BIS

BIS Working Papers

No 1229

Through stormy seas: how fragile is liquidity across asset classes and time?

by Nihad Aliyev, Matteo Aquilina, Khaladdin Rzayev and Sonya Zhu

Monetary and Economic Department

November 2024 (revised May 2025)

JEL classification: G10, G12, G14.

Keywords: Liquidity, market resiliency, high frequency trading, fragmentation.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Through stormy seas: How fragile is liquidity across asset classes?

Nihad Aliyev [a], Matteo Aquilina [b,c], Khaladdin Rzayev [d,e,f], Sonya Zhu [b*]

[a] *University of Technology Sydney*
[b] *Bank for International Settlements*
[c] *Macquarie University*
[d] *The University of Edinburgh*
[e] *Koc University*
[f] *Systemic Risk Centre, London School of Economics*

Abstract

Liquidity has improved across global markets, but fragility concerns remain. We study the distribution of bid-ask spreads across equities, bonds, and foreign exchange (FX) in the US, Europe and Japan. While average and standard deviation of spreads have decreased since 1990s, skewness and kurtosis have increased, especially in bond and most equity markets, but not FX. We identify structural breaks in the mean and skewness and map them to macroeconomic events, market structure changes, and regulatory reforms. Simulations show that increased skewness raises trading costs—up to $1 billion annually in US equities—even when few trades require urgent execution.

JEL classification: G10, G12, G14
Keywords: Liquidity, trading cost, liquidity distribution, market fragility.

Date: May 14, 2025.
* 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 Thierry Foucault, Robert Gaudiosi (discussant), Denis Gorea, Shize Li (discussant), Marco Lombardi, Albert Menkveld, Benoit Mojon, Andreas Schrimpf, Vladislav Sushko, 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

Periods of market stress often reveal that liquidity, while abundant in normal times, can vanish abruptly when most needed. Such episodes—whether systemic, like the 2008 global financial crisis, or brief but disruptive, like the 2010 Flash Crash—highlight the fragile nature of liquidity. By *fragility*, we refer to the tendency of liquidity to dry up suddenly during periods of market stress, even when conditions appear stable just before. As traders often say, *"liquidity is a coward; it disappears at the first sign of trouble."*

Although liquidity is commonly viewed as a feature that lowers transaction costs, its significance extends far beyond pricing: it affects capital allocation, funding costs, and ultimately, the real economy. When liquidity dries up, the consequences can be far-reaching. In his influential book, Persaud (2003) argues that episodes in which liquidity disappears suddenly can "*destroy companies, cause significant economic contraction, bring down governments, rip the social fabric and steer capital away from certain markets more permanently*." These outcomes illustrate that liquidity is not just a market microstructure issue—it is a macroeconomic concern. Yet despite its central role, most of the literature has focused on average liquidity, and the distributional properties, particularly how liquidity behaves in the tails, remain underexplored.

In this paper, we go beyond average measures of liquidity and examine how higher-order moments of its distribution, specifically skewness and kurtosis, have evolved over time. Our aim is to understand how liquidity behaves not only in normal times but also in the tails during periods of stress when it matters most. We focus on three major asset classes: stocks, government bonds, and foreign exchange (FX). These markets play a critical role in capital allocation, monetary policy transmission, and risk management for governments, corporations, and investors. We study the distribution of bid-ask spreads (and other liquidity measures for robustness) across the key developed markets of the United States, Europe, and Japan.[1] To this end, we compile a high-frequency dataset of relative bid-ask spreads, which we aggregate into monthly and annual metrics to capture the distribution of liquidity over time, across asset classes, and across regions.

---

[1] The countries in our study collectively account for approximately 42% of global GDP and 57% of global market capitalization. See, for example, https://data.worldbank.org/indicator/CM.MKT.LCAP.CD. Europe in our sample is represented by the combination of Germany, France, UK, and Italy in different markets. In the FX market, we focus on USD exchange rates against the Euro, British Pound, and Japanese Yen, as currency pairs do not map directly to regions.

Our empirical analysis comprises three main components:

i. we document key facts about the distribution of liquidity in different asset classes and regions,
ii. we identify potential drivers of these facts, and
iii. we explore their implications for market participants.

We begin by documenting several empirical facts, with a particular focus on the first and third moments of the bid-ask spread (mean and skewness), which capture both typical liquidity conditions and how markets behave under stress. Over the past three decades, the average bid-ask spread, as well as the variation around that average (standard deviation), have decreased across all asset classes. Although the absolute and relative magnitude of the decline varies across asset classes and geographic regions, a consistent downward trend is observed across most markets.

In contrast, the higher-order moments of the distribution (skewness and kurtosis) tell a different story. In the US and Japanese stock markets, as well as in all government bond markets, both measures have increased over time, indicating more frequent and severe episodes of illiquidity. In European stock markets, however, skewness has decreased slightly, and kurtosis does not follow a clear pattern. FX markets also show a different pattern: skewness generally follows an inverted-U shape, and kurtosis fluctuates without a clear trend. This suggests that European stock and FX markets have improved average liquidity without increasing the likelihood of extreme illiquidity events.

Next, we identify potential drivers of these facts. We do this in two ways. First, we detect breaks in the time series of the mean and skewness of bid-ask spread using the Bai-Perron multiple structural break test (e.g., Bai and Perron, 1998, 2003). This test allows us to endogenously identify breakpoints by dividing the time series into segments within which mean or skewness remain statistically stable, and flags points where they change significantly. We apply this procedure separately to the mean and skewness of spread for each asset class and region.

Having statistically identified the structural breakpoints in the mean and skewness of liquidity, we then map these breakpoints to real-world events to better understand the potential drivers. While it is difficult to establish clear causal relationships, many breaks in the mean of bid-ask spreads coincide with macroeconomic events, such as the global financial crisis (GFC) and Abenomics (a pro-growth policy package introduced by Prime Minister Shinzō Abe starting in late

2012) in Japan. In contrast, breaks in skewness of spread are more often associated with changes in market structure, such as the rise of high-frequency trading, the introduction of Autoquote in equity markets, and the rollout of Application Programming Interfaces (APIs) in FX markets, and to a lesser extent, with regulatory reforms, including Reg NMS, MiFID II, and the Global FX Code. These findings suggest that changes in the average level of liquidity tend to be associated with macroeconomic conditions, whereas changes in its fragility (captured by skewness) are more associated with market design and regulatory developments.

To further examine potential drivers of fragile liquidity, we construct a comprehensive panel dataset for equity, bond, and FX markets, and estimate monthly regressions of the mean and skewness of bid-ask spreads on a range of market characteristics. In equity markets, we focus on two major developments: the rise of algorithmic trading (AT) and increasing market fragmentation. Both are associated with a decrease in average spreads, consistent with the idea that technological advancements and competition lower trading costs and improve liquidity. However, we also find that AT and fragmentation are associated with higher skewness, suggesting a connection to more fragile liquidity conditions.

We observe a similar pattern in FX markets, where AT is associated with lower average spreads and higher skewness. In contrast, we find no significant relationship between AT and skewness in bond markets, consistent with the fact that AT is less prevalent in bond markets. Across all asset classes, volatility—both at the instrument level (mid-price return volatility) and market level (e.g., VIX, MOVE, or JP Morgan FX volatility index)—is a strong predictor of both mean and skewness. Taken together, these findings suggest that while trading activity and volatility are key drivers of average liquidity, changes in market structure and technology have contributed to the growing fragility of liquidity.

Finally, we examine the implications of increasing skewness of bid-ask spreads for market participants. To do this, we develop a simulation model calibrated to US equity market data. The model features a simplified trading environment where spreads follow a skewed distribution. This allows us to vary skewness while keeping the mean and standard deviation constant. At regular intervals, a representative trader faces a potential trade opportunity and is randomly classified as either patient or impatient. If classified as patient, the trader executes only when spreads are favourable; if impatient, the trader executes immediately, regardless of market conditions.

Simulation results show that changes in the skewness of spreads can have meaningful effects on trading costs. We find that moving from the low-level of skewness of the late 1990s to the high-level of US large cap equities in 2023, the trading cost difference between patient and impatient trades becomes significantly larger. In the low-skewness environment, the premium paid for immediate execution is around 6%, but it more than doubles to 13% in the high skewness environment. When the probability of being impatient is low, overall trading costs remain similar across different skewness levels as traders can wait for the spread to revert. However, as the probability of impatience increases—a condition more likely during market stress—trading costs increase quickly in high-skewness regime.

These findings suggest that fragile liquidity imposes substantial costs on traders, especially when they are extremely impatient. Using simulations calibrated to 2023 market conditions, we estimate that increased skewness in US equity spreads adds approximately $400 million in annual trading costs if 20% of volume requires immediate execution, and up to $1 billion if that share rises to 30%. These magnitudes imply that even moderate levels of urgency well within the range observed in today's algorithm-driven, high-speed markets can translate into significant cost burdens when liquidity becomes skewed.

Our findings have important implications for market practitioners and regulators. Even when average spread conditions appear stable, an increase in skewness can increase trading costs substantially for those who require immediate execution. For investors, incorporating skewness into their trading strategies may help reduce trading costs during volatile periods. For regulators, monitoring changes in the higher-order moments of liquidity can provide early warning signs of stress in financial markets.

Several recent papers explore higher-order moments of liquidity. Roll and Subrahmanyam (2010) are among the first to document that bid-ask spreads of US equities has become more right-skewed over time, attributing this to increased competition among algorithmic market makers. Kim and Na (2018) show that skewness and kurtosis in liquidity are priced in the cross-section of stock returns. Aliyev et al. (2025) further argue, through asset pricing tests, that relying on average liquidity measures can give a misleading impression of market liquidity. Their results on the skewness as a more direct measure of fragility are consistent with findings of Pástor and Stambaugh (2003) and Acharya and Pedersen (2005), who show that investors demand a premium for bearing liquidity risk (see also Ben-Rephael et al., 2015; Li et al., 2019). However, all these studies focus primarily on equity markets.

Theoretical models have had limited success in explaining the distribution of liquidity. Foucault et al. (2005) provide one of the few frameworks and link the distribution of bid-ask spread to waiting costs and the mix of patient versus impatient traders, as in our simulations. In their model, skewness increases when the market becomes dominated by impatient traders. However, their model cannot account for the simultaneous decrease in average spreads and increase in skewness observed in data. Other studies such as Baruch and Glosten (2019), Jovanovic and Menkveld (2022), and Aliyev et al. (2022) incorporate liquidity risk more explicitly, by modeling undercutting, bidding costs, and uncertainty about the composition of traders.

This paper makes three main contributions to the literature. First, to our knowledge, we are the first to document how the distribution of liquidity has evolved across three major asset classes (equities, government bonds, and FX) in the major global markets of the US, Europe, and Japan, drawing on a dataset of over 2 billion high-frequency observations. This broader perspective is important because it allows us to compare patterns across markets that differ in structure, participants, and regulation. For example, FX markets—where average spreads have decreased but skewness has not increased—may offer useful insights into mechanisms that preserve the resilience of liquidity. Understanding these differences can help identify which aspects of market design contribute to more stable liquidity and what structural features one market might adopt from another to improve the resilience of liquidity.

Second, we go beyond describing trends and identify potential drivers of both mean and skewness of liquidity and link them to macroeconomic events, market structure, and regulation across different markets and regions. This allows us to disentangle the factors associated with improvements in average liquidity from those contributing to its increased fragility. Third, we quantify the trading cost implications of skewed liquidity using simulations. We show that higher skewness in spreads, even when average conditions are stable, can increase costs for traders who need to execute quickly. Together, our findings provide new evidence on what makes liquidity fragile across different markets and why that matters for market participants.

The rest of the paper is structured as follows. Section 2 introduces our liquidity measure, describes the data, and presents summary statistics for the empirical analysis. Section 3 describes the evolution of the distribution of the bid-ask spread across asset classes and regions. Section 4 investigates potential drivers of changes in the distribution of liquidity using structural break test and regression analysis. Section 5 presents the simulation model. Section 6 focuses on alternative liquidity measures, and Section 7 concludes.

