Through Stormy Seas: How Fragile is Liquidity Across Asset Classes and Time?

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

[†]University of Technology Sydney
[¶]Bank for International Settlements
[#]Macquarie University
[‡]The University of Edinburgh
[*]Koç University
[♭] Systemic Risk Centre, London School of Economics

January 2025

Abstract

Market liquidity across asset classes has considerably increased in recent decades. Our study of stocks, foreign exchange (FX), and government bonds in the US, Europe, and Japan—using 25 years of high-frequency data—reveals a significant decline in both the average and standard deviation of bid-ask spreads across all asset classes. However, we also observe an increase in its skewness and kurtosis in equity and bond markets, indicating more frequent episodes of illiquidity. In contrast, FX markets do not show a significant increase in the higher moments of the distribution of bid-ask spreads. We identify structural breaks in the time series of spread distributions across regions and asset classes, associate these breaks with macroeconomic shocks and changing market conditions, and quantify the cost of this fragility to investors.

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 Robert Gaudiosi (discussant), Denis Gorea, Shize Li (discussant), Marco Lombardi, Benoit Mojon, Andreas Schrimpf, Vladislav Sushko and Nicholas Zarra and participants to the Sydney Banking and Financial Stability and the Australasian Finance & Banking Conferences 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 financial markets, there are episodes where market liquidity vanishes abruptly, leaving traders unable to execute orders quickly or at prevailing prices. 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. Despite the severity of these episodes, much of the attention from practitioners and regulators when evaluating the dynamics of market liquidity has been on how liquid markets are on average. In this paper, we move beyond analysing only the average liquidity level (the first moment of the distribution) and extend the analysis to the higher moments of the distribution of liquidity in stocks, FX and government bonds—key asset classes for both investors and funding activities of corporations and governments—in the most important developed markets: the US, Europe and Japan.

While the average level of liquidity is important, resilience of liquidity (its availability at all times) is equally important. The growing prevalence of liquidity crashes in modern markets has heightened interest from practitioners and regulators on liquidity resilience. A notable example is the response to the Flash Crash in US stocks in May 2010, which led to the establishment of the CFTC-SEC Advisory Committee on Emerging Regulatory Issues. The committee recommended extensive changes to market microstructure, emphasizing 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 US 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). Furthermore, the European Commission undertook substantial work on corporate bond markets, producing studies on the drivers of liquidity (European Commission, 2017).

Another reason why the resilience of liquidity is crucial is that many asset classes, such as real estate, can function effectively with low average liquidity due to their inherent characteristics—participants expect it may take months or even years to sell a property and factor this into their decisions. However, when 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) warns that such phenomena, which he terms liquidity *black holes*, can "destroy compa-

nies, cause significant economic contraction, bring down governments, rip the social fabric and steer capital away from certain markets more permanently." While these statements are probably somewhat hyperbolic, they highlight the importance of liquidity resilience.

Hence, 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 to study the distribution of liquidity. We first document several interesting patterns in the time series 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. Our results can be categorised into three main areas: facts on the distribution of liquidity, potential drivers of these facts and, finally, their implications for traders and markets. We document the following facts:

- 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.[1]

- 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. Hence FX markets have managed to significantly reduce the spread paid by traders, without increasing the prevalence of episodes of illiquidity.

[1] A similar pattern is documented in other studies. See, for instance Duffie et al. (2022).

- A robustness test conducted on equity markets in the United States shows that these patterns are present in other liquidity measures that reflect more directly trading costs. Effective spreads, realized spreads, and price impact exhibit declining averages accompanied by increasing skewness.

Moving to the potential drivers of the patterns described above, we find the following:

- 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. 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 liquid in average, but also more fragile and less resilient. A similar result is observed for AT in the FX market. Consistent with the fact 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]

- We confirm that instrument-level volatility, measured as the absolute value of mid-price return, is a significant determinant of bid-ask spread moments across all markets (Jahan-Parvar and Zikes, 2023). Specifically, higher volatility is strongly associated with higher average bid-ask spreads. A similar relationship is present between volatility and skewness, 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.

There are two main implications of our findings:

- Simulation results demonstrate that changes in the skewness of the bid-ask spread have a direct impact on trading profitability and costs. Specifically, increasing

[2]The acronym "AT" refers to both algorithmic trading and high-frequency trading.
[3]We cannot measure fragmentation for market segments other than equities.

skewness by about 50%—a change similar to that observed in stock markets—while keeping the mean and standard deviation constant, reduces the trading profitability of a simple strategy, which involves buying at the lowest ask price and selling at the highest bid price each day, by approximately 7.2%.

- While trading has become easier than ever, with spreads narrowing significantly, illiquidity episodes are now more frequent than in the past in many asset classes. Given the potentially severe consequences associated with these episodes it is important to understand how they can be mitigated. FX markets, where average spreads declined but skewness has not increased, could provide clues on effective mitigants.

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. The results of our econometric analysis are in Section 4. Section 6 reports the results of the simulation exercise, and Section 7 concludes.

2 Liquidity measure, data and summary statistics

A liquid market is one where trading costs are small, large orders can be executed (ie the market is deep), the time it takes to conduct a trade is not overly long, multiple traders can trades in close succession and so on. Given the multi-faceted nature of market liquidity, there are a number of alternative approaches that can be taken to its measurement. Liquidity measures based on the cost of trading are typically preferred in the literature (Foucault et al., 2013) as they more closely relate to the multifaceted nature of liquidity. The easiest way to measure trading costs is through the quoted spread which can be thought of as the cost of buying and immediately selling an asset (or vice versa) and is defined as the difference between the best ask and the best bid price. To take into account the fact that the absolute spread is not scale invariant, the quoted spread is typically divided by the midpoint and called the quoted relative spread:

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

There are other ways to measure trading costs including the effective and the realized

spread. The former compares the price at which the trade took place to the midpoint at the time of the trade and can account for instances in which there is not enough liquidity available at the best bid and ask. The latter takes into account the fact that the midpoint often moves in response to trades taking place and compares the trade price to the midpoint at some future point. Spread measures are not the only ones that can be used to gauge implicit trading costs; other approaches, often used when quote data are not available, include price impact measures and measures that are based on the autocovariance of returns (Kyle (1985); Roll (1984)). Price impact measures are an intuitive measure of liquidity because, in a perfectly liquid market, the price of an asset does not vary in response to trades but only in response to changes in fundamental information. Hence, the fact that trades do have a price impact can provide useful information on the liquidity of markets. A similar insight underpins measures based on the autocovariance of returns: if buy and sell orders arrive randomly, then trading prices can give an idea of the bid ask spread as buy orders would execute at the ask and sell orders at the bid. Hence traded prices would fluctuate around the midpoint and provide information on the prices at which dealers would be willing to trade. Liquidity measures other than the quoted spread however have more stringent data requirements: in particular they often require tracking what happens to the market in response to trades. Our data do not allow us to do so for most of the instrument we analyse. The quoted spread effectively captures the ease of trading, is widely available and computationally easy to calculate for the breadth of markets that are included in the analysis. It is also easily comparable across asset classes and time. Thus, our analysis relies on the first to fourth moments of relative quoted spread at the instrument and month level. However, in Section 5 we conduct a robustness test using alternative liquidity measures for US equity markets.

