the mean and standard deviation, as well as between skewness and kurtosis, is high, offering limited additional insight from examining all moments. Second, the positive skewness of liquidity is particularly relevant to our study, more so than kurtosis. The distribution of the spread is truncated on the left, as the spread cannot be negative, and if the distribution is particularly *fat* on the left tail—indicating many instances of especially narrow spreads—this would be beneficial for market participants.

4.1 Estimating the structural breaks

To determine structural breaks in the mean and skewness of the relative bid-ask spread, we rely on the methodology developed by Bai and Perron (1998, 2003). Their approach identifies multiple structural breaks that occur at unknown dates within linear regression models estimated by least squares. This methodology is highly flexible, accommodating both scenarios where the number of structural breaks is known in advance and those where it is unknown, as in our case. If necessary, tests can also be performed on the coefficients of a subset of regressors. The general model is expressed as:

$$Y_t = X_t'\beta + Z_t'\delta_j + \varepsilon_t \tag{2}$$

where Xt is a vector of regressors with fixed coefficients and Zt a vector of regressors with coefficients that are subject to change. The break dates are $t = Tj$ for $j = 1, ..., m$ and T is the entire sample size.

The model tests the null hypothesis of the coefficients remaining constant against the alternative hypothesis that the coefficients change over time. The procedure then compares different combinations of partitions of the data to minimise the global residual sum of squares. In a nutshell, it compares a partition of $m-1$ breaks to a partition of m breaks and selects the partition with the overall lower residual sum of squares. For our purposes, we are interested in estimating a mean-shift model for the mean and skewness of the distribution of the spread. Hence, the regression model only includes a potentially shifting constant. Using monthly data, we estimate such mean-shift models for the three geographical units (Europe, Japan and US) and for each sub-asset class separately (FX; large, medium and small cap stocks; 2y, 5y and 10y government bonds) for a total of 21 models.[9]

The results of our estimation are visually summarized in Tables 2 and 3. In the table, we report the direction of the break, i.e. whether the jump in the time series is upwards or downwards, together with the period in which the estimated jump took place. While the model reports the estimated month of the break, as there is uncertainty over its exact timeframe and also for readability purposes, we divide each year into two halves and report the half in which the break is identified. The top panel of the table reports the breaks identified in the mean of the spread, while the bottom panel reports the breaks identified in the skewness. Overall, the breaks identified statistically using the monthly data are aligned with the general path that can be gauged by looking at the yearly graphs presented in the previous section.

There are downward shifts in the mean spread in the early 2000s for large cap stocks and for FX, and additional ones are identified across all asset classes in the mid-2000s and in the early 2010s. Upward shifts in spreads for small cap equities and some government bonds are identified in the late 2000s and across asset classes from 2015 onwards. A particularly interesting pattern that emerges from these tests is the frequent occurrence of upward shifts in skewness shortly after a downward break in the mean spread. For example, in the early 2000s, skewness increases in the stock and foreign exchange markets, while in the mid to late 2010s, similar increases are observed in government bond markets and certain equity markets. This pattern aligns with the theoretical prediction by Roll and Subrahmanyam (2010), who attribute the increase in skewness to more intense competition among market makers, which reduces cross-subsidization across periods.

[9]We use the Yao (1988) Bayesian information criterion and a 95% confidence interval in our estimates.

4.2 Assessing the potential causes of the breaks

An obvious question that arises is whether it is possible to identify potential drivers of the observed breaks. This is by no means an easy task for two main reasons. First, multiple factors can contribute to shifts in the distribution of spreads. Second, empirically disentangling causality is challenging, especially when many breaks are estimated over a relatively long period. However, in real-world scenarios, researchers often rely on observational data to study complex situations as they unfold, capturing a broad range of issues that may influence outcomes. While this type of evidence is imperfect and not definitive, it remains valuable for generating hypotheses and highlighting potential causal relationships that can be studied more rigorously.

To gather such evidence, we conduct a comprehensive literature review to identify potential explanatory factors. Our search includes possible changes in regulation, market microstructure, and macroeconomic shocks that occur around the time of the identified breaks in different markets. To complement our manual search, we also rely on ChatGPT to explore potential reasons why the statistical tests identify these breaks.

Before discussing our findings in more detail, we should address one of the most significant structural changes that characterised financial markets in our period of analysis, namely their *electronification*. This change was accompanied by an increase in the fragmentation of trading and by the birth and subsequent expansion of algorithmic and high frequency trading that now represent a substantial share of liquidity provision.[10] Ideally, one would look at the effect of the entry of HFTs on liquidity and its skewness by conducting differences-in-differences analysis before and after their entry in specific markets.[11] However, the HFT industry is notoriously opaque, and as the only publicly listed HFT company is Virtu Financial, it is particularly difficult to gather precise information on when such firms entered different market segments that could be used in such an analysis.[12] However, approximate estimates of the penetration of HFTs by asset class are available (see, for instance Goldman Sachs (2018), or Chaboud et al. (2014) and (BIS Markets Committee, 2011) for specific estimates in the FX market) and indeed there seems to be a broad correlation in all asset classes between the reduction in spreads and the overall importance of algorithmic trading. However, as noted above, such a reduction

[10]The literature on the impact of HFTs on market quality and liquidity is now vast. See Menkveld (2016) for a review.

[11]There are a number of papers that attempt to do exactly that, but focus on average levels of liquidity rather than its skewness.

[12]In the next section, we attempt to look at the effect of HFTs by using either quote-to-trade ratios, or the overall number of trades as a proxy for their involvement in markets

in spreads is often accompanied by an increase in its skewness. The exception is the foreign exchange market, which now exhibits very limited skewness notwithstanding the very significant role played by HFTs as liquidity providers.

Many of the breaks in the mean and skewness of the bid-ask spread appear to coincide with significant changes in financial markets. Several factors seem particularly relevant: changes in the market's microstructure, regulatory shifts, and macroeconomic shocks. In terms of microstructural changes, in addition to the broad correlation with the increasing relevance of HFTs mentioned earlier, the decimalization[13] of stock prices and the introduction of Autoquote[14] in the early 2000s coincided with the reduction in spreads and the increase in skewness. Indeed, Hendershott et al. (2011) use the introduction of Autoquote as an instrument for HFT activity to highlight the positive impact HFTs had on liquidity.