2. Data and summary statistics

Liquidity measures based on the cost of trading are commonly preferred in the literature because they capture the many dimensions of liquidity (e.g., Foucault et al., 2013). The most straightforward way to measure these costs is through the quoted spread, which represents the cost of buying and immediately selling an asset (or vice versa).[2] It is defined as the difference between the best ask and best bid prices. Since the dollar spread is not scale-invariant, it is usually normalized by the midpoint price to form the relative quoted spread: $s_{it} = \frac{Ask_{it} - Bid_{it}}{MidPrice_{it}}$.

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 both researchers and industry practitioners. The dataset includes trades and quotes from a number of real-time feeds across more than 500 trading venues and captures activity from all types of market participants, not just dealers.[3]

For the government bond and FX markets, we obtain bid and ask quotes at 1-minute intervals. For equities, because the number of stocks is very large (around 10 thousand), we use 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. Using higher frequency data would introduce

[2] Other common trading cost measures include the effective spread, realized spread, and price impact. While these measures provide additional insights into execution costs and market response to trades, they require detailed trade-level data not available for most instruments in our sample. However, we conduct robustness checks using alternative liquidity measures for US equity markets in Section 6.

[3] Tick History has been widely used in liquidity studies across equities (e.g., Comerton-Forde et al., 2019; Ibikunle et al., 2021; Werner et al., 2023; Aquilina et al., 2024), government bonds (Sakiyama and Kobayashi, 2018), and FX markets (Krohn and Sushko (2022))

more microlevel noise, making economic interpretations more challenging. For example, a trade that consumes all liquidity at the best quote can temporarily widen bid-ask spread, leading to an increase in skewness and kurtosis. However, the speed at which modern markets operate ensures that using 1- minute and 5-minute snapshots helps mitigate this effect. In today's markets, reaction times by fast participants are measured in fractions of a second. For example, O'Hara (2015) argues that in a high-frequency trading environment, five minutes is effectively a lifetime, and even five to fifteen seconds are often sufficient to capture the temporary effects of large orders. (e.g., Aquilina et al., 2018), show that even after relatively large price movements, liquidity is typically replenished within a minute in UK stock markets while Bessembinder et al. (2016) show that the temporary price impact of order imbalances in crude oil ETFs is almost entirely reversed within one minute.

We supplement the Tick History data with several additional sources. We source FX volatility indexes from JP Morgan and the TED spread and the VIX from FRED. To control for algorithmic trading (AT), we use the trade-to-quote ratio from LSEG. Since FX spot volume is not available, we proxy trading activity using FX futures volumes. For equity market fragmentation, we construct an inverse Herfindahl–Hirschman Index (HHI) based on trading volumes across 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.

2.2. Filters and data integrity analysis

To ensure data quality, we apply several filters. First, we restrict the sample to regular trading hours and remove records where the order book is crossed. We also check for significant data gaps: if a month has fewer than 9 trading days, we remove it from the sample. We also drop any year with more than 2 months of missing data. We only keep the time series for which at least 5 years of data are available. Finally, to remove the impact of outliers or data errors on our moments estimates, we trim the high-frequency bid-ask spreads at the 99% percentile.

Our main goal is to characterize the distribution of liquidity—with a specific focus on its higher-order moments—across key global financial markets. We group the sample into three major regions: the United States, Europe, and Japan. For equities and bonds, the US and Japan are selected given their out-sized importance in these markets. In Europe, we include German, French, and UK stocks, and German, Italian, and UK government bonds—reflecting the significance of French equities and Italian sovereign debt. 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.

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, which 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). 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 US 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). This pattern does not appear in other datasets focusing on dealer-to-dealer trades (Fleming and Ruela, 2020). We have investigated this issue with LSEG, who confirmed that the data are accurate. 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 US 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. Given the assurances from LSEG, the fact that our study encompasses transactions beyond those involving only dealers, and the likelihood that market participants traded on these quotes, we retain the data in our sample. However, in Section 4.5, we re-run regressions excluding this period to test robustness.

2.3. Summary statistics

We begin our analysis by presenting summary statistics for the distribution of the bid-ask spread across asset classes and regions. Table 1 reports the average moments by asset class and by country. For equities, we calculate the four moments for each trading day using the 5-minute bid ask spread and then group stocks into terciles based on their average market capitalization each year.

In the US, all four moments—mean, standard deviation, skewness, and kurtosis—are strongly related to size. 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). In contrast, skewness and kurtosis decrease with market capitalization. This relationship between market capitalization and the moments of the bid-ask spread is not unique to the US; it holds across all countries in our sample. Across all regions, 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. For US 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, US 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 have less volatility in the spread. Interestingly, the skewness and kurtosis of two-year government bonds are higher than those of ten-year government bonds, except in German bonds.

In the FX market, the costs of trading GBP, Yen, or Euro against the US 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 approximately 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 generally close to zero.

Table 1 here

3. The evolution of the distribution of liquidity

In this section, we examine how the distribution of liquidity has evolved over time across equity, FX, and government bond markets. We do this in two ways. First, we calculate the four daily moments of the high-frequency bid-ask spread—mean, standard deviation, skewness, and kurtosis—for each instrument. We then calculate annual averages of these daily moments to capture long-term trends across markets and regions. Second, we estimate trend regressions to assess whether these patterns are statistically significant.

3.1 Time-series of the distribution of liquidity

3.1.1 Equities

We start with equity markets in the US, Europe (France, Germany, and the UK), and Japan, which collectively account for more than half of global equity market capitalization. Figure 1 shows that average bid-ask spreads decreased substantially from 1996 to 2023 across all regions and firm size groups. Aggregating across regions, the spread for large-cap stocks fell from 60 bps in 1996 to 13 bps in 2023; for mid-caps, from 103 bps to 34 bps; and for small-caps, from 184 bps to 111 bps. The standard deviation also decreased, from 34 bps in 1996 to 19 bps in 2023.

Despite these improvements in average liquidity, bid-ask spreads remain positively skewed across all countries and size groups. Before the GFC, skewness increased across the US, Japan, and Europe. After the GFC, however, it continued to increase in the US, remained stable in Japan, and started to decrease in Europe. Skewness is generally higher for larger stocks, but the relative increase over time is more pronounced among smaller stocks. Kurtosis shows a similar but more extreme pattern. For example, among small-caps, kurtosis increased from 0.23 in 1996 to 6.13 in 2023, with even larger increases for mid- and large-cap stocks.

Figure 1 here

3.1.2 Foreign exchange

Figure 2 shows the distribution of bid-ask spreads for trading the Euro, Japanese Yen and British Pound against the US 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 US dollar accounted for nearly 50% of the total turnover in the foreign exchange spot market (McGuire et al., 2024).

Average bid-ask spreads in the FX market have approximately halved since the mid-1990s. For example, the spread for the Japanese yen decreased from around 6 bps in 1996 to about 3 bps by the mid-2000s. In dollar terms, a $100,000 transaction would have incurred a $60 spread in 1996 and a $30 spread in 2005. Following the GFC, average spreads increased by around one-third across all three pairs and have remained relatively flat since then. The standard deviation followed a similar pattern—decreasing until the GFC, then increasing gradually after 2015. This post-2015 increase is more pronounced for GBP/USD and EUR/USD, and less so for USD/JPY.

Higher-order moments in FX differ from those in equities. Skewness shows an inverted-U pattern, moving closely across all three currencies before 2014 with a generally increasing trend. After 2015, skewness for the GBP and EUR started to decrease. The Japanese yen followed a similar downward trend at first, but its skewness increased sharply in 2019 and has since remained high. Kurtosis was mostly stable over time but spiked around the GFC, then gradually fell into negative territory. After the COVID-19 outbreak, kurtosis started to increase again across currencies.

Figure 2 here

3.1.3 Government bonds

Figure 3 shows the distribution of bid-ask spreads in government bonds across the US, Europe (UK, Germany, and Italy), and Japan for 2-, 5-, and 10-year maturities. The moments of the bid-ask spread co-move closely across maturities. Average spreads generally decreased over time across all three regions. Within Europe, we observe steady improvements in the UK and German markets, and more volatile patterns in Italy, with sharp spikes during the GFC and the European sovereign debt crisis. Across all countries, 2-year bonds tend to be more liquid than 5- and 10-year bonds.

The standard deviation of bid-ask spreads has followed a long-term downward trend in the US and Europe. In Japan, however, it has been increasing since 2010. In the US, the downward trend reversed around 2016, when the standard deviation started to gradually increase. Among the countries we study, Japanese bonds have the highest standard deviation, while US bonds have the lowest. Japanese bonds also show a wider range of standard deviation across maturities. For example, in the US, the standard deviation for the 10-year bond is 0.63 bps, nearly twice that of the 2-year bond (0.38 bps). In Japan, the 10-year bond has a standard deviation of 2.08 bps, about five times higher than that of the 2-year bond (0.52 bps).

Higher moments of the bid-ask spread follow a different pattern. Before the GFC, skewness was stable and fluctuated around zero. After the crisis, skewness increased across all regions and maturities, especially in the US and Japan. Following the COVID-19 shock, skewness has started to decline but remains positive. Kurtosis followed a similar path, increasing rapidly between 2015 and 2018 and only recently beginning to decline. While the mean and standard deviation show a clear relationship with bond maturity, skewness and kurtosis remain similar across maturities. This suggests that the drivers of the lower and higher moments may differ as the lower moments follow a clear term structure, while higher moments do not.

Figure 3 here

3.2. Trend regressions

To further investigate whether these observed patterns are statistically significant, we estimate trend regressions on the mean and skewness of the bid-ask spread. We focus on these two moments for two reasons. First, mean is highly correlated with standard deviation, and skewness is highly correlated with kurtosis, so including all four moments offers limited additional insight. Second, skewness provides a more meaningful signal for the asymmetry in the distribution of liquidity and serves as a direct measure of fragility. Because bid-ask spreads are bounded below by zero, increases in skewness signal a greater likelihood of extreme illiquidity events, while decreases suggest more symmetric and stable liquidity conditions.

To formally assess these trends, we estimate the following time-trend regressions:

$$Mean_t = \alpha + \beta X_t + \epsilon_t \tag{1}$$

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

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 estimate the model separately for each asset class and region, consistent with the time series figures (Figures 1, 2 and 3). Standard errors are computed using the Newey and West (1987) estimator with 12 lags to account for autocorrelation.

Table 2 here

Table 2 presents the regression results. Broadly, our findings confirm that the patterns documented in Section 3.1 are statistically significant. Across equity markets, the mean bid-ask spread declines significantly in all three regions. However, the skewness of the bid-ask spread increases significantly only in the US and Japan. In European equity markets, skewness shows no significant change in small and mid-cap stocks, with only a weakly significant decrease observed for large-cap stocks.

To understand why skewness evolves differently in European equity markets compared to those in the US and Japan, we need to consider the factors that drive skewness. In Section 4, we explore potential drivers of skewness. Our analysis suggests that the prevalence of AT/HFT, as well as market fragmentation, is positively correlated with skewness in equity markets. Other factors—such as trading volume, market capitalization, and volatility—also appear to play a role. When we compare these characteristics across regions, the most notable difference emerges in

AT/HFT activity—proxied by the frequency of quote updates. This measure is significantly higher in the US and Japan than in the EU. In contrast, we do not observe a similar cross-regional pattern in other variables that could explain the difference in skewness. This also aligns with industry reports and academic studies showing that HFT activity has consistently been more prevalent in US equity markets than in European equity markets (e.g., Bouveret et al., 2014).

In addition, while HFT activity arrived in the EU before Japan, HFT activity in Japan 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 decreasing 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 equity markets.

In government bond markets, the mean bid-ask spread decreases across most regions, although the decline is not statistically significant for European bonds with 2- and 5-year maturities. For European 10-year bonds, the trend is slightly positive and marginally significant at the 10% level. Skewness increases significantly in all regions, though the increase is relatively smaller in Japan. These results suggest that while average trading costs have improved, the likelihood of extreme illiquidity events has increased—particularly in US and European sovereign bond markets.