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, because the number of stocks is very large (about 10 thousand), we rely on 5-minute snapshots for computational reasons.

These frequencies provide sufficient time for market participants to replenish the order book after a trade, while still being granular enough to capture the prevailing conditions throughout the trading day. Data at a higher frequency would introduce more micro-level noise, making economic interpretations more challenging. For example, if the arrival of a market order consumes all the liquidity available at the best quote, the bid-ask spread would temporarily widen mechanically and, therefore, increase its skewness and kurtosis. However, the speed at which modern markets operate ensures that using 1-minute and 5-minute snapshots minimizes the risk of this behavior affecting our results. In today's markets, reaction times by fast participants are measured in fractions of a second (Desagre et al., 2024), and even after relatively large price movements, liquidity is typically replenished within a minute, at least in stock markets (Aquilina et al., 2018).

With these intraday data, we calculate daily moments (mean, standard deviation, skewness and kurtosis) of the relative bid-ask spread and then average them into monthly frequency to investigate longer-term trends in the distribution of liquidity.

To ensure the integrity of the data we have implemented a number of steps to clean them. 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 are available. 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. In particular, 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

[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.

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 empirical analysis, 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 volume 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. Our sample period for the equity market mostly ranges from January 1996 to December 2023, except for German stocks, whose tick data became available in December 1997. For the FX markets, tick data have been available for trading Euro (against US dollars) from May 1998, for trading the British pounds or the Japanese Yen from January 1996. The high-frequency trading data for western government bonds have become widely available in late 90s, with data for the US, German, British, and Italian government bonds starting from June 1998, May 1997, January 1999, and March 2001 respectively. In comparison, tick data for Japanese government bonds have only been available since January 2005.

High-frequency data are 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

bid and ask quotes can be susceptible to matching errors when multiple trading venues are involved, 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 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 4.5, 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, we calculate the four moments for each trading day using the 5-minute bid ask spread. Then, we group stocks into terciles based on their average market capitalization for each year. In the U.S. market, all four 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

[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.

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 for large stocks, 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 in large-cap stocks.

The trend in the kurtosis of bid-ask spreads mirrors that of skewness, but 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.

3.2 Foreign exchange

Figure 2 reports the distribution of bid-ask spreads 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.

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

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 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 government bonds are least exposed to liquidity risk among the countries that we study will the Japanese ones are the most

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

exposed to it. 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 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 Empirical analysis

In the previous section, we document the behavior of the bid-ask spread distribution across various asset classes and geographical areas. In this section, we conduct an in-depth analysis of the distribution of liquidity. First, to put things in context, we describe the main changes that the three asset classes we study went through over the period of analysis. We then analyse statistically the distribution of liquidity and perform three separate investigations.

- run trend regressions to assess if the patterns observed in Section 3 are statistically significant.
- Identify statistically structural changes in the moments of these distributions and hypothesize about their potential drivers.[9]

[9]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

- Explore the potential determinants of the distribution of the bid-ask spread.

4.1 The evolution of the relevant markets

Before moving to the details of how we conduct the estimation and the results it is worth summarising the main changes that the three asset classes went through over the period of analysis.

Up to the late 1990s and 2000s, trading in equities was done mainly by humans on trading floors and pits; while dealer desks intermediated most sovereign bond and foreign exchange transactions which were agreed upon over the phone. Since then all these markets experienced a substantial increase in their *electronification* and the entry of new players aiming to provide liquidity or extract profits from potential arbitrage opportunities. In equity markets, regulatory changes and technological advancements have brought two major shifts: the rise of AT/HFT and market fragmentation (Menkveld, 2014). The second of these shifts, increase in market fragmentation, has been largely driven by regulation.[10]

AT refers to the use of algorithms to make and execute trading decisions. Traders utilize two types of algorithms: (i) those designed to achieve long-term position changes at low cost and (ii) proprietary algorithms aimed at generating profits from short-term price fluctuations, commonly referred to as HFTs (Menkveld, 2014). According to SEC (2010), HFT employs sophisticated computer programs to generate a large number of orders and trades daily, rapidly liquidating positions and concluding the trading day with minimal holdings. HFT firms also invest heavily in microwave networks and locate their servers near stock exchanges' matching engines to gain microsecond, or even nanosecond, advantages in speed (Brogaard et al., 2015; Shkilko and Sokolov, 2020; Rzayev et al., 2023). The volume of HFT initially grew rapidly, peaking in the early 2010s, before stabilizing at around 52% in the US and 35% in Europe (Zaharudin et al., 2022).

The second major change in equity markets is the increasing fragmentation. As noted earlier, this shift began with global regulatory changes such as Reg-NMS and MiFID. Fragmentation can be categorized into two types: (i) on-exchange or lit market fragmentation and (ii) off-exchange fragmentation, including dark pools. Lit markets refer to

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.

[10]For instance, Reg-NMS in the US and MiFID in Europe facilitated the entry of new trading venues to enhance competition in financial markets. However, the effectiveness of these regulations depended on the development of electronic trading, which drastically reduced search costs (Menkveld, 2016).

regulated exchanges that operate limit order books and provide real-time dissemination of order prices and volumes. In contrast, off-exchange venues are not required to display real-time order book information, leading to a lack of pre-trade transparency.[11]

Equity markets have become increasingly fragmented over the past two decades, as evidenced by the significant loss of market share by traditional exchanges. For instance, the NYSE's share of trading volume for NYSE-listed stocks dropped from 82% in 2004 to 27% in 2018 (Baldauf and Mollner, 2021). A similar trend has been observed in Europe (Hagströmer, 2022). Another important empirical observation is the significant rise in off-exchange trading. In 2023, off-exchange trading accounted for 43.4% of the total market share in the US, meaning more than 4 out of 10 trades were executed outside traditional exchanges.[12] In Europe, the increase in off-exchange trading followed a slightly different trajectory. After MiFID, traditional exchanges lost market share to off-exchange venues until 2012. However, by 2022, lit markets had regained much of their share due to two factors: (i) the introduction of dark volume caps (DVC) under MiFID II[13], and (ii) the growing popularity of periodic and particularly closing auctions, driven by increasing interest from index investors (Hagströmer, 2022).