Another market change that seems to have affected both average spreads and skewness is the introduction of the Euro, particularly the introduction of Euro notes and coins in 2002. From January 1999 to December 2001, the Euro was an *invisible* currency, serving as the unit of account in 12 countries and being used in electronic payments, but without any physical coins or notes. This changed in January 2002. All the currency pairs in our sample showed declines in spreads and increases in skewness after the introduction of the Euro, with significant jumps for the EUR/USD pair occurring after the introduction of notes and coins.

Moving onto regulatory changes, MiFID I and MiFID II in Europe, REG NMS in the US, the Financial Instruments and Exchange Act in Japan and the period associated with the development of the global code in foreign exchange markets are associated with many of the identified breaks. In the US equity market, downward breaks in spreads pre-date the introduction of REG NMS, but the increase in skewness for mid-cap equities takes place at approximately the same time. This is potentially an indication that the additional fragmentation that was brought about by it may have resulted in an increase in skewness. In Europe, it is difficult to disentangle the effects of MiFID I as its implementation coincides with the onset of the global financial crisis. However, spreads in bond markets increased after the introduction of MiFID II, as did the skewness - albeit only in some

[13]The decimalization of pricing was ultimately driven by a change introduced by the SEC, but in the preceding years, several exchanges began planning the move and introduced pilot programs. Hence, we categorize decimalization as a market change rather than a regulatory one.

[14]Autoquote was software that automatically disseminated all changes in the best quotes to market participants. Previously, market makers had to manually update the best quotes. This innovation allowed algorithmic traders to receive information much more quickly. See Abergel et al. (2012) for details.

segments. MiFID II expanded some of the provisions of MiFID I to markets other than equities. In FX, the global code was discussed for a few years, with the principles initially published in 2016 and the first version of the code in May 2017. This period is associated with upward breaks in mean spread (in late 2015) and downward shifts in skewness.

The last category of events associated with the identified breaks in the series includes macroeconomic shocks and the interventions by authorities in response to them. Beginning with the financial crisis, bond market spreads increased in Europe but not in the U.S. and Japan, while skewness rose in Europe and Japan but remained unaffected in the U.S. In equity markets, the financial crisis is linked to increased spreads in small-cap equities across all regions and in mid-cap equities in Japan. Upward breaks in skewness are observed in U.S. large caps and Japanese small caps, while downward breaks are identified in European equities. As noted earlier, since MiFID I came into force in late 2007, it is challenging to separate its effects from those of the financial crisis.

Mario Draghi's *"whatever it takes"* speech, which effectively marked the end of the European sovereign bond crisis, was followed by a reduction in spreads in European bond markets, though no breaks in skewness were observed. In Japan, the advent of Abenomics in 2012 is associated with downward shifts in spreads in equity markets, but there is no noticeable effect on skewness or in bond markets.

In summary, a wide range of factors can be associated with the identified jumps in average bid-ask spreads and their skewness, ranging from broad economic and geopolitical events to market-specific changes and developments. Overall, it is easier to link the identified breaks in the mean of the bid-ask spread to macroeconomic shocks, while skewness appears to respond more strongly to changes in the underlying market structure and, to a somewhat lesser extent, to regulatory changes.

5 Regression analysis of mean and skewness

So far, we have discussed the time-series evolution of four moments of the bid-ask spread and identified periods when breaks occur in the evolution of the mean and skewness across different asset classes. In this section, we explore the potential determinants of these moments by estimating panel regressions. For this, we focus solely on the mean and skewness, similar to the break analysis, due to the high correlation observed between the mean and standard deviation, as well as between the skewness and kurtosis, which suggests that their variations are closely related.

We start our analysis by focusing on the equity markets and estimating the following regression model:

$$Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Frag_{i,m} + \beta_3 Volume_{i,m} + \beta_4 MCap_{i,m} + \beta_5 Volatility_{i,m} + \beta_6 VIX_m + \beta_7 TED_m + \varepsilon_{i,m} \quad (3)$$

$$Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Frag_{i,m} + \beta_3 Volume_{i,m} + \beta_4 MCap_{i,m} + \beta_5 Volatility_{i,m} + \beta_6 VIX_m + \beta_7 TED_m + \beta_8 Mean_{i,m} + \varepsilon_{i,m} \quad (4)$$

where $Mean_{i,m+1}$ and $Skewness_{i,m+1}$ are the mean and skewness of the equity bid-ask spread for stock i and month $m+1$. Across all specifications, we employ the first lag of independent variables to reduce the endogeneity concern. $Algo_{i,m}$ is the proxy for algorithmic trading (AT), calculated as the number of quotes divided by the number of trades for stock i and month m (Hendershott et al., 2011). The number of trades and quotes for each stock and hour are sourced from LSEG. The monthly average of the hourly ratio of the number of quotes to the number of trades is then used as our AT proxy. The second market quality characteristic, market fragmentation, is denoted by $Frag_{i,m}$. To calculate this measure, we collect the trading volume for each stock i on day d across different trading venues from LSEG. $Frag_{i,m}$ is then computed as the monthly average of the daily $\frac{1}{HHI}$ index, where the HHI index is the sum of the squares of the fraction of shares for stock i traded on a venue on a given day. For the US, we employ trading volume for all trading venues provided by LSEG. For Europe, we use data from Xetra Germany, SIX Swiss Exchange, London Stock Exchange, Euronext Paris, CBOE (both BATS and Chi-X), and Euronext Amsterdam. For Japanese equities, we source data from the Osaka Stock Exchange, Tokyo Stock Exchange, Nagoya Stock Exchange, Fukuoka Stock Exchange, and Sapporo Stock Exchange.

The rise in AT and fragmentation are considered two of the most important technological advancements in the modern history of equity markets. Hence, we include them as our main variables. In addition to these variables, we also control for total trading volume ($Volume_{i,m}$), market capitalization ($MCap_{i,m}$), the absolute value of midpoint return ($Volatility_{i,m}$), VIX (VIX_m), and TED rate (TED_m). Controlling for these characteristics allows us to interpret the association between AT/fragmentation and the bid-ask spread moments in a more robust way. $Volume_{i,m}$ is the monthly (m) average of the daily total number of shares traded for stock i, representing overall trading activities. Market capitalization, denoted by $MCap_{i,m}$, is the monthly (m) average of daily market capitalization for stock i, capturing firm size. To control for stock- and market-level volatility, we include $Volatility_{i,m}$ and VIX_m, respectively. $Volatility_{i,m}$ is the monthly average of the absolute value of daily midpoint returns. We also include the TED_m spread as a measure of funding stress. The TED_m index was discontinued in 2022. For the months without the TED index, we replace it with the difference between the 3-month Treasury yield and the Secured Overnight Financing Rate. In addition to these variables, in Equation (4), we also control for the mean of the bid-ask spread to ensure that the mechanical correlation between mean and skewness does not impact the association between skewness and explanatory variables.