In FX markets, the trend in mean bid-ask spreads is negative and statistically significant, consistent with the time series patterns shown in Figure 2. However, we do not observe a statistically significant trend in skewness. As shown in Figure 2, FX skewness increased until around 2005, then stabilized and eventually declined. We interpret this non-monotonic pattern as reflecting structural changes in FX market design, particularly the introduction of two major APIs in the early 2000s—one for banks in 2004 and another for non-bank participants (principal trading firms and HFTs) in 2005. The main distinction between these APIs is their functionality. Bank APIs focus on automated execution of human-led trading decisions, while non-bank APIs enable full automation of both decision-making and execution, often using more aggressive strategies. This distinction likely influences skewness.

The initial increase in skewness coincides with the broader adoption of electronic trading, while the subsequent decline may reflect targeted mitigation efforts. Around 2013–2014, major FX platforms introduced speed bumps, batching, and randomization to reduce the impact of non-bank APIs and curb latency arbitrage and aggressive HFT activity (Chaboud et al., 2014). These interventions may help explain why FX markets—unlike equity and bond markets—do not show a persistent upward trend in skewness over time.

In summary, while the mean bid-ask spread has declined across almost all asset classes and regions, the evolution of skewness shows more heterogeneity. It has increased in most equity and bond markets, particularly in the US and Japan, but remained flat or even decreased in European equities and FX markets. Based on the available data, literature, and market structure insights, we argue that differences in the nature and intensity of electronic trading—especially HFT—are a key part of this story. However, we acknowledge that these interpretations are suggestive and should be viewed with caution, given the lack of formal theoretical models explicitly linking market design features to skewness in liquidity.

4. Drivers of changes in the distribution of liquidity

In this section, we explore what drives changes in the distribution of bid-ask spreads across asset classes over time. In particular, we aim to understand what makes liquidity more fragile across different asset classes. We take a three-step approach. First, we describe key changes in market design and trading technology that could plausibly affect the distribution of liquidity, particularly its skewness. Second, we use the Bai-Perron method to formally identify points in time when the mean or skewness of spreads changed significantly and examine whether these changes coincide with key macroeconomic events, changes in market structure, or regulatory reforms. Third, we run panel regressions to test whether market characteristics—such as volatility, algorithmic trading, or fragmentation—can explain changes in the mean and skewness of spreads across asset classes and over time.

4.1 Key changes in market design

Over the past three decades, financial markets have undergone significant structural and technological transformations. Among these, the shift from manual to electronic trading stands out as the most profound, as it triggered a range of subsequent changes, including regulatory interventions, and is also the most relevant for the purpose of our study.

4.1.1 Equities

The equity market was the first to undergo a major transformation. Beginning in the early 2000s, traditional voice and floor-based trading was gradually replaced by electronic platforms. While this transition was largely driven by advances in technology, it was also shaped by regulations aimed at improving market quality—typically by enhancing price discovery and reducing transaction costs. Although these changes succeeded in many respects, they also potentially altered the broader liquidity landscape. For example, while electronification enabled faster price discovery and reduced average transaction costs (Menkveld, 2016), it also introduced a degree of fragility, as liquidity could now vanish rapidly during stress events (Easley et al., 2012).

In equity markets, the combined effects of electronification and regulatory reforms brought two significant changes: the rise of algorithmic and high-frequency trading (AT/HFT) and increased market fragmentation. AT refers to the use of computer programs to automate order execution and trading decisions. Traders typically employ two types of algorithms: (i) those designed to execute large trades or adjust positions gradually over time with minimal market impact, and (ii) proprietary algorithms aimed at generating profits from short-term price fluctuations—commonly referred to as high-frequency trading (HFT) strategies (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. To achieve this, they invest heavily in technologies such as microwave networks and co-location with exchange servers (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).

Fragmentation, on the other hand, refers to the proliferation of trading venues. Regulatory changes such as Reg NMS in the US and MiFID in Europe lowered barriers to entry and encouraged competition among exchanges and alternative platforms (Menkveld, 2016). As a result, trading activity dispersed across multiple exchanges and off-exchange venues and traditional exchanges dramatically lost market share. For example, the NYSE's share of trading volume in its own listings dropped from 82% in 2004 to 27% in 2018 (Baldauf and Mollner, 2021).

In parallel, off-exchange trading grew significantly, now representing over 40% of US equity volume (Rosenblatt Securities, 2023). In Europe, traditional exchanges regained market share in the 2020s following the introduction of dark volume caps under MiFID II and the growing use of periodic and closing auctions by institutional investors (Hagströmer, 2022). While fragmentation improved access and competition, it also increased complexity and reduced transparency, particularly in off-exchange venues that do not display pre-trade information. These changes have contributed to greater asymmetry in the availability of liquidity (Degryse et al., 2015).

4.1.2 Foreign exchange

FX markets experienced a similar technological shift but followed a more market-driven path, given the lower degree of regulatory intervention relative to equity and sovereign bond markets. Initially dominated by interdealer voice trading, the FX market began to change in the early 1990s with the emergence of two electronic platforms—EBS and Refinitiv (formerly Reuters)—which introduced central limit order books for interdealer trading (Chaboud et al., 2023). In the early 2000s, multi-dealer platforms expanded the dealer-to-customer segment, allowing clients to request quotes from multiple dealers simultaneously. Additionally, the distinction between interdealer and dealer-to-client trading began to blur, and a broader set of financial institutions became active participants. Algorithmic activity accelerated in 2005, when EBS permitted non-bank participants to connect directly to its matching engine via APIs, which had been available to bank participants since the previous year.

By the 2010s, electronic trading had become the dominant mode of execution in most FX markets. While many institutional investors still rely on execution algos with human oversight, HFTs operate fully automated systems to arbitrage small price differentials across venues. This rise in HFT activity raised concerns over latency arbitrage and predatory trading, where faster participants exploit slower ones (BIS Markets Committee, 2011; Budish et al., 2015). In response, several FX venues introduced mitigation mechanisms, including speed bumps, randomization, and order batching, beginning in 2013–2014. This is perhaps the most important difference between FX and equity markets in terms of how they approached electronification. FX markets were quicker to implement interventions aimed at curbing the potentially destabilizing effects of HFT activity.

4.1.3 Government bonds

The sovereign bond market has undergone a slower transition to electronic trading compared to equity and FX markets. A recent report by the FSB (2022) provides detailed insights into the structure of sovereign bond markets across the jurisdictions included in our sample, up to late 2021. A key finding is that while secondary markets in all countries maintain a division between interdealer and dealer-to-customer segments, the level of electronification varies widely. In countries like the US, UK, and Italy, interdealer trading is often conducted anonymously via central limit order books operated by interdealer brokers (IDBs) using electronic platforms, with HFT playing a meaningful role only in the US. In contrast, Japan and Germany, still rely heavily on voice trading even among dealers, showing that manual execution remains common in some jurisdictions.

The dealer-to-customer segment remains heavily intermediated by dealers, with trading protocols dominated by request-for-quote systems and traditional voice trading. Most of the dealers active in this space are affiliated with banks and provide a broad range of services to their clients. Non-bank dealers are present but constitute only a small portion of total market activity. In many jurisdictions, dealers are formally registered with the National Debt Management Office, which confers specific obligations and privileges, particularly concerning participation a central role in primary issuance. HFTs have very limited involvement in the US dealer-to-customer segment and virtually no presence elsewhere. Overall, both the degree of electronification and the participation of HFTs in sovereign bond markets remain considerably lower than in equity and FX markets in our sample.

4.2 Breaks in the time series of mean and skewness

The time series plots and trend regressions in Section 3 show how the mean and skewness of bid-ask spreads have changed over time. While trend regressions help capture gradual trends formally, the time series plots suggest that some changes in liquidity are sudden—possibly triggered by shocks to the macroeconomy, market structure, or regulation. However, visual inspection alone cannot confirm whether these changes are statistically significant or precisely when they occur.

To formally identify and date discontinuities in liquidity, we test for structural breaks in the time series of the mean and skewness of bid-ask spreads. We use the method developed by Bai and Perron (1998, 2003), which estimates linear regression models and searches for multiple breakpoints—points where the underlying pattern in the data changes in a statistically significant way. This approach is flexible and can be used even when the number of breaks is unknown, as is the case in our analysis.

The method works by dividing the time series into segments and estimating the model separately within each segment. It then tests the null hypothesis that the regression coefficients remain constant over time against the alternative that they change at one or more unknown breakpoints. To determine whether a structural break has occurred, the method compares different ways of splitting the sample and selects the partition that minimizes the total residual sum of squares. Put differently, it compares a model with $m - 1$ breaks to one with m breaks and chooses the version that fits the data better, based on overall model fit. The general form of the model is:

$$Y_t = \beta X_t + \delta_j Z_t + \epsilon_t, \qquad \text{for} \quad T_{j-1} < t \leq T_j, \quad j = 1, \dots, m+1 \tag{3}$$

where Y_t is the variable of interest—either the mean or skewness of the bid-ask spread, X_t includes variables whose effects (β) are assumed to be constant over time, and Z_t contains variables whose effects (δ_j) are allowed to change at unknown breakpoints T_j.

In our application, we estimate a simple version of this model known as a *mean-shift* model. That is, we only include an intercept in the regression and allow this intercept to change at one or more points in time. This allows us to capture sudden level changes in the mean or skewness of the spread, without requiring explanatory variables. For example, suppose the skewness of spread is stable for several years but then jumps in mid-2008 due to a market shock. The model would estimate one constant before 2008 and a different one afterward, if doing so improves the overall fit of the model.

We estimate separate models for each of the three regions (US, Europe, and Japan) and each sub-asset class—FX; large-, mid-, and small-cap equities; and 2-, 5-, and 10-year government bonds—resulting in a total of 21 specifications. All models are estimated using monthly data. While the method returns a precise breakpoint at the monthly level, we aggregate the results to half-year intervals for clarity and comparability across series. We also record the direction of the shift (upward or downward) in each case. This approach allows us to systematically detect significant changes in the mean and skewness of spreads and to compare their timing across asset classes and regions.

Tables 3 and 4 here

The results of our structural break analysis are summarized in Tables 3 (for the mean) and 4 (for skewness). Broadly, the estimated timing and direction of these breaks are consistent with the patterns observed in the annual time series plots discussed in Section 3. Downward shifts in the mean bid-ask spread appear in the early 2000s for large-cap equities and FX, with further declines observed across asset classes in the mid-2000s. Upward shifts in spreads for small-cap equities and some government bonds are identified in the late 2000s (around GFC), with further increases observed across asset classes from 2015 onwards. For skewness, we observe increases in equity and FX markets starting in the early 2000s, followed by similar increases in government bond markets in the US and Japan, as well as in US equity markets, during the mid to late 2010s. In contrast, skewness declines in European and Japanese equity markets between 2010 and 2013 and in most FX markets from 2014 onward.

The natural next step is to ask whether these identified breaks can be linked to specific events or developments. Establishing clear causal links is challenging for two main reasons. First, multiple factors—including market structure, regulation, and macroeconomic conditions—can interact to affect liquidity. Second, break tests identify statistically significant changes, but they do not explain what caused them. Nonetheless, examining the historical context around the break dates can offer suggestive evidence and help generate hypotheses for future work. To this end, we conduct a literature review to identify major changes in regulation, market microstructure, and macroeconomic conditions that coincide with the breaks. Table 5 summarizes these events.

Table 5 here

Several breaks coincide with important changes in market microstructure. For example, the decimalization of stock prices and the introduction of Autoquote in the early 2000s are associated with simultaneous reductions in average spreads and increases in skewness. Decimalization refers to the change from quoting stock prices in fractions (e.g., 1/8 or 1/16 of a dollar) to decimal units ($0.01 increments), which increased price competition and contributed to narrower spreads. Autoquote was a software tool that automated quote dissemination and allowed market participants to receive pricing updates more rapidly benefiting AT. Hendershott et al. (2011) use the adoption of Autoquote as an instrument for HFT activity to study its impact on liquidity.