Foreign exchange markets have evolved in a way that is similar to equity markets, but with some differences. As FX markets have historically been subject to less regulatory scrutiny than equity and sovereign bond markets their evolution has been mostly determined by market forces. Chaboud et al. (2023) provide a detailed description of the changes experienced by these markets. Once mostly used by companies to hedge their currencies exposures using a network of dealers, the distinction between the interdealer and dealer-to-customer segments is now much more blurry. The relative importance of bank dealers and of non-financial customers has diminished substantially since the early 2000s, while other financial institutions are now much more important. The advent of two electronic brokers (EBS and Refinitiv) in the early 1990s was the catalyst of such change. These brokers were organised as central limit order books and quickly became the main trading venues in this segment for interdealer trades.

In the early 2000s a number of dealer to customer platforms which allowed customers

[11]For the sake of brevity and because it is beyond the scope of this paper, we do not provide a detailed discussion on off-exchange venues. These venues include dark pools, periodic auctions, and systematic internalisers, each with distinct trading mechanisms and regulations that warrant separate consideration.

[12]https://www.rblt.com/market-structure-reports/let-there-be-light-us-edition-54

[13]According to DVC, if a dark pool's market share exceeds 4% of a given stock over the last 12 months, trading on that venue is suspended for the following six months. Similarly, if the volume of all dark pools exceeds 8% of total trading in a stock over the last 12 months, all such trading is suspended for six months.

to submit requests for quotes to multiple dealers at the same time entered the scene, considerably increasing the complexity of the market. AT and HFT-involvement started growing from 2005 when EBS, allowed non-banks's computers to directly interact with its platform through an application programming interface (API).[14]

Today, electronic trading is nearly ubiquitous in FX markets, but while HFTs tend to use fully automated algorithm, other participants often leave the execution to algos, but the decision to trade is still taken by a human. The advent of HFTs in the market was also accompanied by worries of predatory behaviour by some of them. In particular the fast HFTs were believed to exploit slower participants in what has become to be known as latency arbitrage (Budish et al., 2015; BIS Markets Committee, 2011). In response to these worries, many trading venues in the market have introduced changes to avoid these shortcomings: in 2013 and 2014 speedbumps, batching of orders and randomisation have been introduced in a number of venues.

The sovereign bond market has experienced a somewhat lower degree of electronification. A recent report by the FSB (2022) provides detailed information on the characteristics of sovereign bond markets in the jurisdictions that make up our sample up to late 2021. The main insight is that while stock markets have a very similar market structure across jurisdictions, there are significant differences in the sovereign bond market.

In all jurisdictions, the secondary market can be separated into an interdealer and a dealer-to-customer segment. The former is characterised by dealers that trade mainly on-the-run bonds with each other on electronic platforms, often on an anonymous basis through a broker (referred to as inter-dealer-brokers, or IDBs). In Italy, the UK and the US, a central limit order book set-up is prevalent, but in Japan and Germany, voice broking is still substantial even in this segment.

Within the dealer to customer cash market, dealer intermediation remains predominant and trading protocols include request for quotes as well as pure voice trading. Most of the dealers active in the inter-dealer market are bank-affiliated and offer a wide range of products and services to their clients. Dealers that are not affiliated with banks exist but represent a very small share of the overall market. In many jurisdictions, dealers are registered with the National Debt Management Office and have some obligations and privileges, especially with respect to participating in the primary market.

HFTs are an important participant in the US inter-dealer segment of the market but are not significant actors elsewhere. They have a very limited footprint in the US dealer-to-customer segment and are insignificant in all other jurisdictions. Overall therefore, the

[14]This had been possible for banks since the previous year.

degree of electronification of the sovereign bond market and the level of participation of HFTs is significantly lower compared to the equity and FX market, especially if looking at the entire time period of our analysis.

4.2 Trend Regressions

Our previous discussion suggests that, overall, the mean bid-ask spread decreases across all markets (the only exception being the lack of a clear trend in EU bonds with 10-year maturities). On the other hand, while the skewness of the bid-ask spread increases in equity markets in the US and Japan, as well as in all government bond markets, equity markets in the EU and FX markets show no clear trend. To further investigate these patterns, we estimate formal trend regressions on the mean and skewness of the bid-ask spread in this section. This approach complements our discussion in Section 3 and, more importantly, allows us to assess the statistical significance of changes in the moments of the bid-ask spread. We focus on the mean and skewness for two main reasons: first, the high correlation between the mean and standard deviation, as well as between skewness and kurtosis, limits the additional insight gained from analyzing all moments. Second, positive skewness in liquidity is particularly relevant to our study, as it offers more meaningful implications than kurtosis. Since the spread distribution is truncated on the left (as the spread cannot be negative), a left-skewed distribution—implying frequent occurrences of particularly narrow spreads—would be advantageous for market participants.

To formally assess these trends, we estimate the following regression:

$$Mean_t = \alpha + \beta X_t + \epsilon_t, \tag{2}$$

$$Skewness_t = \alpha + \beta X_t + \epsilon_t, \tag{3}$$

where $Mean_t$ and $Skewness_t$ are the mean and skewness of the bid-ask spread for month t. These moments are calculated daily based on intraday data, and then averaged across the month for use in the regression. The variable X_t is a time indicator that equals 1 for the first month and increases by 1 with each subsequent month. We compute standard errors using the Newey and West (1987) method with 12 lags. We estimate Equations (2) and (3) separately for each jurisdiction and asset class, consistent with the figures in Section 3.

The results of the trend regression are presented in Table 2. These findings largely

confirm our discussion in Section 3 and show that time-series changes in the bid-ask spread are statistically significant in most cases. However, there are some differences across asset classes and jurisdictions that warrant further discussion. First, while the mean bid-ask spread decreases significantly in equity markets across all jurisdictions, the skewness of the bid-ask spread increases only in the US and Japan. In contrast, skewness in the EU shows no significant change in small and mid-cap stocks, with only a weakly significant decrease observed for large-cap stocks. This pattern is interesting and merits further attention.

To understand why the EU exhibits different behaviour in terms of skewness evolution, it is crucial to explore what factors drive skewness. In Section 4.5, we examine potential determinants of skewness based on the relevant literature. Our analysis suggests that the prevalence of AT/HFT, along with market fragmentation, is positively correlated with skewness in equity markets. Additionally, we suggest that trading volume, market capitalization, and volatility also influence skewness. Interestingly, when we compare these characteristics across different jurisdictions, our proxy for AT/HFT activity—measured by the messaging frequency of HFTs—is significantly higher in the US and Japan compared to the EU. We do not observe a similar pattern among the other variables that could correlate to the increase in skewness in the US and Japan relative to the EU. Consistent with this observation, industry reports and academic studies suggest that HFT activity in the US has consistently been higher than in the EU.[15] Additionally, while HFT activity arrived in the EU before Japan, HFT activity in Japan have surpassed that in the EU since 2015 due to technological upgrades implemented by the Japan Exchange Group (Kiuchi, 2022). This is particularly interesting as, according to Figure 1, the evolution of skewness in the EU and Japan initially follows a similar path; however, post-2015, skewness increases in Japan while declining in the EU. This suggests that the prevalence of HFT activities may help explain the cross-sectional differences in the evolution of the skewness of the bid-ask spread between the US/Japan and EU markets.