The results of Equations (3) and (4) are presented in Panel A of Table 4. We include only stock fixed effects because VIX_m and TED_m values are the same across different stocks for a given month, which prevents the inclusion of time fixed effects. It is also important to note that all variables have been standardized, as we are interested in comparing the magnitude of the impact of each characteristic.

As mentioned, the impact of AT and market fragmentation on market quality has been a topic of interest in recent years. Therefore, we focus our discussion on the relationship between these two variables and the bid-ask spread moments. Our results suggest that there is a negative correlation between AT/fragmentation and the mean of the bid-ask spread, while the respective correlations are positive for skewness. This implies that while an increase in AT and market fragmentation corresponds to a decline in the mean of the bid-ask spread, it is associated with an increase in the skewness of the bid-ask spread.

This result is interesting and can be explained based on the existing literature. Hendershott et al. (2011) and Brogaard et al. (2015) show that high-frequency trading (HFT), a subset of AT, reduces the average bid-ask spread because high speed allows high-frequency market makers to update their quotes quickly, reducing their adverse selection and inventory management risks. Similarly, Degryse et al. (2015) find that (lit) fragmentation reduces the overall bid-ask spread and improves liquidity by increasing competition between liquidity providers. While there is no explicit literature on the relationship between AT/fragmentation and the skewness of the bid-ask spread, which makes our results interesting, the findings can be explained based on various streams of literature. For instance, Brogaard et al. (2018) and Aquilina et al. (2018) demonstrate that HFT may contribute to extreme price movements by reducing their liquidity provision and increasing their liquidity demand. A sudden decrease in liquidity provision can increase order imbalance and force market makers to charge extreme bid-ask spreads, which may increase skewness.

The positive association between fragmentation and skewness can be explained by the competition between market makers mechanism proposed by Roll and Subrahmanyam (2010). The idea is as follows: monopolistic market makers elevate spreads during periods of high information asymmetry while subsidizing them in low information asymmetry periods. This strategic practice enables them to charge lower spreads during high information asymmetry periods, offsetting losses incurred with informed investors. However, in competitive markets, market makers lack the ability to set spreads above a minimum during low information asymmetry periods. As a result, competitive market makers struggle to offset losses in high information asymmetry periods, leading to more extreme bid-ask spread observations and increased skewness. With this intuition, increased market fragmentation can make the market-making process more competitive (as also suggested by Degryse et al. (2015)) and lead to increased skewness. Related to this, Van Kervel (2015) shows that competition between trading venues, i.e., highly fragmented markets, can generate extreme illiquidity in one market because trades on one venue are followed by significant cancellations of limit orders on competing venues.

In the second test, we focus on the FX markets. Similar to Equations (3) and (4), we estimate the following regression model:

$$\begin{aligned} Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Volume_{i,m} &+ \beta_3 Volatility_{i,m} \\ &+ \beta_4 JPVIX_m + \beta_5 TED_m + \varepsilon_{i,m} \end{aligned} \tag{5}$$

$$\begin{aligned} Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Volume_{i,m} &+ \beta_3 Volatility_{i,m} \\ &+ \beta_4 JPVIX_m + \beta_5 TED_m + \beta_6 Mean_{i,m} + \varepsilon_{i,m} \end{aligned} \tag{6}$$

where $Mean_{i,m+1}$, $Skewness_{i,m+1}$, $Volatility_{i,m}$, and TED_m are as previously defined. The number of trades and trading volume in spot markets is not publicly available for FX instruments. Instead, we use daily trading volume in futures markets to capture trading volume. Hence, $Volume_{i,m}$ is the monthly average of daily FX futures trading volume. Linked to this, our AT proxy ($Algo_{i,m}$) for FX instruments is the monthly ratio of the number of quotes (obtained from LSEG) to futures volume. Additionally, consistent with the literature, instead of using VIX to capture market-level volatility, we use the JP Morgan FX volatility index ($JPVIX_m$) for G10 countries (Ranaldo and de Magistris, 2022).

The results are reported in Panel B of Table 4. Interestingly, the relationship between

AT and the bid-ask spread moments is consistent with observations in equity markets, notwithstanding the different overall pattern followed by skewness in this marker segment. Specifically, a one-standard-deviation increase in AT is associated with a 5.4% decrease in the average bid-ask spread and, more surprisingly, a substantial 71% increase in the skewness of the bid-ask spread.

In the final test, we explore the determinants of the bid-ask spread moments in the government bond markets. This analysis is particularly interesting because, unlike in equity and FX markets, AT and HFT are less common in the bond markets. This is largely due to bond trading being primarily dealer-driven. In Europe and Japan, nearly all bond trading is conducted exclusively by dealers. While HFT is more prevalent in the US government bond markets, the extent of HFT in US bond markets is significantly smaller compared to that in equity and FX markets.[15] Therefore, we expect to find a less pronounced effect of AT on the bid-ask spread moments in the bond markets:

$$\begin{aligned} Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} &+ \beta_2 Volume_{i,m} + \beta_3 Volatility_{i,m} \\ &+ \beta_4 VIX_m + \beta_5 TED_m + \beta_6 MOVE_m + \varepsilon_{i,m} \end{aligned} \tag{7}$$

$$\begin{aligned} Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} &+ \beta_2 Volume_{i,m} + \beta_3 Volatility_{i,m} \\ &+ \beta_4 VIX_m + \beta_5 TED_m + \beta_6 MOVE_m + \beta_7 Mean_{i,m} + \varepsilon_{i,m} \end{aligned} \tag{8}$$

where $Mean_{i,m+1}$, $Skewness_{i,m+1}$, $Volatility_{i,m}$, VIX_m and TED_m are as previously defined. Similar to the FX analysis, we use futures volume for government bonds. The only exception is Japanese bonds, where our resources allow us to obtain spot volume rather than futures volume. In addition to VIX_m, we also use the $MOVE_m$ index to capture implied bond volatility. In this analysis, we restrict our sample to the post-2010 period. This is because, as mentioned in Section 2, bond bid-ask spread data provided by LSEG is not consistent with the data described in previous studies. The results are qualitatively similar when we use the whole sample.

We report the results in Panel C of Table 4. Consistent with our expectations, the association between AT and the bid-ask spread moments is weak and not statistically significant in the bond markets.