Another development that coincides with a reduction in mean spreads and an increase in skewness is the introduction of the Euro. The Euro was launched as an electronic currency in January 1999 and became fully operational with the introduction of physical notes and coins in January 2002. While often viewed as a macroeconomic event (due to the unification of monetary policy across the Eurozone), the euro also brought significant market microstructure changes by replacing multiple national currencies with a single, deep currency market that concentrated liquidity and reduced cross-border frictions. The introduction of the Euro corresponds with changes in FX liquidity across all relevant currency pairs in our sample, with average spreads decreasing and skewness increasing—particularly for the EUR/USD pair following the launch of physical currency. In FX markets, the introduction of API connectivity for both bank and non-bank participants is associated with significant declines in both the mean and skewness of spreads, while later measures, such as speed bumps to limit latency arbitrage, appear to have slightly raised mean spreads.

Regulatory changes also appear around several key breakpoints. These include MiFID I and MiFID II in Europe, Reg NMS in the US, the Financial Instruments and Exchange Act in Japan, and the development of the FX Global Code. In the US, the decline in average spreads predates Reg NMS, but the increase in skewness, particularly in mid-cap equities, occur around its implementation, suggesting that increased fragmentation may have contributed to the increase in skewness. In Europe, the impact of MiFID I is harder to isolate as it mostly coincides with the GFC. However, MiFID II is associated with increases in both average spreads and skewness in some bond segments. In FX markets, the publication of the FX Global Code in 2016–2017 is preceded by an upward break in spreads (in late 2015) and followed by a decline in skewness.

Macroeconomic shocks and policy responses represent a third category of events that coincide with breaks. During the GFC, average spreads increased in European bond markets but remained more stable in the US and Japan. Skewness increased in European and Japanese bond markets but was largely unchanged in the US. In equity markets, the average spreads increased for small caps across all regions and for Japanese mid-caps. At the same time, we observe upward breaks in skewness for US large caps and Japanese small caps, and downward breaks in European equities.

Other important turning points include Mario Draghi's "whatever it takes" speech in 2012, which helped stabilize European bond markets and is followed by a drop in average spreads but no change in skewness. In Japan, the introduction of Abenomics (a set of economic policies introduced by Japanese Prime Minister Shinzō Abe starting in late 2012) coincides with downward breaks in average spreads in Japanese equity markets, although we observe no major changes in bond markets or in skewness.

Taken together, a variety of factors, including structural, regulatory, and macroeconomic changes, coincide with the observed breaks in the distribution of bid-ask spreads. While macroeconomic shocks appear to drive many of the changes in mean spreads, changes in market structure and regulation are more closely associated with skewness. These findings suggest that market fragility, as captured by skewness, may not always move in tandem with average liquidity measures and is more sensitive to how markets are structured and governed. Understanding these distinctions is important for designing policies that promote not just liquid markets, on average, but also resilient ones.

4.3 Regression analysis of mean and skewness

In this section, we explore the potential determinants of mean and skewness of spreads by estimating panel regressions. We start our analysis by focusing on the equity markets and estimating the following regression model:

$$\begin{aligned} Mean_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Frag_{i,m} + \beta_3 Volume_{i,m} + \beta_4 MCap_{i,m} \\ +\beta_5 Volatility_{i,m} + \beta_6 VIX_m + \beta_7 TED_m + \varepsilon_{i,m} \end{aligned} \tag{4}$$

$$\begin{aligned} Skewness_{i,m+1} = \alpha_i + \beta_1 Algo_{i,m} + \beta_2 Frag_{i,m} + \beta_3 Volume_{i,m} + \beta_4 MCap_{i,m} \\ +\beta_5 Volatility_{i,m} + \beta_6 VIX_m + \beta_7 TED_m + \beta_8 Mean_{i,m} + \varepsilon_{i,m} \end{aligned} \tag{5}$$

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 use lagged values of independent variables to mitigate endogeneity concerns. We include only stock fixed effects because VIX_m and TED_m values are the same across different stocks for a given month. All variables are standardized, as we are interested in comparing the magnitude of the impact of each characteristic.

$Algo_{i,m}$ is the proxy for 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 use 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.

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). $Volume_{i,m}$ is the monthly (m) average of the daily total number of shares traded for stock i, representing overall trading activities. Market capitalization ($MCap_{i,m}$) 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 the skewness regression, we also control for the mean of the bid-ask spread to ensure the relationship between skewness and explanatory variables are not confounded by overall changes in spread levels.

Table 6 here

The results of the model in Equations (4) and (5) are presented in Panel A of Table 6. The results indicate that AT/HFT is negatively correlated with the mean of the bid-ask spread but positively correlated with its skewness. In other words, a greater presence of AT is associated with lower average spreads but with a more right-skewed bid-ask spread distribution, indicating more frequent occurrences of extreme spread values. The negative correlation between AT/HFT and the mean bid-ask spread is well-documented in the literature, suggesting 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, which leads to tighter spreads (Hendershott et al., 2011; Menkveld, 2013; Brogaard et al., 2015).

In contrast, the positive association between AT/HFT and the skewness of the spread is less understood, as there is little theoretical work directly addressing this relationship. An exception is Foucault et al. (2005), who link spread skewness to the composition of traders by their level of impatience. Their model suggests that markets with a higher share of impatient traders are less resilient, leading to right-skewed spread distributions. Fast markets may exhibit this characteristic: while generally more liquid, they may be less resilient in absorbing temporary imbalances. Our findings are consistent with this view that greater AT/HFT activity is associated with higher skewness in spreads. Another possible channel through which AT/HFT may affect skewness relates to the dual role of high-frequency traders as liquidity suppliers and demanders. For example, Aquilina et al. (2018) show that HFTs can amplify market stress by switching from supplying to demanding liquidity. Similarly, Brogaard et al. (2018) find that HFTs typically supply liquidity during extreme price movements in individual stocks but become net demanders of liquidity when price movements are more systemic. Such shifts can lead to sharp drops in liquidity provision, creating order imbalances that result in wider, more dispersed spreads and, ultimately, greater skewness.

The negative correlation between market fragmentation and the mean bid-ask spread is also in line with prior research. O'Hara and Ye (2011) and Degryse et al. (2015) show that (lit) fragmentation tends to reduce average spreads by intensifying competition among liquidity providers. However, the positive relationship between fragmentation and skewness is less explored. One explanation is the cross-subsidization channel discussed in Roll and Subrahmanyam (2010). In less competitive environments, market makers can maintain slightly high spreads when information asymmetry is low. This helps them absorb losses during high information asymmetry, smoothing out the distribution of spreads (i.e., fewer extreme spreads during high information asymmetry periods). In contrast, in highly competitive settings, such as fragmented markets, market makers are forced to price liquidity at all times. This limits their ability to cross-subsidize and may force them to post very wide spreads when adverse selection risk is high, leading to a more skewed distribution.

Supporting this mechanism, Van Kervel (2015) show that increased competition across trading venues can result in extreme illiquidity in one market, as trades on a given venue often trigger large-scale cancellations of limit orders on competing venues. This dynamic indicates the idea that fragmentation, while reducing average transaction costs, may increase the frequency of extreme illiquidity events—captured by increased skewness in bid-ask spreads.

In the second test, we focus on FX markets and estimate Equations (4) and (5) for the three currency pairs in our sample, with a few modifications. First, instead of using the VIX_m, we use the $JPVIX_m$, the JP Morgan FX Volatility Index for G10 currencies, which is specifically designed to capture volatility in FX markets (Ranaldo and de Magistris, 2022). To proxy trading volume, we use FX futures volume, as we do not have access to trade data for spot FX markets. Related to this, our proxy for AT in FX is constructed as the ratio of the number of quotes in the spot market (available through LSEG) to FX futures trading volume. We do not include a fragmentation proxy for FX markets, as FX trading does not operate with centralized exchanges with clear venue-level market share. Additionally, we do not include market capitalization, as this is not a firm-level analysis. All other variables are as previously defined.

The results are reported in Panel B of Table 6. The relationship between AT and the bid-ask spread moments in FX markets is consistent with the findings for equity markets: increases in AT are associated with lower average bid-ask spreads but higher skewness. This is interesting, especially since, unlike equities, the skewness of spread in FX does not increase over time.

In the final test, we explore the determinants of the mean and skewness of bid-ask spread in the government bond markets. 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 (e.g., Harkrader and Puglia, 2020), the extent of HFT in US bond markets is significantly smaller compared to that in equity and FX markets.

We estimate a modified version of Equations (4) and (5) for government bonds. First, as in the FX analysis, we use futures volume as a proxy for trading activity in government bond markets.

The only exception is Japanese bonds, for which spot volume data is available and used instead. Second, in addition to VIX_m, we include the $MOVE_m$ index to capture implied bond market volatility. Third, we restrict the sample to the post-2010 period, as the pre-2010 bid-ask spread data for the US bonds from LSEG is not fully consistent with earlier studies for the pre-2010 period (Fleming and Ruela, 2020). However, the results are qualitatively similar when the full sample is used. Finally, we do not include market capitalization, as this is not a firm-level analysis.

The results are reported in Panel C of Table 6. As expected, the association between AT and the bid-ask spread moments is weak and not statistically significant in government bond markets. In contrast, volatility is a significant explanatory factor for both the mean and skewness of the bid-ask spread. In particular, $MOVE_m$ is positively correlated with $Mean_{i,m+1}$ and negatively correlated with $Skewness_{i,m+1}$.

Overall, the results in this section show that two major changes in equity markets, AT/HFT and market fragmentation, increase skewness of bid-ask spreads and make liquidity more fragile. We observe a similar result for AT/HFT in FX. The association between AT/HFT and the skewness of bid-ask spread is much weaker in bond markets, which is consistent with the fact 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. Implications for trading costs

The observed increase in the skewness of the bid-ask spread in some markets suggests that traders today face a higher likelihood of experiencing periods of extreme illiquidity. This has at least two implications. First, market resiliency may be undermined if traders, accustomed to stable and low spreads, are faced with sudden and unexpected illiquidity, potentially amplifying financial instability (Persaud, 2003). Second, such episodes impose direct economic costs on market participants. In this section, we focus on the second issue and use a simulation to quantify these costs for a representative trader.

5.1 Simulation model

We simulate a simplified trading environment in which bid-ask spreads are drawn from a gamma distribution. This distribution is well-suited for our purposes: it fits the empirical spread data reasonably well and allows us to vary skewness while keeping the mean and standard deviation constant. The results are consistent when using other skewed distributions, such as the skew-normal distribution. We calibrate the simulation using US large-cap equity market data, and set the mean spread to 17 bps and the standard deviation to 11 bps.

We consider three skewness regimes: low (0.82), medium (1.42), and high (2.7). These values are derived by adjusting the shape parameter of the gamma distribution (and recalibrating the scale parameter to keep the mean and standard deviation constant) to span a plausible range from nearly symmetric to heavily right-skewed distributions. The low-skewness regime captures to conditions observed in US large-cap stocks earlier in our sample, while the high-skewness regime reflects the level seen in 2023. The medium-skewness regime captures intermediate market conditions. Spreads are drawn every minute between 9:30 and 16:00 (450 observations per day), and trade opportunities arise every 10 minutes. The results are not sensitive to this assumption. For example, allowing trades every 5 minutes obtains similar findings.

Although trade opportunities arise at regular intervals, not all result in execution. At each interval, a single representative trader arrives, with either patient (with probability $1-\mu$) or impatient trading demand (with probability μ), where μ represents the probability of requiring immediate execution. This setup is similar in spirit to the sequential trading model of Glosten and Milgrom (1985), where trader types are sequentially drawn from a fixed distribution. But in our setting rather than modelling many agents, we use a single trader whose behaviour probabilistically captures the presence of both patient and impatient participants in the market. This allows us to isolate the cost implications of execution urgency while keeping the model tractable.

If the trading demand is impatient, the trader executes immediately regardless of the prevailing spread. If demand is classified as patient, the trader observes the current bid-ask spread and trades only if it is below a predetermined threshold, defined as the median spread plus 2.57 standard deviations. If the spread exceeds this threshold, the trader refrains from trading. This setup captures realistic trading behaviour—investors with flexibility wait for more favourable conditions, while those with urgent needs execute trades regardless of cost.