The reduction in the mean bid-ask spread is also statistically significant in bond and FX markets, with the exception of European bond markets, where no clear trend is observed. For European bonds with 2- and 5-year maturities, the average bid-ask spread decreases, though not significantly, while for 10-year bonds, the increase is only marginally significant at the 10% level.

Skewness in government bond markets increases significantly, similar to equity markets, however with a relatively smaller increase in Japan. However, in FX markets, we

[15] https://www.esma.europa.eu

do not observe any significant change in skewness. Looking at this result in conjunction with Figure 2, the skewness increases until 2005, then stabilizes before eventually declining. As a result, our trend regression does not show any statistically significant change in skewness over time. We conjecture that this could be related to the unique nature of automation in FX markets. Since the early 2000s, FX markets have increasingly shifted toward electronic trading, driven by two key APIs: the Bank API in 2004 and the Non-Bank API (used by PTFs and HFTs) in 2005. The main distinction between these APIs is their functionality—Bank APIs are focused on automated execution, where human traders make decisions and algorithms execute orders, while Non-Bank APIs automate both trading decisions and execution, often employing more aggressive trading strategies.

This distinction may have important implications for the skewness of the bid-ask spread. Around the early 2000s, we observe an increase in the skewness of bid-ask spreads, which could be linked to the rise of overall electronic trading. A critical question is why FX markets do not exhibit the same sustained increase in skewness seen in equity and bond markets. One plausible factor is the balanced market quality impact between Bank and Non-Bank APIs, as suggested by Chaboud et al. (2014). Additionally, FX markets have adapted swiftly to mitigate the impact of HFT activities through measures such as "speed bumps" introduced by major players and trading venues in 2013-2014. These measures have reduced the influence of Non-Bank APIs, making FX markets less susceptible to aggressive HFT activities compared to equity markets, which could be related to the skewness of the bid-ask.

In summary, while the mean bid-ask spread decreases across almost all asset classes and jurisdictions, with the exception of 10-year bonds in the EU, there is greater heterogeneity in the skewness of the bid-ask spread. Specifically, in European equity and FX markets, skewness either remains unchanged or decreases slightly, whereas it increases significantly in most other markets. Based on our analysis of the data, academic literature, and industry reports, we attribute this to the differing nature of automation and HFT in these markets. However, this explanation should be interpreted with caution because our discussion is rather speculative, given the lack of formal theory investigating the determinants of skewness.[16]

[16] An exception is Foucault et al. (2005), which is discussed in Section 4.5.

4.3 Estimating the structural breaks

The regressions above can only tell us about the general trend followed by the moments of the bid ask spread, but it is of course possible that these moments do not simply change over time monotonically. Hence we look to determine statistically if the time series exhibits structural breaks.

We focus on the mean and skewness of the relative bid-ask spread and 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:

$$Y_t = X_t'\beta + Z_t'\delta_j + \varepsilon_t \tag{4}$$

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.[17]

The results of our estimation are visually summarized in Tables 3 and 4. 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

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

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.

4.4 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 conducted a comprehensive literature review to identify potential explanatory factors. Our search included 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 relied on ChatGPT to explore potential reasons why the statistical tests identify these breaks.

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[18] of stock prices and the introduction of Autoquote[19] 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. In FX markets the introduction of the bank and non-bank APIs were associated with a decrease in both the mean and skewness of the bid ask spread. Interestingly, while there seem to be substantial lags the introduction of speedbumps and other methods to reduce predatory HFT behaviour is associated with a some increase in the mean but a reduction in skewness.

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 segments. MiFID II expanded some of the provisions of MiFID I to markets other than

[18]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.

[19]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.

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.

4.5 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:

$$\begin{aligned} Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Frag_{i,m} + \beta_3 Volume_{i,m} + \beta_4 MCap_{i,m} \\ + \beta_5 Volatility_{i,m} + \beta_6 VIX_m + \beta_7 TED_m + \varepsilon_{i,m} \end{aligned} \tag{5}$$

$$\begin{aligned} Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Frag_{i,m} + \beta_3 Volume_{i,m} + \beta_4 MCap_{i,m} \\ + \beta_5 Volatility_{i,m} + \beta_6 VIX_m + \beta_7 TED_m + \beta_8 Mean_{i,m} + \varepsilon_{i,m} \end{aligned} \tag{6}$$

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 (6), 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 (5) and (6) are presented in Panel A of Table 5. 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. First, our results suggest a negative correlation between AT/HFT and the mean of the bid-ask spread, while the correlation with skewness is positive. This suggests that an increase in AT corresponds to a reduction in the average bid-ask spread but is associated with greater skewness in its distribution. The negative correlation between AT/HFT and the mean of the bid-ask spread is relatively straightforward to explain, as it is in line with extensive literature. For instance, Hendershott et al. (2011) and Brogaard et al. (2015) show that AT/HFT 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 (Menkveld, 2013).

In contrast, the positive correlation between AT/HFT and the skewness of the bid-ask spread is less well understood, as there is no explicit theory on this aspect. An exception is Foucault et al. (2005), which link skewness in the bid-ask spread to the composition of traders based on their level of impatience. The study suggests that when the proportion of impatient traders is relatively high, markets become less resilient, leading to a right-skewed distribution of spreads. Comparing fast and slow markets, the study indicates that while fast markets are generally more liquid, they are also less resilient. This implies that bid-ask spreads may be right-skewed in fast markets due to their reduced resilience. Our findings support this view, showing a positive correlation between fast traders (AT/HFT) and the skewness of the bid-ask spread.

Another potential link between AT/HFT and skewness may relate to the role of

HFTs as liquidity providers or demanders. For instance, Aquilina et al. (2018) show that HFTs can contribute to extreme price movements by reducing liquidity provision and increasing liquidity demand. Similarly, Brogaard et al. (2018) find that while HFTs typically supply more liquidity than they demand during extreme price movements in individual stocks, they tend to demand more liquidity than they provide when such movements affect multiple stocks. A sudden reduction in liquidity provision can lead to order imbalances, causing market makers to impose wider and more extreme bid-ask spreads, which may increase skewness.