Overall, the results in this section show that two major changes in equity markets, namely AT/HFT and market fragmentation, may potentially make markets more frag-

[15] Harkrader and Puglia (2020), estimate that PTFs are responsible for 21% of all trades in US Trasury cash markets.

ile and less resilient by increasing liquidity skewness. We observe a similar result for AT/HFT when we examine the FX market. The association between AT/HFT and the bid-ask spread moments is much weaker in bond markets, which is consistent with the notion that AT/HFT is less prevalent in bond markets compared to equity and FX markets. Regarding other characteristics, both instrument-level and market-level volatility are significant determinants of bid-ask spread moments across different markets.

6 Implications of Skewness Changes on Trading Profits

The observed increase in the bid-ask spread skewness in our study suggests that, over time, traders may face a higher probability of experiencing abnormally high trading costs during periods of extreme illiquidity. This raises at least two important issues. The first deals with the resiliency of the market as traders used to low bid ask spreads may be be surprised by episodes of illiquidity and destabilise the financial system (Persaud (2003)). The second, and more easily measurable one relates to the direct economic cost of such occurrences. We focus on this second issue in this section. We explore the potential implications of increased skewness for end-users by simulating bid-ask spreads with varying levels of skewness and applying a trading strategy to the simulated data. There are multiple ways to approach this problem, but our view is that simplicity is key. A complex model may be more related to specific trading strategies of investors, but may fail to capture the gist of our analysis.

The simulation spans 252 trading days, with 7 data points per day (hourly data from 9:30 to 16:00). The simulation process involves several key steps. Specifically, we fix the midpoint of the security price at \$100. The mean and standard deviation of the bid-ask spread are fixed at 0.53 bps and 0.19 bps, respectively, based on 2023 equity market data, with an initial skewness of 0.61. This initial skewness level is defined as a skewness factor of 1 and skewness factor is then manipulated by factors ranging from 0.5 to 1.5 in increments of 0.1, effectively decreasing and increasing skewness from 0% to 50%. For each skewness factor, bid-ask spreads are drawn from a gamma distribution. The shape and scale parameters of the gamma distribution are calculated based on the target mean, standard deviation, and skewness. The gamma distribution is used due to its flexibility in modeling skewness while maintaining constant mean and standard deviation. After generating the bid-ask spreads, the simulated data is adjusted to ensure that the mean and standard deviation align with the target values of 0.53 bps and 0.19 bps, respectively.

The adjusted spreads are then employed to generate the bid and ask prices. The ask price is calculated as the midpoint (fixed at $100) plus half the spread, while the bid price is calculated as the midpoint minus half the spread. We assume a simple trading strategy of buying at the lowest ask price and selling at the highest bid price during the day. This strategy allows traders to execute trades at the most favourable prices; more importantly, it effectively captures the impact of extreme spread changes. The cumulative profit for each skewness factor is calculated by aggregating the daily profits over 252 trading days.

The simulation results reported in Table 5 show that changes in the skewness of the bid-ask spread have a direct impact on trading profits. As skewness increases, the cumulative profit from the trading strategy decreases, underscoring the increased trading cost during periods of higher skewness. Specifically, increasing skewness by about 50% (from skewness factor 1 to 1.5) reduces the profitability of the trading strategy by about 7.2%. Conversely, reducing skewness leads to improved profitability. While the simulation itself is designed to be straightforward, we believe the results suggest that increased skewness may have significant implications for end-users, such as institutional and retail investors. Hence, it is important for investors and regulators to monitor skewness in the bid-ask spreads as an indicator of liquidity conditions and potential trading costs in financial markets.

7 Conclusions

In this paper, we study the evolution of market liquidity across equities, government bonds, and FX markets in the world's most significant jurisdictions over the past quarter-century.

We gather detailed intraday data on the bid-ask spreads faced by market participants and analyze its distribution. Our findings show that most modern financial markets are, on average, significantly more liquid than they were 25 years ago, with bid-ask spreads having declined substantially over this period. However, while equity and government bond markets have become more liquid on average, they also experience more frequent episodes of illiquidity, as indicated by the higher moments of the bid-ask spread distribution. In contrast, the FX market does not display this increased fragility. Metaphorically, market participants are navigating a sea that is often much calmer than in the past but one that is also increasingly prone to sudden and significant storms.

We conduct several additional analyses to explore the potential causes of these phenomena and perform simulations to highlight their practical implications for market participants. Our findings suggest that step changes in the mean and skewness of the bid-ask spread are linked to broad economic and geopolitical events, as well as market-specific changes and developments.

In equity markets, where detailed data allow us to proxy for AT and fragmentation, we find that these two factors are associated with lower average spreads and higher skewness. We observe a similar association for AT in the FX markets. In contrast, reflecting the lower prevalence of AT in sovereign bond markets, we do not find any significant association between our algorithmic trading measure and the bid-ask spread moments in these markets. Finally, we highlight that if the reduction in bid-ask spreads had been achieved without the accompanying increase in skewness, trading profitability would have been 7.2% higher.

References

Abergel, F., J.-P. Bouchaud, T. Foucault, M. Rosenbaum, and C.-A. Lehalle (2012) "Market microstructure: confronting many viewpoints",Technical report.

Aquilina, M., B. Eyles, J. Shao, and C. Ysusi (2018) "How do participants behave during flash events? evidence from the uk equity market", *Financial Conduct Authority Research Note, London.*

Aquilina, M., S. Foley, P. O'Neill, and T. Ruf (2024) "Sharks in the dark: Quantifying hft dark pool latency arbitrage", *Journal of Economic Dynamics and Control*, 158, p. 104786.

Bai, J. and P. Perron (1998) "Estimating and testing linear models with multiple structural changes", *Econometrica*, 66 (1), pp. 47–78.

Bai, J. and P. Perron (2003) "Computation and analysis of multiple structural change models", *Journal of Applied Econometrics*, 18 (1), pp. 1–22.

BIS Markets Committee (2011) "High-frequency trading in the foreign exchange market".

BIS Markets Committee (2016) "Electronic trading in fixed income markets".

Brogaard, J., A. Carrion, T. Moyaert, R. Riordan, A. Shkilko, and K. Sokolov (2018) "High frequency trading and extreme price movements", *Journal of Financial Economics*, 128 (2), pp. 253–265.

Brogaard, J., B. Hagströmer, L. Nordén, and R. Riordan (2015) "Trading fast and slow: Colocation and liquidity", *The Review of Financial Studies*, 28 (12), pp. 3407–3443.

Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega (2014) "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market", *Journal of Finance*, 69 (5), pp. 2045–2084.

Comerton-Forde, C., V. Grégoire, and Z. Zhong (2019) "Inverted fee structures, tick size, and market quality", *Journal of Financial Economics*, 134 (1), pp. 141–164.

Degryse, H., F. De Jong, and V. v. Kervel (2015) "The impact of dark trading and visible fragmentation on market quality", *Review of Finance*, 19 (4), pp. 1587–1622.

European Commission (2017) "Drivers of corporate bond market liquidity in the European Union".

Fleming, M. J., B. Mizrach, and G. Nguyen (2018) "The microstructure of a u.s. treasury ecn: The brokertec platform", *Journal of Financial Markets*, 40, pp. 2–22.

Fleming, M. and F. Ruela (2020) “Treasury market liquidity during the covid-19 crisis”, *Liberty Street Economics* (hal-00872398).

Goldman Sachs (2018) “Liquidity, volatility, fragility”, *Top of Mind* (68).

Harkrader, J. C. and M. Puglia (2020) “Principal trading firm activity in treasury cash markets”, *FED Notes*.

Hendershott, T., C. Jones, and A. Menkveld (2011) “Does algorithmic trading improve liquidity?”, *Journal of Finance*, 66 (1), pp. 1–33.

Ibikunle, G., M. Aquilina, I. Diaz-Rainey, and Y. Sun (2021) “City goes dark: Dark trading and adverse selection in aggregate markets”, *Journal of Empirical Finance*, 64, pp. 1–22.

Krohn, I. and V. Sushko (2022) “Fx spot and swap market liquidity spillovers”, *Journal of International Money and Finance*, 120, p. 102476.

Kyle, A. S. (1985) “Continuous auctions and insider trading”, *Econometrica*, 53 (6), pp. 1315–1335.

McGuire, P., G. von Peter, and S. Zhu (2024) “International finance through the lens of bis statistics: the global reach of currencies”, *BIS Quarterly Review*, p. 1.

Menkveld, A. J. (2016) “The economics of high-frequency trading: Taking stock”, *Annual Review of Financial Economics*, 8, pp. 1–24.

Menkveld, A. J. and B. Z. Yueshen (2018) “The flash crash: A cautionary tale about highly fragmented markets”, *Management Science*, 65 (10), pp. 4470–4488.

Mizrach, B. and C. J. Neely (2006) “The transition to electronic communications networks in the secondary treasury market”, *Federal Reserve Bank of St. Louis Review*.

Persaud, A. (2003) *Liquidity Black Holes*, London, Risk Books.

Ranaldo, A. and P. S. de Magistris (2022) “Liquidity in the global currency market”, *Journal of Financial Economics*, 146 (3), pp. 859–883.

Roll, R. (1984) “A simple implicit measure of the effective bid-ask spread in an efficient market”, *The Journal of finance*, 39 (4), pp. 1127–1139.

Roll, R. and A. Subrahmanyam (2010) “Liquidity skewness”, *Journal of Banking & Finance*, 34 (10), pp. 2562–2571.

Sakiyama, T. and S. Kobayashi (2018) “Liquidity in the jgb cash market: An evaluation from detailed transaction data”,Technical report, Bank of Japan Reports and Research Papers.

Van Kervel, V. (2015) “Competition for order flow with fast and slow traders”, *The Review of Financial Studies*, 28 (7), pp. 2094–2127.

Vayanos, D. and J. Wang (2013) “Market liquidity—theory and empirical evidence”, *Handbook of the Economics of Finance*, 2, Elsevier, pp. 1289–1361.

Werner, I. M., B. Rindi, S. Buti, and Y. Wen (2023) “Tick size, trading strategies, and market quality”, *Management Science*, 69 (7), pp. 3818–3837.

Yao, Y.-C. (1988) “Estimating the number of change-points via schwarz’ criterion”, *Statistics Probability Letters*, 6 (3), pp. 181–189.

Figure 1: **Moments of bid ask spread in the equity market.** This graph shows the mean, standard deviation, skewness and kurtosis of bid ask spread in the US, Japan, and Europe stock markets. In each year, stocks in each market are classified into terciles by their market capitalization (i.e, the large-, medium-, and small-sized stocks). Mean and standard deviations are expressed in basis points. Panel C is the average values across three European markets, namely the UK, Germany, and France stock market.