The trader starts with no position and alternates between buy and sell orders. If a buy order is executed—either because the trader is impatient or because the spread is below the pre-specified threshold for a patient type—the next eligible trade order will be a sell, and vice versa. Each executed trade incurs a cost equal to the prevailing spread times a fixed trade size.

This setup lends itself to a simple analytical interpretation that clarifies the role of skewness. Suppose the bid-ask spread S follow a positively skewed distribution with mean $E[S]$, and let τ denote a threshold spread below which a patient trader is willing to execute. When the trader has impatient trading demand, they execute at every opportunity and thus incur an expected cost of $E\left[\text{Cost}_{\text{impatient}}\right] = E[S]$. In contrast, when the trader has patient trading demand, they execute only when $S \leq \tau$, and hence the expected cost conditional on trading is $E\left[\text{Cost}_{\text{patient}}\right] = E[S|S \leq \tau]$. The difference,

$$\Delta C = E\left[\text{Cost}_{\text{impatient}}\right] - E\left[\text{Cost}_{\text{patient}}\right] = E[S] - E[S|S \leq \tau], \tag{6}$$

captures the premium paid for immediacy. Because the distribution is right-skewed, rare but extreme spreads increase $E[S]$ far above the conditional mean $E[S|S \leq \tau]$, so ΔC increases with skewness. This analytical derivation underpins the simulation framework and guides our interpretation of the simulation results.

The simulations allow us to isolate how skewness in the distribution of spreads affects transaction costs under different market compositions. If the trader were always patient, skewness would be largely irrelevant—wide spreads could simply be avoided. However, when immediacy is sometimes required, right-skewed spreads lead to higher trading costs, as the probability of encountering unusually wide spreads increases. This also aligns with the theoretical insights of Foucault et al. (2005), where impatient traders submit market orders to obtain immediacy, and with empirical findings from Menkveld et al. (2017), who show that time-sensitive traders often choose venues with more reliable execution but higher transaction costs. The simulations enables us to quantify these urgency-driven costs and shows how skewness can amplify the burden on traders.

5.2 Simulation results

5.2.1 Skewness, impatience and trading costs

To evaluate the impact of skewness on trading costs—measured as the prevailing bid-ask spread prior to trade execution—we employ two key metrics. First, we calculate the percentage difference in trading costs between impatient and patient trades, which captures the premium paid for immediacy across different skewness regimes. To calculate the first metric, we run simulations across a range of impatience probabilities (μ) and, for each of the three skewness regimes, calculate the cost per share for impatient trades and compare it with the corresponding average for patient trades. The resulting difference between impatient and patient trades captures the premium paid for immediacy, which we then average across all μ to obtain an estimate for each skewness regime.

The second metric is the percentage increase in trading costs as the probability of impatience increases relative to a benchmark with low impatience probability ($\mu = 0.1$) in three different skewness regimes. For each skewness regime, we simulate trading by varying the probability of impatience (from $\mu = 0.1$ to $\mu = 0.9$ with an increment of 0.1) and compare the resulting average cost per share to that in the low-impatience benchmark ($\mu = 0.1$). This approach quantifies the cumulative cost of increased impatience under different skewness regimes and allows us to investigate whether trading costs increase uniformly with impatience, or whether they increase more sharply in markets with higher skewness. While the first metric summarizes the average cost of immediacy, the second metric shows how the cost burden builds due to skewness as the market becomes more impatient. For both measures, the results are averaged across 1,000 simulation runs and reported with 95% confidence intervals.

Figures 4 and 5 here

As reported in Figure 4, the premium paid for immediate execution increases monotonically with skewness. In the low-skewness regime, the average difference in trading costs between impatient and patient trades is 6.3%. This increases to 10.3% in the medium-skewness regime and to 13.1% in the high-skewness regime. The premium for immediacy more than doubles from low to high skewness regime. These results indicate that as the distribution of bid-ask spreads becomes increasingly skewed, the cost of immediate execution increases sharply. In each skewness regime, the cost difference between impatient and patient trades is statistically significant, based on standard errors of the cost differences across 1,000 simulations per regime. Furthermore, the size of this cost difference (the premium paid for immediacy) varies significantly across regimes

($p < 0.01$ for all pairwise comparisons).[4]

Figure 5 illustrates how the second metric evolves as μ increases relative to the low-impatience benchmark ($\mu = 0.1$), and how this relationship varies across different skewness regimes. At low levels of impatience, trading costs remain relatively stable across skewness regimes, consistent with the idea that skewness does not materially affect overall trading costs. In such cases, traders can wait for spreads to return to more favorable levels and avoid transacting during periods of extreme illiquidity. However, as the level of impatience grows, trading costs increase sharply in high skewness regime. For example, when $\mu = 0.9$, trading costs increase by approximately 9.6% in the high-skewness regime, compared to only 2.6% in the low-skewness regime, relative to the low-impatience benchmark.

These results can be understood analytically. Since with probability μ the trader has impatient trading demand and executes at the prevailing spread, and with probability $1 - \mu$ has patient demand and waits for $S \leq \tau$, the expected cost is

$$\begin{aligned} E[\text{Cost}] &= \mu \cdot E\big[\text{Cost}_{\text{impatient}}\big] + (1-\mu) \cdot E\big[\text{Cost}_{\text{patient}}\big] \\ &= \mu \cdot E[S] + (1-\mu) \cdot E[S|S \leq \tau]. \end{aligned} \tag{7}$$

Differentiating with respect to μ obtains

$$\frac{\partial E[\text{Cost}]}{\partial \mu} = E[S] - E[S|S \leq \tau] = \Delta C. \tag{8}$$

That is, the marginal cost of increased impatience equals the premium paid for immediacy, which increases with skewness. This explains why trading costs increase more sharply with μ when skewness is high, as shown in Figure 5. These results hold under broad conditions and are not specific to any one distribution. The simulation illustrates them using the gamma distribution and allows us to quantify the effects under empirically calibrated market conditions.

These findings suggest that skewness increases the average cost of immediacy for traders and increases the cumulative cost of traders as their impatience increases. Although the simulation abstracts from many real-world complexities, it carries important implications. Even when the

[4] Statistical significance is tested using independent-samples t-tests comparing the mean percentage cost differences between impatient and patient trades across skewness regimes. For each regime, we run 1,000 simulations and compute the cost difference between impatient and patient trades in each run. The average of these 1,000 values gives the mean cost difference for that regime. We then test whether these regime-level averages differ significantly across regimes. Pairwise comparisons are conducted between Low–Medium, Medium–High, and Low–High skewness regimes, under the null hypothesis that the mean differences are equal.

mean and standard deviation of bid-ask spreads remain fixed, increased skewness can significantly increase execution costs—particularly for traders who are unable to delay execution. This is especially relevant because the level of impatience tends to rise during periods of market stress or ahead of major information releases (e.g., Menkveld et al., 2017), precisely when liquidity is most needed. Our results suggest that in such conditions, high skewness can exacerbate market fragility.

5.2.2 Total amounts at stake

Using the information in Figure 5, we can also estimate the dollar value of trading costs attributable to bid-ask spread skewness across different levels of trading impatience. Using the average bid-ask spread of 17 bps observed in our data for the US large cap stocks as the base spread (corresponding to $\mu = 0.1$) and applying it to the \$75.5 trillion annual trading volume reported for US markets in 2024 by the World Federation of Exchanges, we can quantify the economic impact at different impatience levels as $Base\,Spread \times \%Cost\,Increase \times Trading\,Volume$.

At a 20% impatience level, the additional cost in low skewness regimes is negligible (approximately \$0.7 million), while in high skewness regimes it reaches \$400 million. When impatience increases to 30%, the costs increase to \$300 million and \$1 billion in low and high skewness regimes, respectively. At a 50% impatience level, costs increase further to \$0.8 billion and \$3.3 billion in low and high skewness regimes, respectively, resulting in a cost difference of \$2.5 billion.

These figures indicate that even at a 20% (resp. 30%) impatience level, the cost associated with increased skewness, holding the mean and standard deviation of the bid-ask spread constant, can amount to approximately \$0.4 (resp. \$1) billion under current skewness conditions observed in US large-cap stocks. Importantly, the assumed 20% or 30% share of trading requiring immediate execution is not only plausible but may in fact be conservative. For example, Foucault et al. (2005) argue that arbitrageurs, technical traders, and indexers, who attempt to replicate the return on a specific stock or index, often act as impatient traders by aggressively demanding execution. Basak and Croitoru (2006) estimate that stock index arbitrage alone may account for 7% to 9% of total trading volume. More recently, Aquilina et al. (2022) document that 22% of trading in FTSE 100 stocks arises from high-frequency latency arbitrage, a different type of impatient demand and show in sensitivity analyses that this share can be as high as 44%.

Taken together, our findings offer relevant insights for both investors and policymakers. For traders, they highlight the importance of considering not just the average level of spreads, but also higher-order moments when designing trading strategies. Algorithms that dynamically adjust execution parameters in response to changing skewness could help reduce trading costs. For regulators, tracking the skewness of bid-ask spreads may serve as an early indicator of declining market quality, complementing traditional measures focused on central moments like the mean.

6. Alternative liquidity measures

So far, we have focused on the relative quoted spread as our primary liquidity measure. This allows for consistent comparisons across asset classes and avoids the data limitations that often constrain alternative measures. However, quoted spreads have some well-known drawbacks. They may not fully capture true trading costs, especially for larger orders that go beyond the best available quotes. In such cases, execution may happen deeper in the order book, making trading more expensive than the displayed spread suggests.

In addition, quoted spreads can be affected by short-lived spikes. One might argue that the observed increase in skewness reflects temporary spikes in quoted spreads following large market orders that consume available depth at the best quotes. In such cases, the distribution of bid-ask spreads may appear more right-skewed. The skewness would be less concerning if spreads revert quickly through competitive quoting, allowing subsequent trades to occur at narrower levels.

For one, the results of our simulation show that this is only true at low levels of impatience. In addition, this concern is partially mitigated by our sampling frequency—five-minute intervals for equity markets and one-minute intervals for FX and government bond markets—which reduces the influence of transient fluctuations. O'Hara (2015), Aquilina et al (2018) and Bessembinder et al. (2016) show that our sampling window is wide enough to allow spreads to normalize, which reduces the likelihood that short-lived spikes drive our skewness results.

In any case, to further address these limitations, we complement our main analysis with the effective spread measure and its components. Due to data availability, this additional analysis is restricted to the US equity market. The effective spread captures actual trading costs by comparing execution prices to the prevailing mid-quote at the time of the trade. It can be decomposed into two parts: *price impact* (pi), which reflects how much prices move in the direction of the trade and is typically associated with adverse selection costs faced by liquidity providers (Brogaard et al., 2015), and *realized spread* (rs), which represents the portion of the effective spread that liquidity providers retain after accounting for price impact.

Calculating pi_{it} and rs_{it} typically requires transaction-level data to determine trade direction and measure execution costs accurately. However, relying on such granular data presents two key challenges in our setting. First, trade-level data would capture temporary liquidity fluctuations that reverse quickly, whereas our focus is on persistent patterns in liquidity, as discussed earlier. Second, processing transaction-level data for the entire US equity market over the 1996–2022 period would require substantial computational resources. To address both issues, we use five-minute data and modify the calculation by using the absolute value of the price difference in the numerator. This approach ensures the resulting measures consistently reflect the positive cost incurred by liquidity demanders. We therefore calculate the two components as follows:

$$pi_{it} = \frac{|Midprice_{it+5} - MidPrice_{it}|}{MidPrice_{it}} \tag{9}$$

$$rs_{it} = \frac{|TradePrice_{it} - Midprice_{it+5}|}{MidPrice_{it}} \tag{10}$$

where pi_{it} and rs_{it} are calculated over a five-minute interval ($t + 5$). The effective spread is the sum of pi_{it} and rs_{it} (Conrad and Wahal, 2020).