The negative correlation between market fragmentation and the mean of the bid-ask spread is also consistent with the literature. For example, O'Hara and Ye (2011) and Degryse et al. (2015) show that total market fragmentation reduces the overall bid-ask spread and improves liquidity by increasing competition between liquidity providers. However, explaining the positive effect of market fragmentation on the skewness of the bid-ask spread is more challenging, as there is limited literature on this topic. One possible explanation can be drawn from the competition mechanism among market makers proposed by Roll and Subrahmanyam (2010). The concept is that market makers set spreads above a minimum during periods of low information asymmetry, allowing them to increase spreads less than they otherwise would during periods of high information asymmetry. Essentially, market makers subsidize their losses during high information asymmetry by charging slightly higher spreads during low information asymmetry. This strategy allows them to offer lower spreads during periods of high information asymmetry than they would otherwise, helping to offset potential losses from trading with informed investors. However, in more competitive markets, market makers cannot maintain spreads above a minimum during low information asymmetry periods. As a result, they struggle to compensate for losses during high information asymmetry, leading them to set excessively wide bid-ask spreads. This results in more extreme spread observations and increased skewness. With this understanding, increased market fragmentation—which fosters greater competition among market makers (as noted by Degryse et al. (2015))—can lead to higher skewness in bid-ask spreads. Supporting this view, 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 (5) and (6), we

estimate the following regression model:

$$Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Volume_{i,m} + \beta_3 Volatility_{i,m} + \beta_4 JPVIX_m + \beta_5 TED_m + \varepsilon_{i,m} \tag{7}$$

$$Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Volume_{i,m} + \beta_3 Volatility_{i,m} + \beta_4 JPVIX_m + \beta_5 TED_m + \beta_6 Mean_{i,m} + \varepsilon_{i,m} \tag{8}$$

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 5. 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.[20] Therefore, we expect to find a less pronounced effect of AT on the bid-ask spread moments in the bond markets:

$$Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Volume_{i,m} + \beta_3 Volatility_{i,m} + \beta_4 VIX_m + \beta_5 TED_m + \beta_6 MOVE_m + \varepsilon_{i,m} \tag{9}$$

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

$$Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Volume_{i,m} + \beta_3 Volatility_{i,m} + \beta_4 VIX_m + \beta_5 TED_m + \beta_6 MOVE_m + \beta_7 Mean_{i,m} + \varepsilon_{i,m} \quad (10)$$

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 5. 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 fragile 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.

5 Alternative Liquidity Measures

The multidimensional nature of liquidity has led to various measurement approaches in the literature, as discussed in Section 2. We employ the relative quoted spread as our primary liquidity measure because it enables cross-asset class analysis without the data requirements that typically constrain other measures. However, we are aware that this approach has limitations. In particular, quoted spreads may not accurately reflect actual trading costs, particularly for large trades. For example, when depth at the best quotes is limited, larger orders often require execution at prices further into the order book, resulting in higher realized trading costs despite narrow displayed spreads at the best prices. Furthermore, the quoted spread may not reflect actual transaction costs accurately when trades execute between the posted bid and ask prices (Petersen and

Fialkowski, 1994).

A second limitation relates to transitory high-frequency liquidity fluctuations. One might argue that increased skewness in quoted spreads results from temporary spikes following large market orders that consume available liquidity at the best quotes. This could lead to more positively skewed distributions of the bid-ask spread, while the liquidity may still be plentiful. Such increased skewness would be less concerning if quoted spreads revert quickly through competitive quote updates, enabling trades to execute at narrower bid-ask spreads. First and foremost, it is important to note that this second limitation is partially mitigated by our sampling frequency—five-minute intervals for equity markets and one-minute intervals for Treasury and FX markets—which ensures that high-frequency spread spikes are unlikely to persist in our data. For example, O'Hara (2015) argues that in a high-frequency market microstructure environment, five minutes are effectively a lifetime, and five to fifteen seconds are sufficient to capture temporary price impacts. This view is further supported by Bessembinder et al. (2016), who document that the temporary price impact of order imbalances in crude oil ETFs is almost entirely reversed within one minute. Given such rapid reversals, our sampling frequencies, particularly the five-minute interval for equities, provide sufficient time for transient spread movements to revert.

In any case, to address these limitations further, we complement our analysis with effective spread measures and their components—realized spread and price impact. Due to data constraints, we conduct this additional analysis only for the equity market in the United States. The effective spread captures actual trading costs by measuring the difference between execution prices and the prevailing mid-quote at the time of trade. This measure can be decomposed into two components. The first component, price impact, measures how much prices move in the direction of the trade and typically reflect the losses incurred by liquidity providers due to adverse selection (Brogaard et al., 2015). The second component, realized spread, represents the portion of the effective spread that liquidity providers retain after accounting for price impact. We calculate price impact and realized spread as follows:

$$pi_{it} = \frac{|MidPrice_{it+5} - MidPrice_{it}|}{MidPrice_{it}}. \tag{11}$$

$$rs_{it} = \frac{|TradePrice_{it} - MidPrice_{it+5}|}{MidPrice_{it}}. \tag{12}$$

where pi_{it} and rs_{it} are price impact and realized spread over a five-minute interval $(t+5)$.

The effective spread is the sum of pi_{it} and rs_{it} (Conrad and Wahal, 2020).

The calculation of pi_{it} and rs_{it} typically requires transaction-level data to determine trade direction and compute these measures accurately. However, using transaction-level data presents two challenges in our context. First, trade-level data would capture temporary liquidity fluctuations that reverse within seconds. As discussed above, our analysis focuses on examining persistent liquidity patterns rather than short-term fluctuations Second, processing such granular data for the entire U.S. market from 1996 to 2022 is computationally infeasible. To address these challenges, we use five-minute data and calculate the measures with the absolute difference in the numerator to ensure they consistently reflect the positive cost incurred by liquidity demanders, akin to the role of trade sign adjustments in transaction-level analysis.

As seen in Figure 4, across all measures—effective spread, realized spread, and price impact—we observe a similar evolution to that of quoted spread: declining averages accompanied by increasing skewness in their distributions. This consistency across different liquidity measures supports our interpretation of the relative quoted spread. While overall trading costs have decreased, markets are experiencing more frequent episodes of severe illiquidity, characterized by extreme price impact and increased compensation demanded by liquidity providers. Moreover, the consistent trend of increased skewness in realized spreads and, particularly, price impact over the five-minute interval suggests that this phenomenon is not solely driven by high-frequency temporary quote movements.

6 Implications of Skewness Changes for Trading Profits

The observed increase in the skewness of the bid-ask spread 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 we run 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 6 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, highlighting 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%. In contrast, 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 25 years.

We gather detailed intraday data on the bid-ask spreads faced by market participants and analyze its distribution. Our findings show that 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.

We conduct a number of tests to explore the potential causes of these phenomena and perform simulations to highlight their practical implications for market participants. Our findings suggest that 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 demonstrate the the cost of increased market fragility by showing that increase in skewness can significantly erode the profitability of trading strategies of end users.

Taken together, our results show that while markets are now on average much more liquid than in the past, they are also more subject to episodes of illiquidity in many cases. 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.

References

Abergel, F., J.-P. Bouchaud, T. Foucault, C.-A. Lehalle, and M. Rosenbaum (2012) *Market microstructure: confronting many viewpoints*, John Wiley & Sons.