Panel A: US

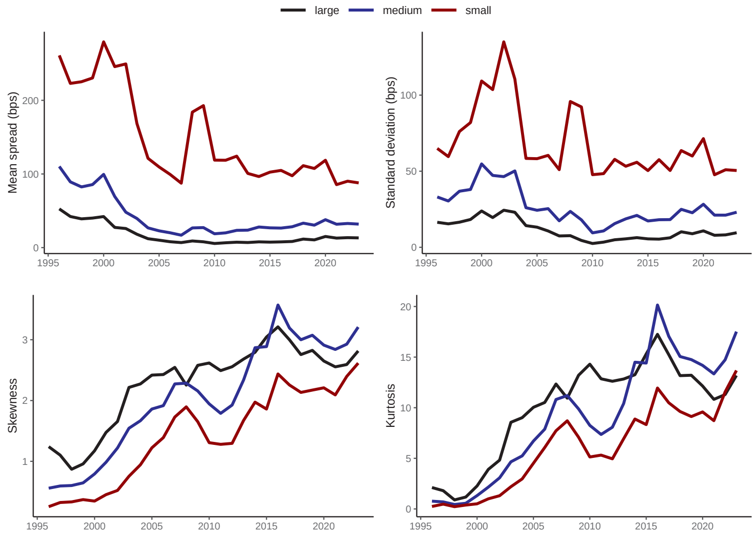

Panel B: Japan

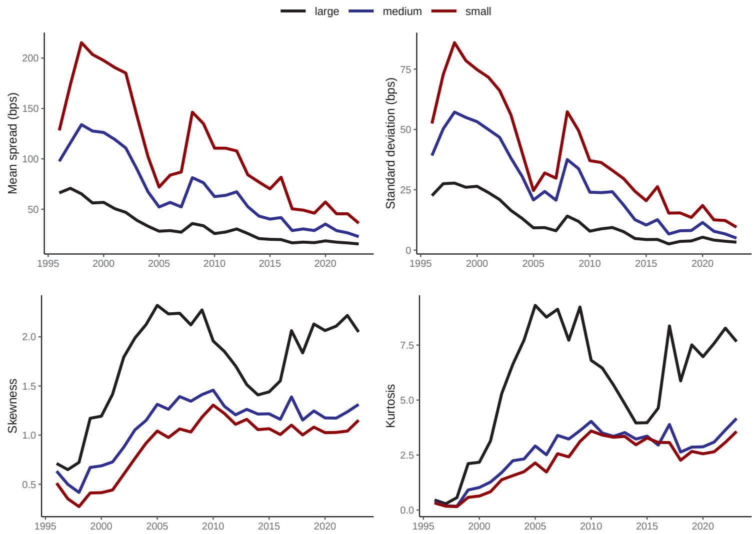

Panel C: Europe

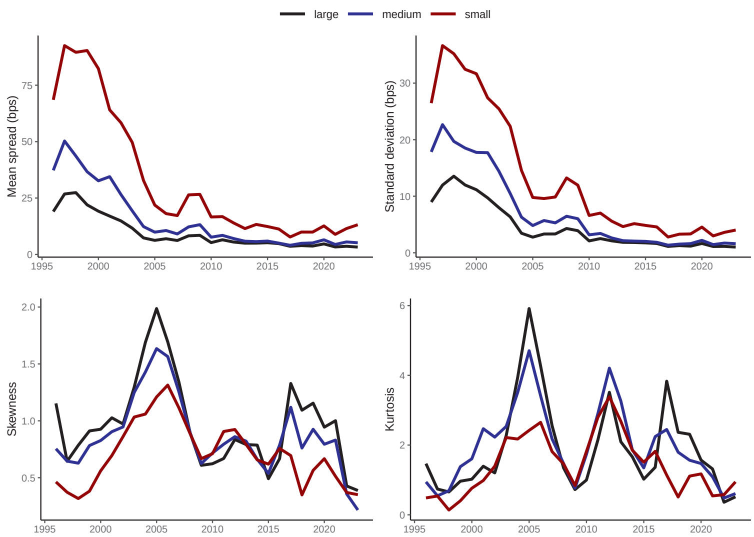

Figure 2: **Moments of bid ask spread in the foreign exchange spot market.** This graph shows the mean, standard deviation, skewness and kurtosis of bid ask spread for trading Japanese Yen (JPY), Euro (EUR), British pounds (GBP) against the US dollars. Mean and standard deviations are expressed in basis points.

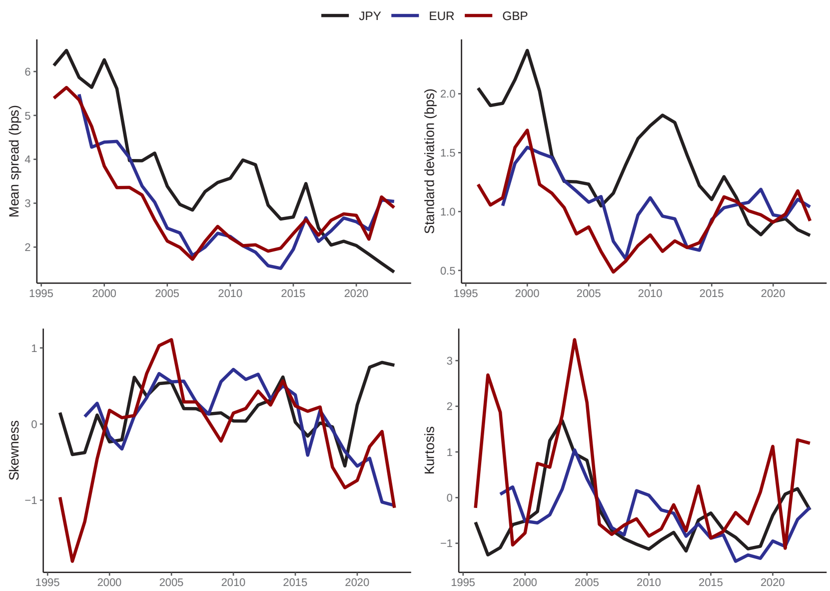

Figure 3: **Moments of bid ask spread in the sovereign bond market.** This graph shows the mean, standard deviation, skewness and kurtosis of bid ask spread for the US, Japan, and Europe government bonds with maturities of 2 years, 5 years, or 10 years. Mean and standard deviations are expressed in basis points. Panel C are the average values across the British, German and Italian government bonds.