Figure 6 here

As shown in Figure 6, all three measures follow a similar pattern to that of the quoted spread: decrease in mean and increase in skewness. The results are also statistically significant. The highest p-value among the trend regressions is 0.0651, corresponding to the skewness of the realized spread. For all other measures, changes in both mean and skewness over time are statistically significant at least at the 5% level. This consistency across different liquidity measures supports our interpretation of the quoted spread findings. While average trading costs have fallen, markets are increasingly characterized by episodes of severe illiquidity, marked by large price impacts and realized spread. Moreover, the consistent increase in skewness for realized spreads and, in particular, price impact at the five-minute level suggests that these patterns are not driven by short-lived quote fluctuations.

7. Conclusion

Financial markets have become more liquid on average, but they have also become more fragile. This paper investigates this growing disconnect between average liquidity and its fragility. We examine the evolution of the distribution of liquidity over the past three decades across three major asset classes (equities, government bonds, and FX) in the world's largest markets (US, Europe, and Japan). We confirm that, on average, most markets have become significantly more liquid, with narrower bid-ask spreads. However, in most equity and government bond markets, there has been a noticeable increase in the skewness of bid-ask spreads, suggesting that episodes of extreme illiquidity have become more common. In contrast, the FX market has managed to reduce the skewness of bid-ask spreads.

To understand the drivers of these patterns, we conduct structural break tests and examine potential links with regulatory, structural, and macroeconomic developments. We observe that many breaks in the mean of bid-ask spreads coincide with macroeconomic shocks, while breaks in skewness are more often associated with changes in market structure, such as the rise of AT, increased fragmentation, and to a lesser extent, regulatory reforms. To further explore these relationships, we run regressions across all markets. In equity markets, AT and market fragmentation are associated with lower average spreads but higher skewness. A similar association is observed for AT in FX markets. In government bond markets, where AT activity is more limited, we find no systematic relationship between AT and the distribution of bid-ask spreads.

Importantly, our simulation results show that skewness is not just a statistical feature; it has real costs. Even when the mean and standard deviation of the bid-ask spreads are held constant, greater skewness leads to significantly higher trading costs, particularly for investors who require immediate execution. For instance, under the skewness conditions observed in US large-cap stocks and based on 2024 US equities trading volume, annual trading costs are between $400 million and $1 billion when 20% to 30% of trades require immediate execution. This shows that skewness imposes costs on traders, especially during periods of market stress when immediacy becomes more important.

These findings have practical implications. For investors, trade execution is not just about average spreads but also about tail risk. Trading algorithms that ignore this dimension may underperform during periods of market stress. For policymakers, our findings suggest that monitoring higher-order moments of liquidity can provide early signals of fragility. Ultimately, improving average liquidity is not enough. The goal should be not only to build liquid markets, but also resilient ones, where liquidity holds up not just in normal times, but during stress, when it matters most.

References

Acharya, V.V., Pedersen, L.H., 2005. Asset pricing with liquidity risk. Journal of Financial Economics 77, 375-410

Aliyev, N., Brogaard, J., Putniņš, T.J., Rzayev, K., 2025. Have Financial Markets Become More Liquid?. Working paper.

Aliyev, N., He, X., Putniņš, T.J., 2022. Learning about adverse selection in markets. Working paper.

Aquilina, M., Budish, E., O'neill, P., 2022. Quantifying the high-frequency trading "arms race". The Quarterly Journal of Economics 137, 493-564

Aquilina, M., Eyles, B., Shao, J., Ysusi, C., 2018. How do participants behave during flash events? evidence from the UK equity market. Working paper.

Bai, J., Perron, P., 1998. Estimating and testing linear models with multiple structural changes. Econometrica, 47-78

Bai, J., Perron, P., 2003. Computation and analysis of multiple structural change models. Journal of Applied Econometrics 18, 1-22

Baldauf, M., Mollner, J., 2021. Trading in fragmented markets. Journal of Financial and Quantitative Analysis 56, 93-121

Baruch, S., Glosten, L.R., 2019. Tail expectation and imperfect competition in limit order book markets. Journal of Economic Theory 183, 661-697

Basak, S., Croitoru, B., 2006. On the role of arbitrageurs in rational markets. Journal of Financial Economics 81, 143-173

Ben-Rephael, A., Kadan, O., Wohl, A., 2015. The diminishing liquidity premium. Journal of Financial and Quantitative Analysis 50, 197-229

Bessembinder, H., Carrion, A., Tuttle, L., Venkataraman, K., 2016. Liquidity, resiliency and market quality around predictable trades: Theory and evidence. Journal of Financial Economics 121, 142-166

BIS Markets Committee, B., 2011. High-frequency trading in the foreign exchange market. Report.

Bouveret, A., Guillaumie, C., Roqueiro, C.A., Winkler, C., Nauhaus, S., 2014. High-frequency trading activity in EU equity markets, Working paper.

Brogaard, J., Carrion, A., Moyaert, T., Riordan, R., Shkilko, A., Sokolov, K., 2018. High frequency trading and extreme price movements. Journal of Financial Economics 128, 253-265

Brogaard, J., Hagströmer, B., Nordén, L., Riordan, R., 2015. Trading fast and slow: Colocation and liquidity. Review of Financial Studies 28, 3407-3443

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

Chaboud, A., Rime, D., Sushko, V., 2023. The foreign exchange market. In: Research Handbook of Financial Markets. Edward Elgar Publishing, pp. 253-275.

Chaboud, A.P., Chiquoine, B., Hjalmarsson, E., Vega, C., 2014. Rise of the machines: Algorithmic trading in the foreign exchange market. Journal of Finance 69, 2045-2084

Conrad, J., Wahal, S., 2020. The term structure of liquidity provision. Journal of Financial Economics 136, 239-259

Degryse, H., De Jong, F., Kervel, V.v., 2015. The impact of dark trading and visible fragmentation on market quality. Review of Finance 19, 1587-1622

Easley, D., López de Prado, M.M., O'Hara, M., 2012. Flow toxicity and liquidity in a high-frequency world. Review of Financial Studies 25, 1457-1493

Fleming, M., Ruela, F., 2020. Treasury market liquidity during the COVID-19 crisis. Working paper.

Foucault, T., Kadan, O., Kandel, E., 2005. Limit order book as a market for liquidity. Review of Financial Studies 18, 1171-1217

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

FSB, 2022. Liquidity in core government bond markets. Report.

Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. Journal of Financial Economics 14, 71-100

Hagströmer, B., 2022. Market fragmentation in Europe. Working paper.

Harkrader, J.C., Puglia, M., 2020. Principal trading firm activity in treasury cash markets. Working paper.

Hendershott, T., Jones, C.M., Menkveld, A.J., 2011. Does algorithmic trading improve liquidity? Journal of Finance 66, 1-33

Jovanovic, B., Menkveld, A.J., 2022. Equilibrium bid-price dispersion. Journal of Political Economy 130, 426-461

Kim, S., Na, H., 2018. Higher-moment liquidity risks and the cross-section of stock returns. Journal of Financial Markets 38, 39-59

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

Li, H., Novy-Marx, R., Velikov, M., 2019. Liquidity risk and asset pricing. Critical Finance Review 8, 223-255

McGuire, P., von Peter, G., Zhu, S., 2024. International finance through the lens of BIS statistics: the global reach of currencies. In: BIS Quarterly Review, pp. 1-16

Menkveld, A.J., 2013. High frequency trading and the new market makers. Journal of Financial Markets 16, 712-740

Menkveld, A.J., 2014. High-frequency traders and market structure. Financial Review 49, 333-344

Menkveld, A.J., 2016. The economics of high-frequency trading: Taking stock. Annual Review of Financial Economics 8, 1-24

Menkveld, A.J., Yueshen, B.Z., Zhu, H., 2017. Shades of darkness: A pecking order of trading venues. Journal of Financial Economics 124, 503-534

Newey, W.K., West, K.D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703-708

O'Hara, M., 2015. High frequency market microstructure. Journal of Financial Economics 116, 257-270

O'Hara, M., Ye, M., 2011. Is market fragmentation harming market quality? Journal of Financial Economics 100, 459-474

Pástor, Ľ., Stambaugh, R.F., 2003. Liquidity risk and expected stock returns. Journal of Political Economy 111, 642-685

Persaud, A., 2003. Liquidity Black Holes, London.

Ranaldo, A., de Magistris, P.S., 2022. Liquidity in the global currency market. Journal of Financial Economics 146, 859-883

Roll, R., Subrahmanyam, A., 2010. Liquidity skewness. Journal of Banking & Finance 34, 2562-2571

Rosenblatt Securities, S., 2023. Let There Be Light: US Edition #54. Report.

Rzayev, K., Ibikunle, G., Steffen, T., 2023. The market quality implications of speed in cross-platform trading: Evidence from Frankfurt-London microwave networks. Journal of Financial Markets 66, 100853

SEC, 2010. Concept release on equity market structure. Report.

Shkilko, A., Sokolov, K., 2020. Every cloud has a silver lining: Fast trading, microwave connectivity, and trading costs. Journal of Finance 75, 2899-2927

Van Kervel, V., 2015. Competition for order flow with fast and slow traders. Review of Financial Studies 28, 2094-2127

Vayanos, D., Wang, J., 2013. Market liquidity—theory and empirical evidence. In: Handbook of the Economics of Finance. Elsevier, pp. 1289-1361.

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

Figure 1. Time series of the moments of bid ask spread distribution in the equity market

This figure presents the mean, standard deviation, skewness, and kurtosis of 5-minute bid-ask spread distributions in the US, Japanese, and European equity markets. For each stock-day, we calculate the moments of relative spreads and average them annually. Each year, stocks within each market are sorted into terciles based on market capitalization (large-, medium-, and small-cap stocks). The mean and standard deviation are reported in basis points. Panel C presents average values across three European markets: the UK, Germany, and France.