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.

Baldauf, M. and J. Mollner (2021) "Trading in fragmented markets", *Journal of Financial and Quantitative Analysis*, 56 (1), pp. 93–121.

Bessembinder, H., A. Carrion, L. Tuttle, and K. Venkataraman (2016) "Liquidity, resiliency and market quality around predictable trades: Theory and evidence", *Journal of Financial economics*, 121 (1), pp. 142–166.

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.

Budish, E., P. Cramton, and J. Shim (2015) " The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response", *The Quarterly Journal of Economics*, 130 (4), pp. 1547–1621.

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

Chaboud, A., D. Rime, and V. Sushko (2023) "The foreign exchange market", *Research Handbook of Financial Markets*, Edward Elgar Publishing, pp. 253–275.

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.

Conrad, J. and S. Wahal (2020) "The term structure of liquidity provision", *Journal of Financial Economics*, 136 (1), pp. 239–259.

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.

Desagre, C., F. Laly, and M. Petitjean (2024) "Revisiting the trading activity of high-frequency trading firms around flash events", *Financial Innovation*, Forthcoming.

Duffie, D., T. Foucault, L. Veldkamp, and X. Vives (2022) *Technology and Finance*, Centre for Economic Policy Research.

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", *Federal Reserve Bank of New York Staff Reports*.

Foucault, T., O. Kadan, and E. Kandel (2005) "Limit order book as a market for liquidity", *The Review of Financial Studies*, 18 (4), pp. 1171–1217.

Foucault, T., M. Pagano, and A. Röell (2013) *Market liquidity: theory, evidence, and policy*, Oxford University Press, USA.

FSB (2022) "Liquidity in core government bond markets",Technical report, FSB.

Hagströmer, B. (2022) "Market fragmentation in Europe", *Available at SSRN 4274523*.

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

Hendershott, T., C. Jones, and A. Menkveld (2011) "Does algorithmic trading improve liquidity?", *The 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.

Jahan-Parvar, M. R. and F. Zikes (2023) "When Do Low-Frequency Measures Really Measure Effective Spreads? Evidence from Equity and Foreign Exchange Markets", *The Review of Financial Studies*, 36 (10), pp. 4190–4232.

Kiuchi, T. (2022) "High-frequency trading in japan: A unique evolution", *The Future of Financial Systems in the Digital Age*, p. 159.

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*, pp. 1–16.

Menkveld, A. J. (2013) "High frequency trading and the new market makers", *Journal of Financial Markets*, 16 (4), pp. 712–740.

Menkveld, A. J. (2014) "High-frequency traders and market structure", *Financial Review*, 49 (2), pp. 333–344.

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*.

Newey, W. K. and K. D. West (1987) "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix", *Econometrica*, 55 (3), pp. 703–708.

O'Hara, M. and M. Ye (2011) "Is market fragmentation harming market quality?", *Journal of Financial Economics*, 100 (3), pp. 459–474.

O'Hara, M. (2015) "High frequency market microstructure", *Journal of financial economics*, 116 (2), pp. 257–270.

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

Petersen, M. A. and D. Fialkowski (1994) "Posted versus effective spreads: Good prices or bad quotes?", *Journal of Financial Economics*, 35 (3), pp. 269–292.

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.

Rzayev, K., G. Ibikunle, and T. Steffen (2023) "The market quality implications of speed in cross-platform trading: Evidence from frankfurt-london microwave", *Journal of Financial Markets*, 66, p. 100853.

Sakiyama, T. and S. Kobayashi (2018) "Liquidity in the JGB cash market: an evaluation from detailed transaction data", *Bank of Japan Reports and Research Papers*.

SEC (2010) "Concept Release on Equity Market Structure; File No. S7-02-10",Technical report.

Shkilko, A. and K. Sokolov (2020) "Every cloud has a silver lining: Fast trading, microwave connectivity, and trading costs", *The Journal of Finance*, 75 (6), pp. 2899–2927.

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.

Zaharudin, K. Z., M. R. Young, and W.-H. Hsu (2022) "High-frequency trading: Definition, implications, and controversies", *Journal of Economic Surveys*, 36 (1), pp. 75–107.

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.