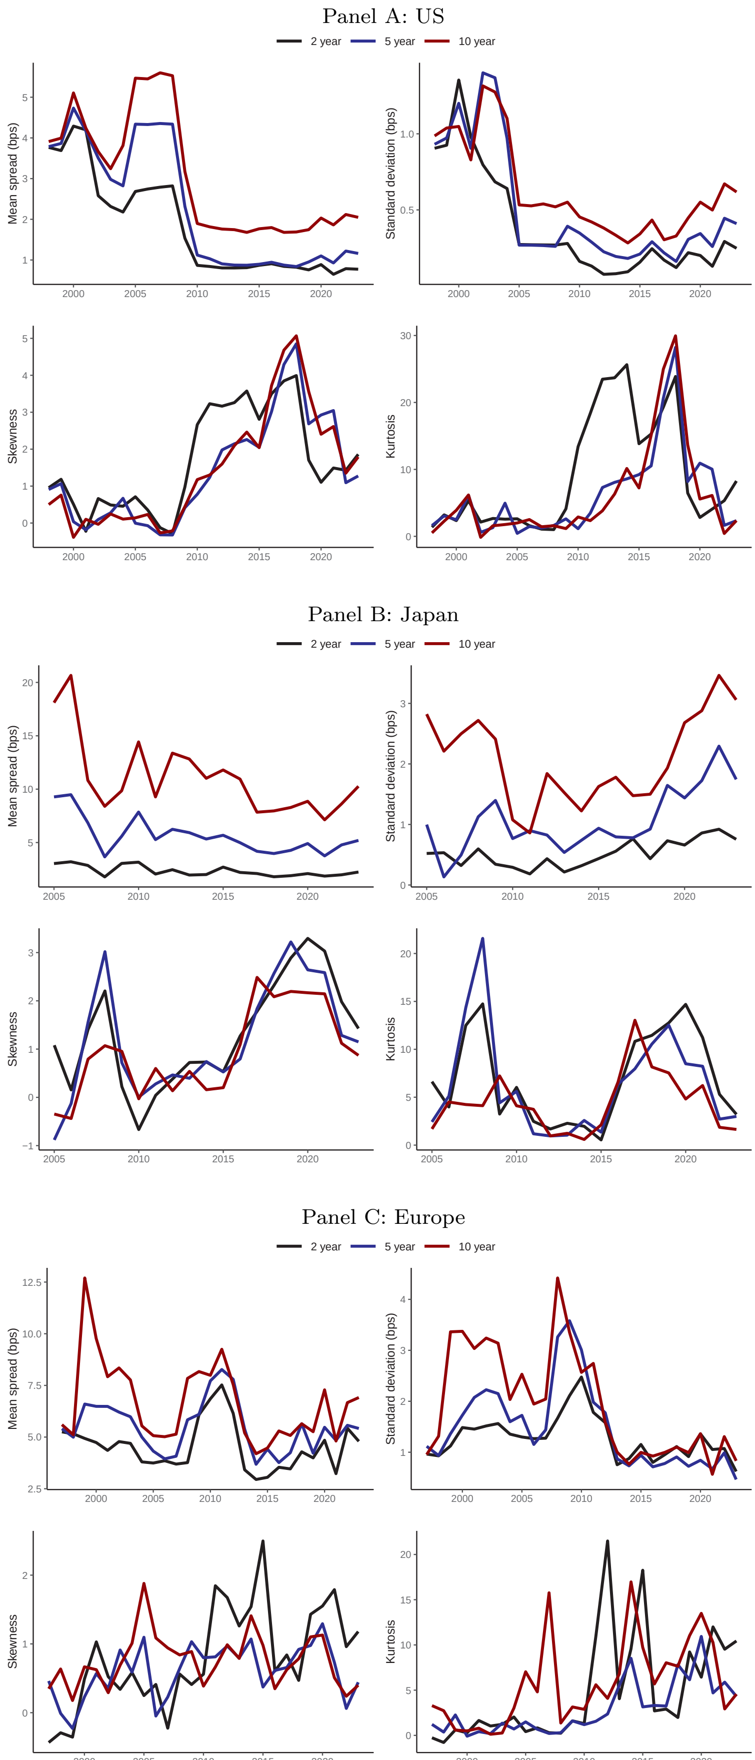

Table 1: **Moments of the bid-ask spread distribution across asset classes and regions**
This table reports the average of the first four moments (mean, standard deviation, skewness, and kurtosis) of the distribution of bid-ask spread across different asset classes (equity, bond, and foreign exchange) and regions (United States, United Kingdom, Japan, Germany, France, Italy).

	Mean (bps)	Standard deviation (bps)	Skewness	Kurtosis
Panel A: Equity market				
Large				
United States	17.02	11.08	2.28	9.97
United Kingdom	28.18	11.75	1.00	2.51
Japan	33.12	11.82	1.74	5.75
Germany	9.66	4.24	0.76	1.39
France	8.93	4.52	1.16	2.43
Medium				
United States	41.38	26.50	2.06	9.11
United Kingdom	87.72	26.45	0.42	0.97
Japan	67.14	26.33	1.10	2.61
Germany	16.70	6.95	0.72	1.84
France	13.26	7.12	1.02	2.15
Small				
United States	147.93	68.66	1.42	5.99
United Kingdom	243.41	53.91	0.08	0.78
Japan	108.41	39.10	0.90	2.22
Germany	41.05	15.33	0.59	1.19
France	22.84	10.68	0.82	1.70
Panel B: Government bond market				
Two year				
United States	1.81	0.38	1.66	8.99
United Kingdom	4.12	0.98	0.82	7.19
Japan	2.34	0.52	1.31	6.90
Germany	3.49	1.13	0.47	1.87
Italy	6.07	1.83	1.41	7.00
Five year				
United States	2.28	0.50	1.39	6.07
United Kingdom	5.39	1.18	0.32	3.34
Japan	5.64	1.06	1.20	6.34
Germany	4.11	1.43	0.75	1.15
Italy	7.12	1.92	0.86	5.14
Ten year				
United States	3.03	0.63	1.44	5.99
United Kingdom	7.34	2.09	0.60	8.73
Japan	11.08	2.08	0.94	4.44
Germany	4.63	1.56	0.82	2.35
Italy	8.11	2.25	0.93	7.15
Panel C: Foreign exchange				
GBP/USD	2.92	0.96	-0.08	0.24
JPY/USD	3.60	1.41	0.18	-0.41
EUR/USD	2.77	1.06	0.10	-0.43

Table 2
Breaks - means
This table visualises the breaks in the time series of the mean bid ask spread identified using the Bai and Perron (1998) procedure discussed in Section 4. Upward shifts are in green and downward ones are in red. The test identifies the month in which the break takes place. We split the period in six-months chunks for readability

Mean	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Europe										GFC									MiFID II				
Equity																							
Large Cap		↓					↓						↓										
Mid cap					↓				MiFID I			↓											
Small cap		↑				↓				↑				↓						↑			
Bonds														Draghi's speech									
2Y						↓				↑				↓							↑		
5Y						↓				↑				↓						↑			
10Y						↓			↑					↓							↑		
Japan																							
Equity								FIEA						Abenomics									
Large cap		↓				↓									↓								
Mid cap					↓				↑					↓					↓				
Small cap					↓				↑					↓				↓					
Bonds																							
2Y												↓											
5Y									↓									↓					
10Y									↓				↑					↓					
United States																							
Equity			Decimalization		Autoquote			Reg NMS															
Large cap			↓				↓													↑			
Mid cap			↓				↓									↑							
Small cap				↓					↑					↓									
Bonds																							
2Y				↓							↓												
5Y						↑					↓												
10Y						↑					↓												
FX	Euro			Euro notes															FX global code				
USDEUR				↓				↓						↓			↑				↑		
USDGBP		↓				↓											↑						
USDJPY			↓				↓				↑				↓				↓				

Table 3

Breaks - skewness

This table visualises the breaks in the time series of the skewness of the bid ask spread identified using the Bai and Perron (1998) procedure discussed in Section 4. Upward shifts are in green and downward ones are in red. The test identifies the month in which the break takes place. We split the period in six-months chunks for readability