Panel A: US

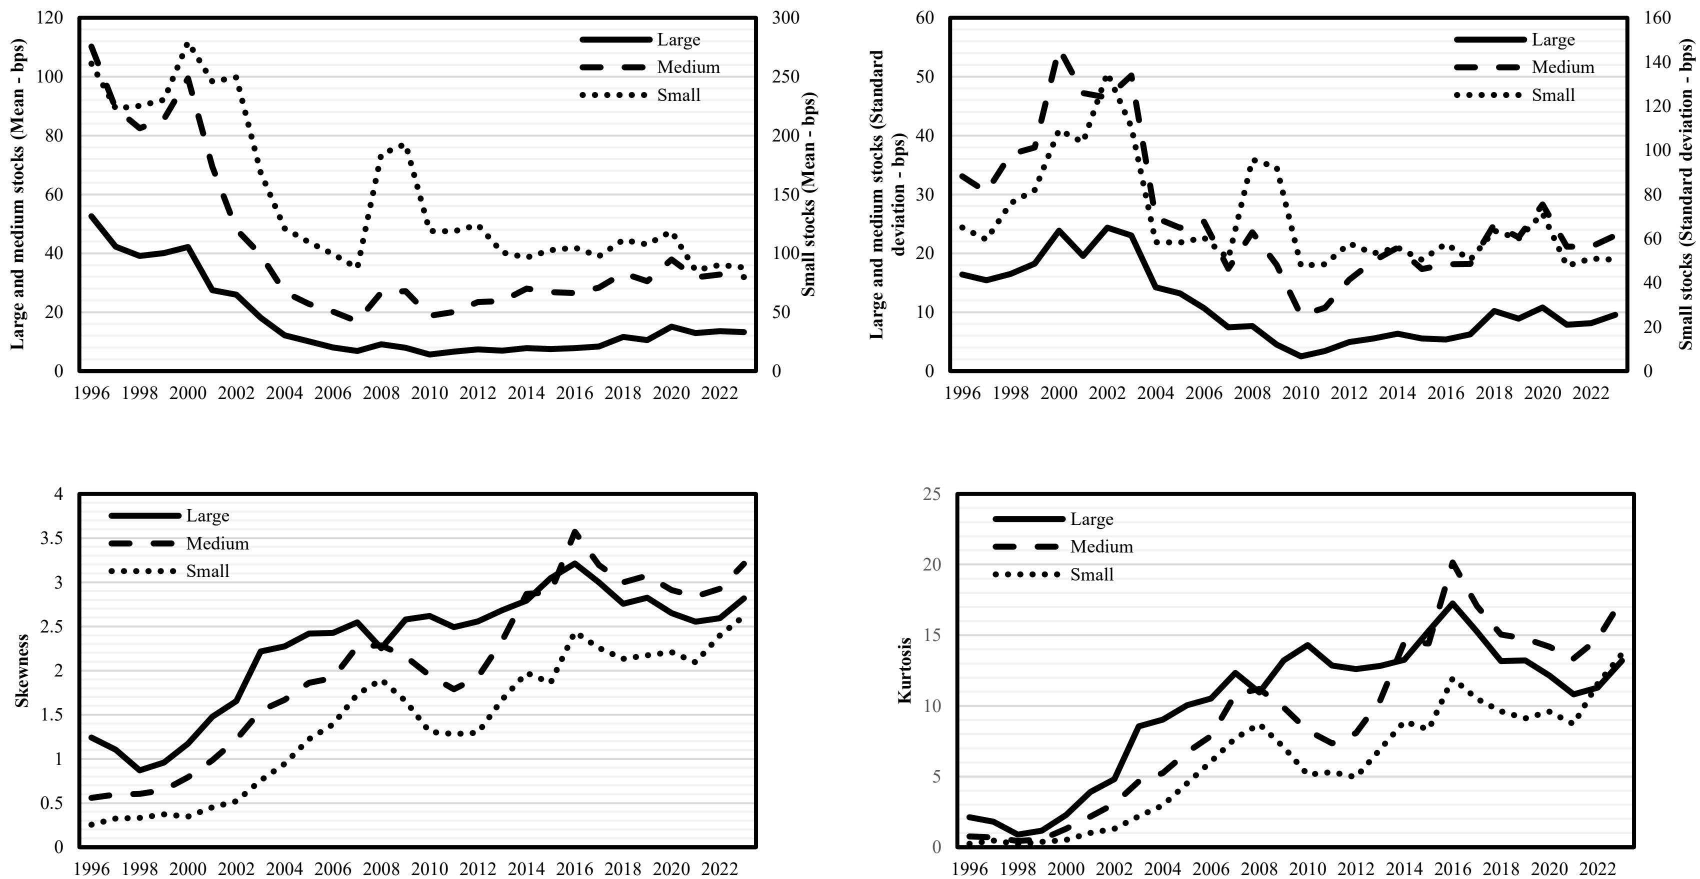

Panel B: Japan

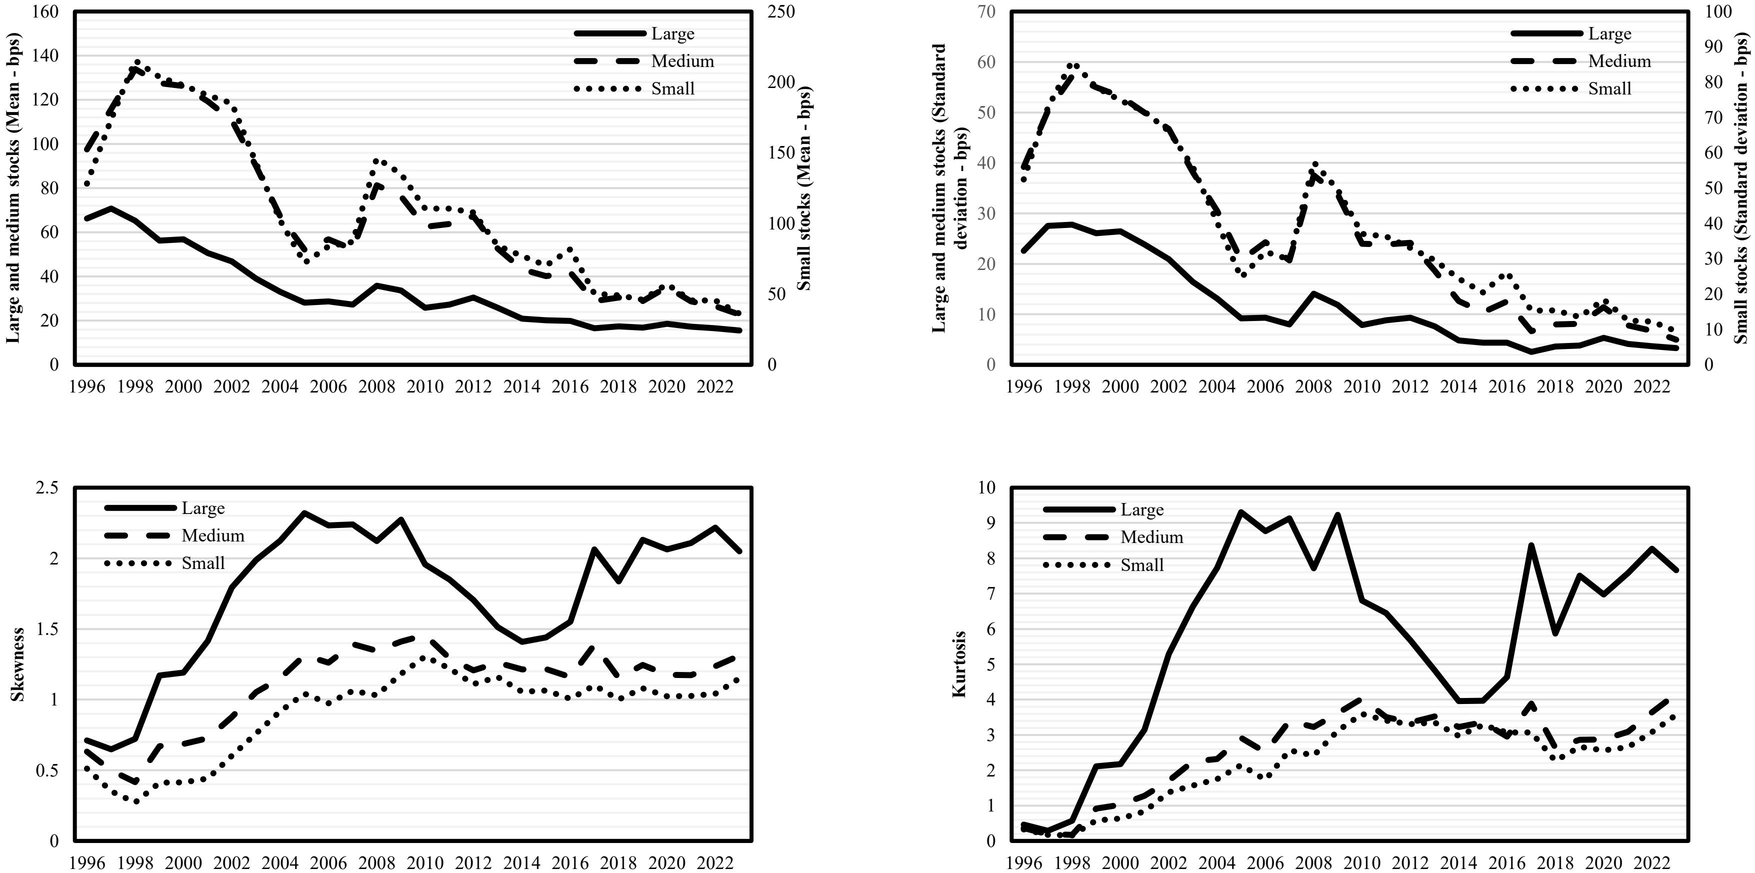

Panel C: Europe

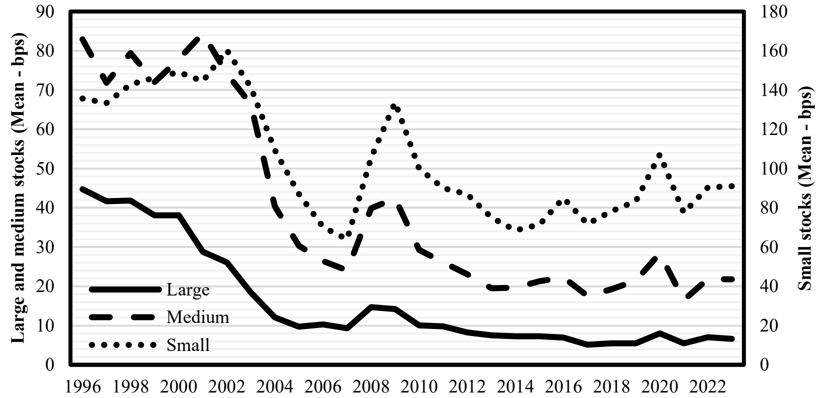

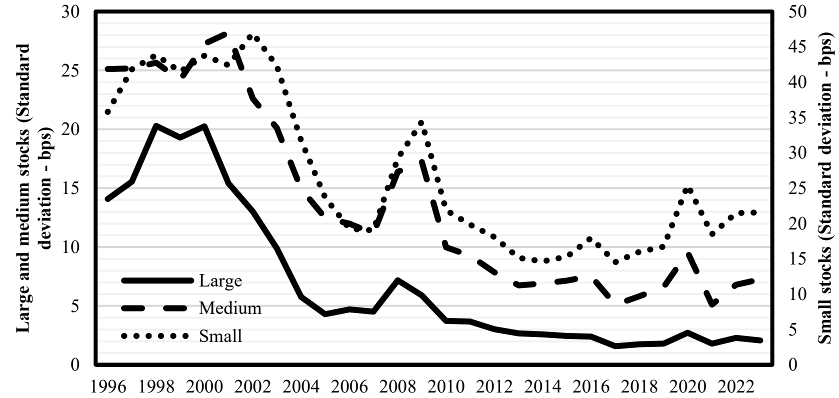

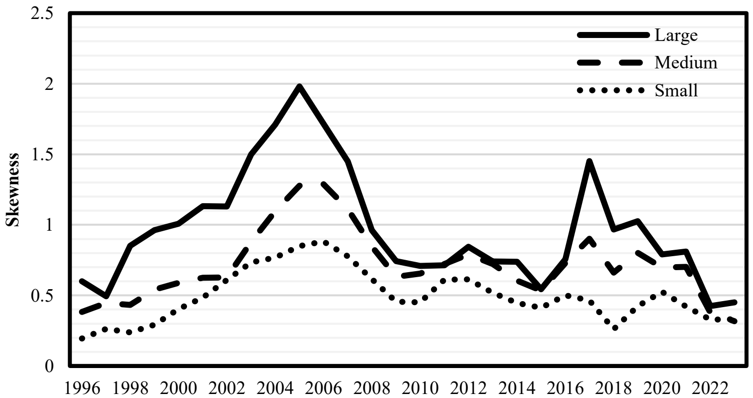

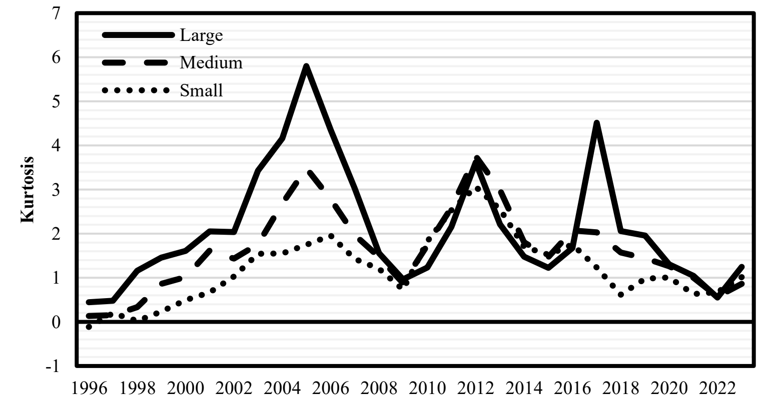

Figure 2. Time series of the moments of bid ask spread distribution in the foreign exchange spot market
This figure presents the mean, standard deviation, skewness, and kurtosis of 1-minute bid-ask spread distributions for trading Japanese Yen (JPY), Euro (EUR), and British Pound (GBP) against the US Dollar (USD). For each currency and day, we calculate the moments of relative spreads and average them annually. The mean and standard deviation are reported in basis points.