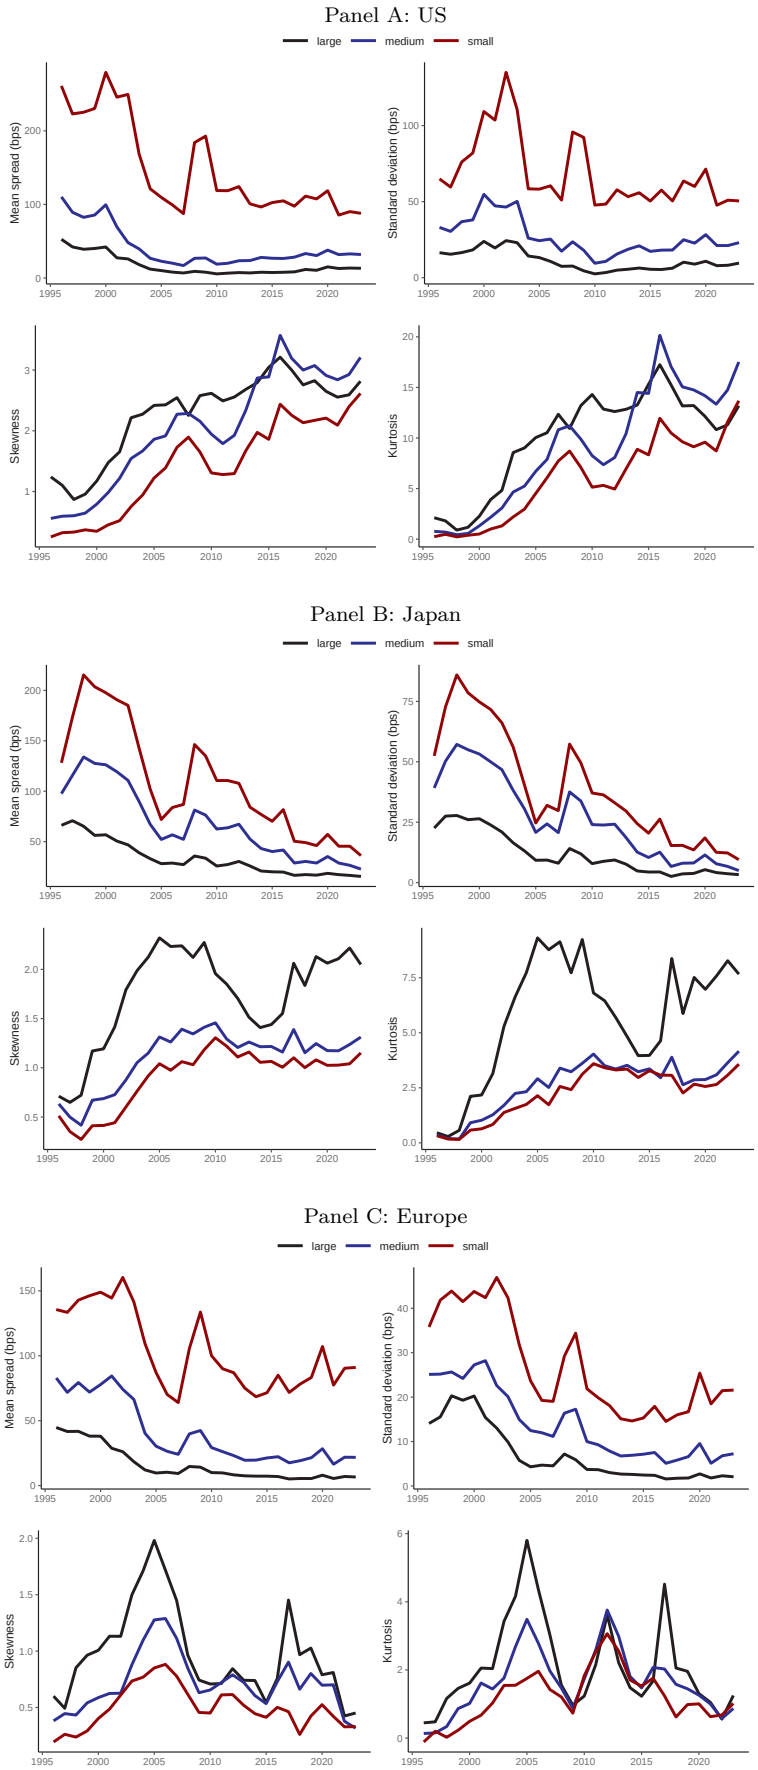


Figure 2: **Moments of bid ask spread in the foreign exchange spot market.** This graph shows the mean, standard deviation, skewness and kurtosis of bid ask spread for trading Japanese Yen (JPY), Euro (EUR), British pounds (GBP) against the US dollars. Mean and standard deviations are expressed in basis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

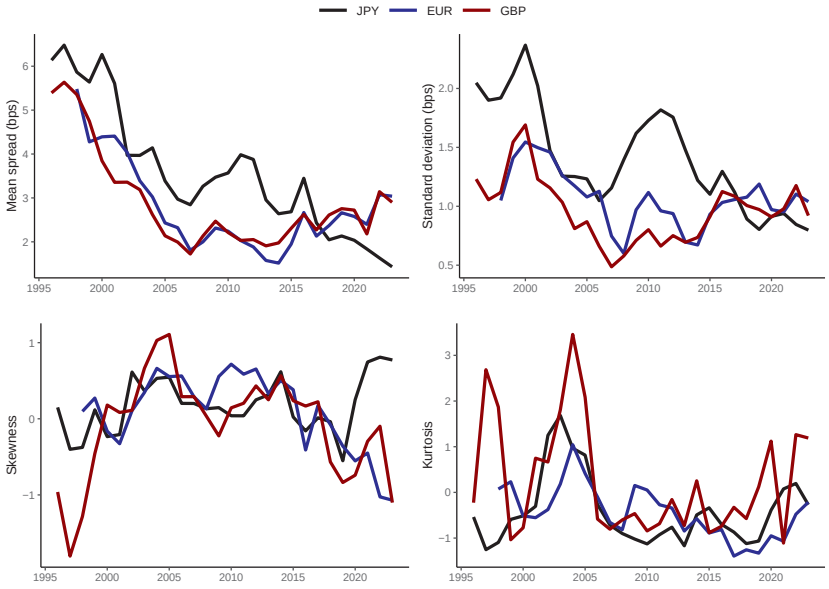


Figure 3: **Moments of bid ask spread in the sovereign bond market.** This graph shows the mean, standard deviation, skewness and kurtosis of bid ask spread for the US, Japan, and Europe government bonds with maturities of 2 years, 5 years, or 10 years. Mean and standard deviations are expressed in basis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