Skewness	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021
Europe										GFC									MiFID II				
Equity																							
Large Cap				↑					↓								↑				↓		
Mid cap		↑				↑				↓					↓								
Small cap		↑					↓		MiFID I			↓									↑		
Bonds														Draghi's speech									
2Y												↑					↓				↑		
5Y			↑							↑													
10Y																							
Japan																							
Equity								FIEA						Abenomics									
Large Cap		↑				↑							↓					↑					
Mid cap				↑				↑					↓										
Small cap					↑						↑				↓								
Bonds																							
2Y											↓			↑						↑			↓
5Y									↑			↓						↑					↓
10Y									↑									↑					↓
United States																							
Equity			Decimalization		Autoquote			Reg NMS															
Large Cap				↑							↑				↑				↓				
Mid cap				↑				↑							↑								
Small cap		↑				↑					↓				↑						↑		
Bonds																							
2Y												↑								↓			
5Y												↑					↑				↓		
10Y												↑						↑				↓	
FX	Euro			Euro notes															FX global code				
USDEUR				↑													↓				↓		
USDGBP		↑						↓					↑						↓				
USDJPY			↑				↓										↓				↑		

Table 4

Regression analysis: the determinants of bid-ask spread moments

This table presents the results of the regression analysis, which examines the relationship between bid-ask spread moments and various explanatory variables. Our dependent variables are the mean ($Mean_{i,m+1}$) and skewness ($Skewness_{i,m+1}$) of the bid-ask spread. We first calculate the daily moments of spread using high-frequency data and then use the monthly average of daily values in the regression specifications. $Algo_{i,m}$ is the proxy for algorithmic trading, calculated as the number of quotes divided by the number of trades for stock i and month m for equities, as the number of quotes divided by futures trading volume for FX pair i and month m for FX instruments, and as the number of quotes divided by futures trading volume for bond i and month m for government bonds. $Frag_{i,m}$ is then computed as the monthly average of the daily $\frac{1}{HHI}$ index, where the HHI index is the sum of the squares of the fraction of shares for stock i traded on a venue on a given day. $Volume_{i,m}$ is the monthly (m) average of the daily total number of shares traded for stock i for equities and is the monthly average of daily futures trading volume for FX instruments and government bonds. $MCap_{i,m}$, is the monthly (m) average of daily market capitalization for stock i, $Volatility_{i,m}$ is the monthly average of the absolute value of daily midpoint returns, VIX_m is the monthly average of daily VIX index, TED_m is the monthly average of daily TED spread. The TED_m index was discontinued in 2022. For the months without the TED index, we replace it with the difference between the 3-month Treasury yield and the Secured Overnight Financing Rate. $JPVIX_m$ is the monthly average of the daily JP Morgan FX volatility index for G10 countries, and $MOVE_m$ is the monthly average of the daily MOVE index. Across all specifications, we include instrument and month fixed effects. The standard errors used to compute the t-statistics (in brackets) are double clustered by instrument and month. *, **, and *** denote the significance at 10%, 5%, and 1%, respectively.

Panel A: Equity		
	$Mean_{i,m+1}$	$Skewness_{i,m+1}$
$Algo_{i,m}$	-0.09*** (-10.16)	0.02** (2.49)
$Frag_{i,m}$	-0.17*** (-13.18)	0.39*** (23.81)
$Volume_{i,m}$	-0.05*** (-8.20)	0.11*** (10.92)
$Mcap_{i,m}$	0.03*** (3.97)	-0.07*** (-5.65)
$Volatility_{i,m}$	0.05*** (3.92)	0.03*** (4.12)
VIX_m	0.08*** (7.73)	-0.03*** (-2.82)
TED_m	0.02* (1.82)	-0.00 (-0.03)
$Mean_{i,m}$		-0.18*** (-17.21)
Stock FE	Yes	Yes
Month FE	Yes	Yes
N obs.	811,638	807,941
R^2	7.2%	15.2%

Panel B: FX		
	$Mean_{i,m+1}$	$Skewness_{i,m+1}$
$Algo_{i,m}$	-0.25***	0.05***
	(-7.44)	(2.92)
$Volume_{i,m}$	-0.41	0.03
	(-1.08)	(0.32)
$Volatility_{i,m}$	0.32***	0.03
	(4.15)	(0.19)
$JPVIX_m$	-0.07	0.21
	(-0.87)	(1.44)
TED_m	0.20	-0.09
	(1.63)	(-0.62)
$Mean_{i,m}$		-0.56***
		(-3.02)
FX pair FE	Yes	Yes
Month FE	Yes	Yes
N obs.	378	375
R^2	17.4%	36.1%

Panel C: Government Bonds		
	$Mean_{i,m+1}$	$Skewness_{i,m+1}$
$Algo_{i,m}$	-0.01	-0.07
	(-0.26)	(-1.14)
$Volume_{i,m}$	-0.02	-0.10
	(-0.22)	(-1.15)
$Volatility_{i,m}$	0.16*	0.02
	(1.87)	(0.48)
VIX_m	0.09	0.03
	(1.53)	(0.61)
TED_m	-0.01	0.03
	(-0.39)	(0.55)
$MOVE_m$	0.06**	-0.17***
	(2.03)	(-3.01)
$Mean_{i,m}$		-0.15
		(-1.52)
Bond FE	Yes	Yes
Month FE	Yes	Yes
N obs.	2,061	2,059
R^2	8.1%	6.9%

Table 5
Cumulative Profit for Different Skewness Factors
This table presents the cumulative profit from a trading strategy applied to simulated bid-ask spreads with varying levels of skewness. The simulation is based on historical data from the 2023 equity markets, with a fixed midpoint security price of $100. The mean and standard deviation of the bid-ask spread are set to 0.53 bps and 0.19 bps, respectively, with an initial skewness of 0.61. The skewness factor is varied from 0.5 to 1.5 in increments of 0.1, representing a range from decreased to increased skewness. The trading strategy involves buying at the lowest ask price and selling at the highest bid price during the day. The table shows the cumulative profit over 252 trading days for each skewness factor, the change in profit compared to the baseline skewness factor of 1.0, and the percentage change in profit.

Skewness Factor	Profit	Change in Profit	% Change
0.5	-0.7156	-0.7156 - (-0.7557) = 0.0401	5.3%
0.6	-0.7260	-0.7260 - (-0.7557) = 0.0297	3.93%
0.7	-0.7388	0.0169	2.24%
0.8	-0.7423	0.0134	1.77%
0.9	-0.7499	0.0058	0.77%
1.0	-0.7557	0	0%
1.1	-0.7680	-0.0123	-1.63%
1.2	-0.7761	-0.0204	-2.70%
1.3	-0.7882	-0.0325	-4.30%
1.4	-0.7929	-0.0372	-4.92%
1.5	-0.8100	-0.0543	-7.18%