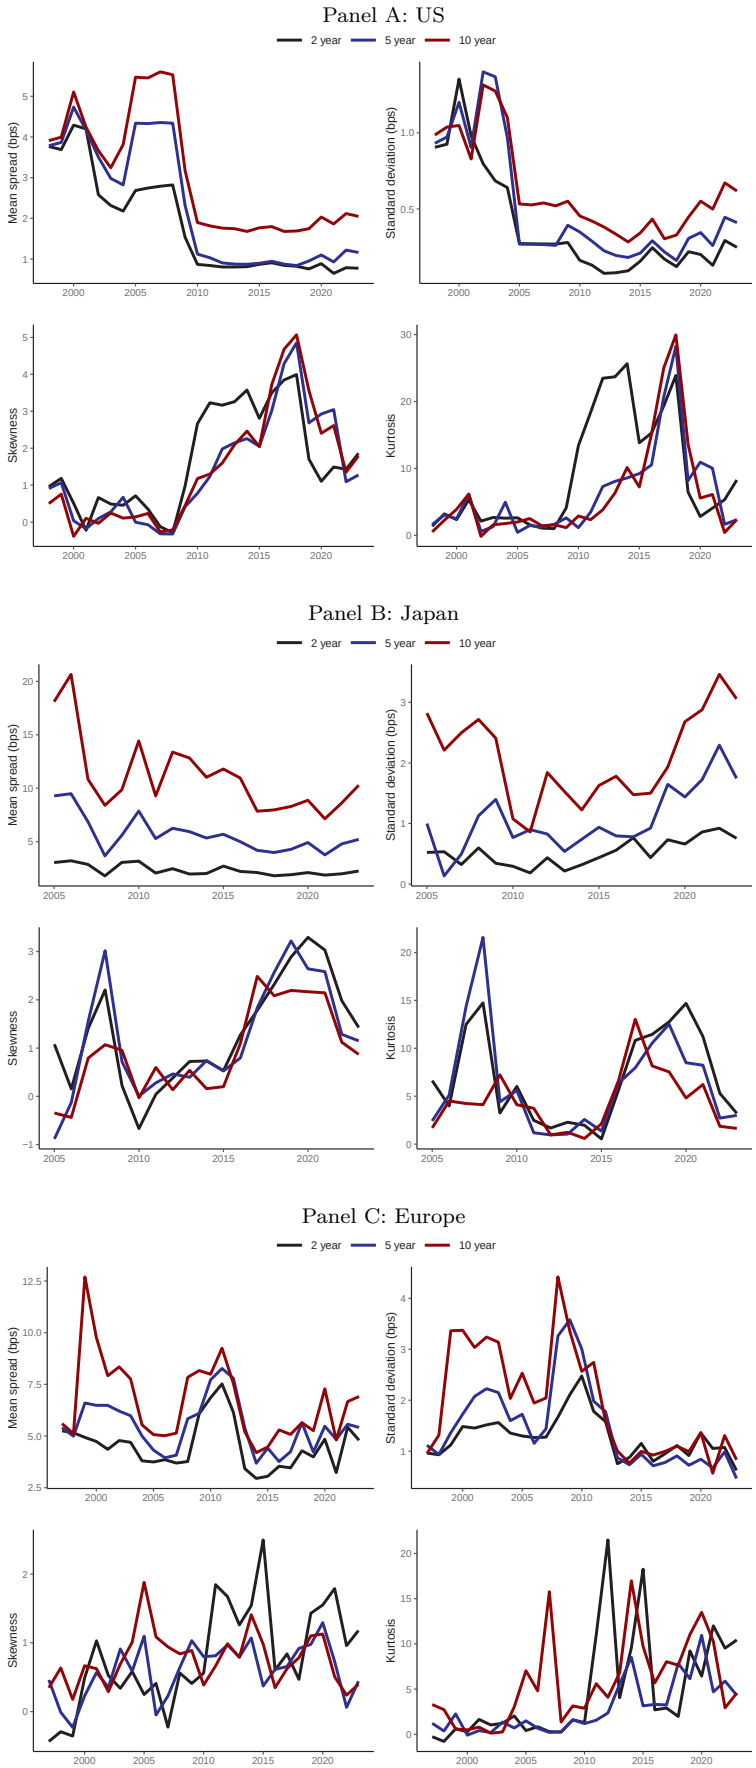


Figure 4: Alternative measures of market liquidity

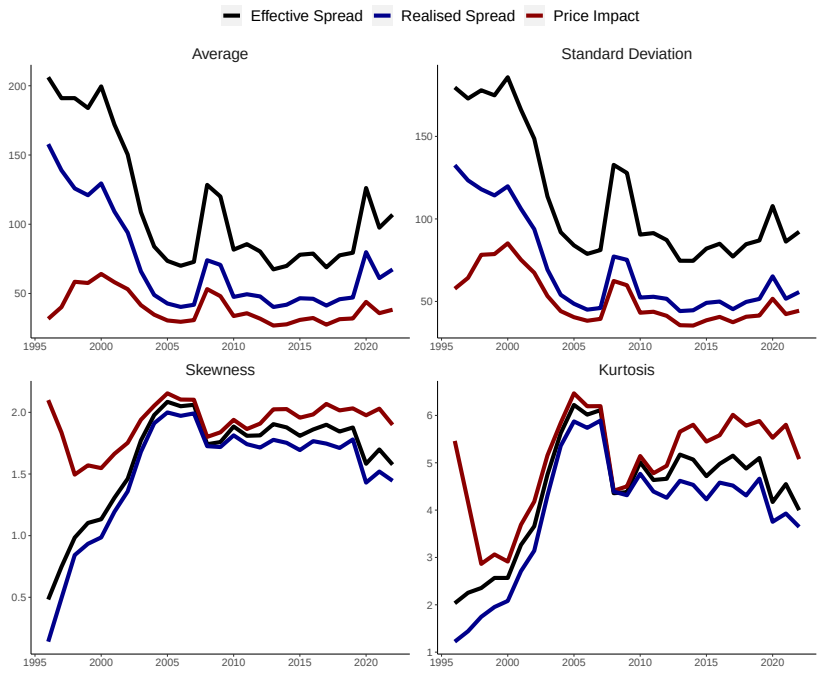


Table 1: **Moments of the bid-ask spread distribution across asset classes and regions**
This table reports the average of the first four moments (mean, standard deviation, skewness, and kurtosis) of the distribution of bid-ask spread across different asset classes (equity, bond, and foreign exchange) 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
Trend Regression
This table presents the results of the trend regression analysis for the mean and skewness of the bid-ask spread. Specifically, we estimate the following regression model:

$$Moment_t = \alpha + \beta X_t + \epsilon_t,$$

where $Moment_t$ is either the mean ($Mean_t$) or skewness ($Skewness_t$) of the bid-ask spread for month t. These moments are calculated daily using intraday data, then averaged across each month for the regression. The variable X_t is a time indicator, starting at 1 for the first month and increasing by 1 each month. Standard errors are computed using the Newey and West (1987) method with 12 lags. We estimate the trend regression separately for each jurisdiction and asset class. *, **, and *** denote the significance at 10%, 5%, and 1% respectively.

	$Mean_t$		$Skewness_t$	
	Coefficient (bps)	t-stat	Coefficient	t-stat
Panel A: Equity market				
Large				
United States	-8.55***	-4.69	0.006***	7.48
EU	-10.56***	-7.75	-0.001*	-1.94
Japan	-15.34***	-11.50	0.003***	4.30
Medium				
United States	-16.62***	-4.58	0.009***	15.84
EU	-19.46***	-9.71	-0.000	-0.26
Japan	-31.57***	-11.94	0.002***	5.21
Small				
United States	-36.27***	-7.70	0.007***	22.31
EU	-18.46***	-6.59	-0.000	-0.52
Japan	-46.51***	-9.43	0.002***	6.49
Panel B: Government bond market				
Two year				
United States	-1.17***	-11.14	0.011***	3.31
EU	-0.08	-0.49	0.006***	6.73
Japan	-0.43***	-3.47	0.008*	1.89
Five year				
United States	-1.44***	-11.79	0.012***	4.12
EU	-0.004	-0.02	0.003**	2.47
Japan	-1.63***	-3.09	0.007	1.51
Ten year				
United States	-1.27***	-9.16	0.013***	4.84
EU	0.39*	1.85	-0.001	-1.28
Japan	-3.34***	-2.93	0.008***	3.22
Panel C: Foreign exchange				
GBP/USD	-0.74***	-3.84	0.001	0.34
JPY/USD	-1.38***	-13.44	0.001	1.35
EUR/USD	-0.64***	-3.36	-0.003**	-2.55

Table 3

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 **??**. 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



Table 4

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 **??**. 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



Table 5

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 fixed effect. 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***	0.02**
	(-10.16)	(2.49)
$Frag_{i,m}$	-0.17***	0.39***
	(-13.18)	(23.81)
$Volume_{i,m}$	-0.05***	0.11***
	(-8.20)	(10.92)
$Mcap_{i,m}$	0.03***	-0.07***
	(3.97)	(-5.65)
$Volatility_{i,m}$	0.05***	0.03***
	(3.92)	(4.12)
VIX_m	0.08***	-0.03***
	(7.73)	(-2.82)
TED_m	0.02*	-0.00
	(1.82)	(-0.03)
$Mean_{i,m}$		-0.18***
		(-17.21)
Stock 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
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
N obs.	2,061	2,059
R^2	8.1%	6.9%

Table 6

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%