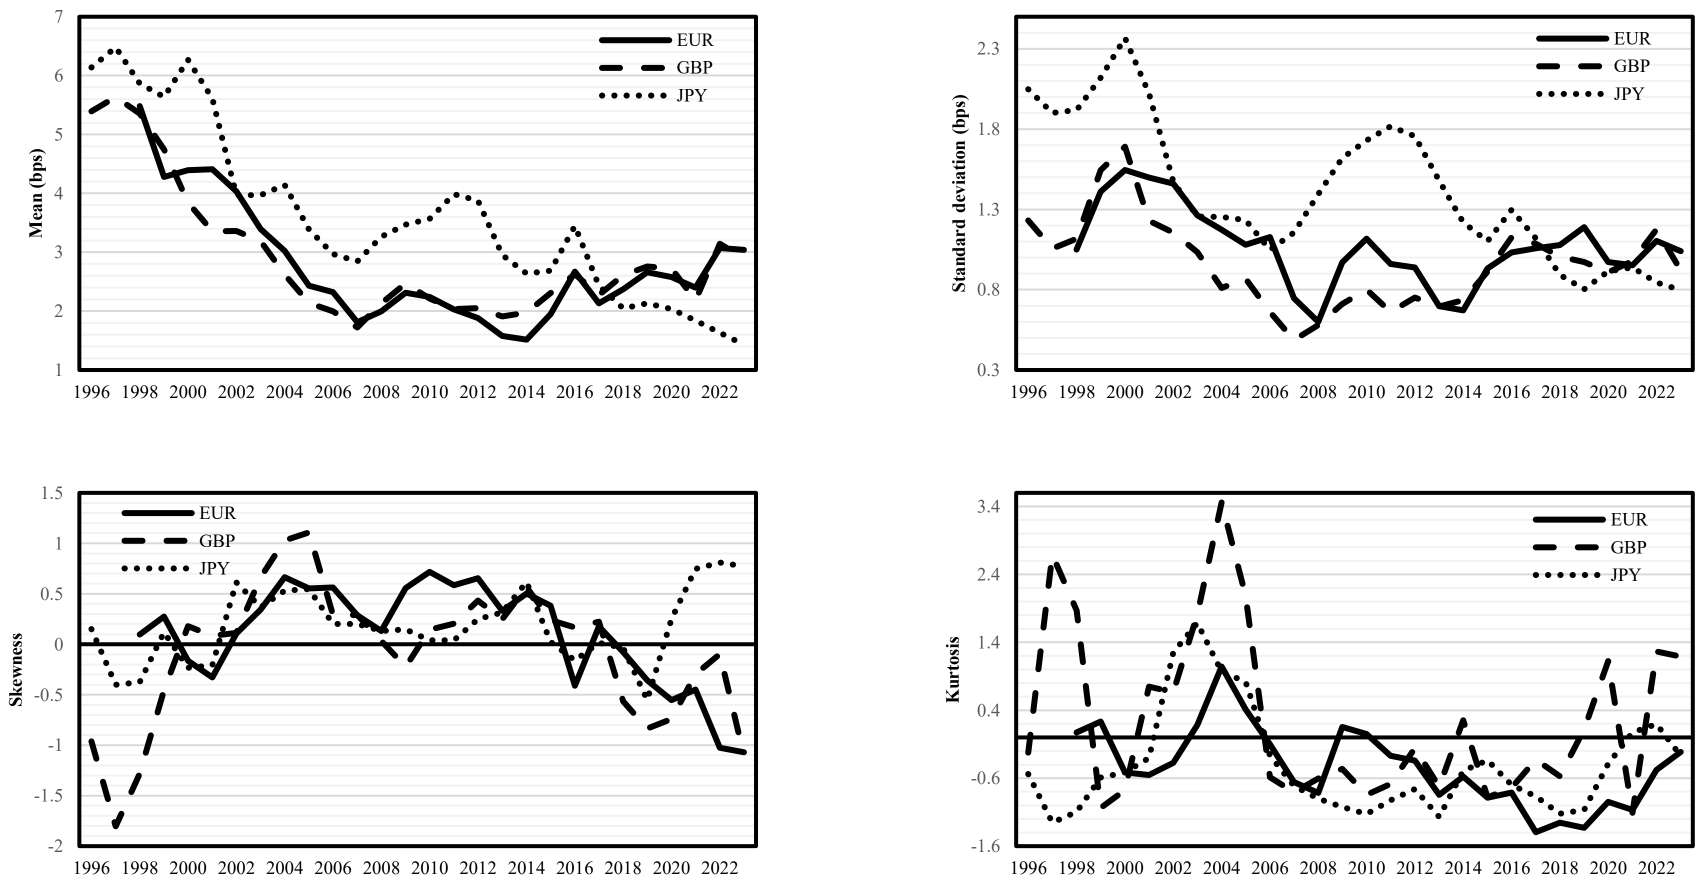

Figure 3. Time series of the moments of bid ask spread distribution in the sovereign bond market

This figure presents the mean, standard deviation, skewness, and kurtosis of 1-minute bid-ask spread distributions for government bonds in the US, Japan, and Europe, with maturities of 2, 5, and 10 years. For each bond and day, we calculate the moments of relative spreads and average them annually. The mean and standard deviation are reported in basis points. Panel C presents average values across three European government bonds: the UK, Germany, and Italy.

Panel A: US

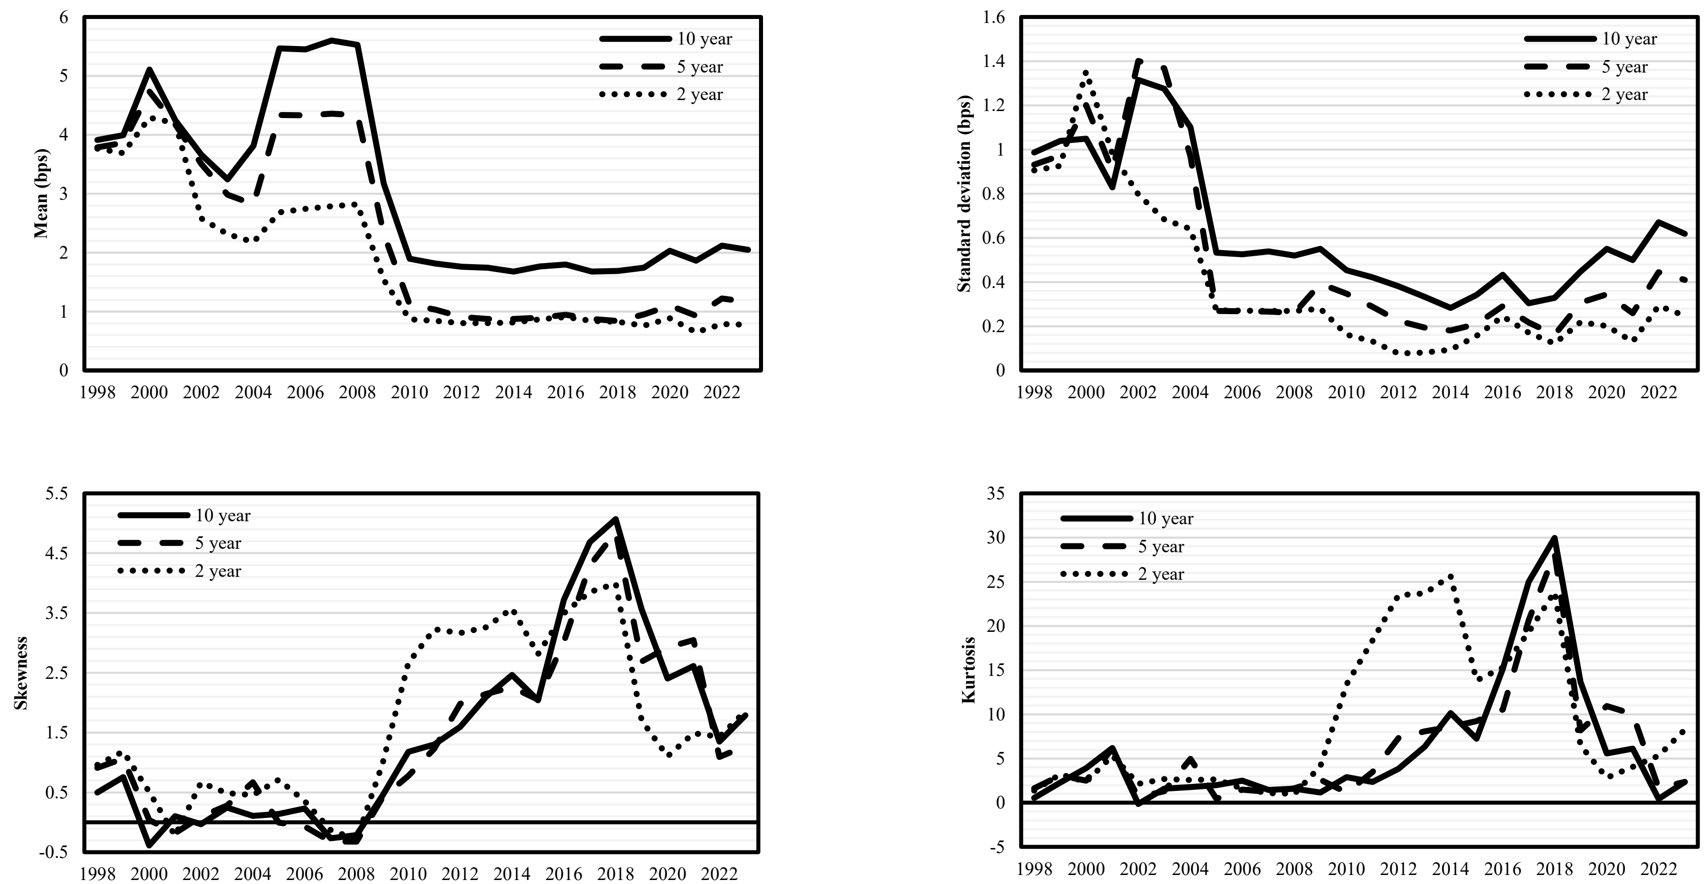

Panel B: Japan

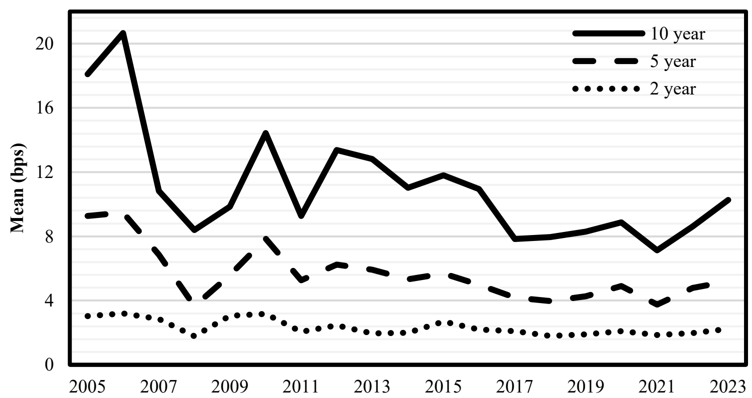

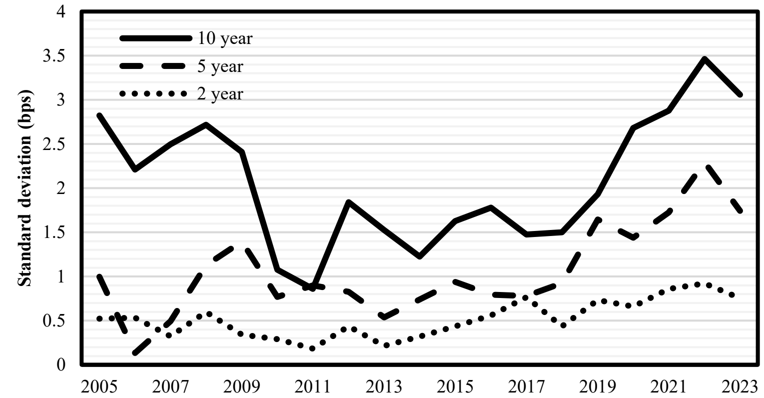

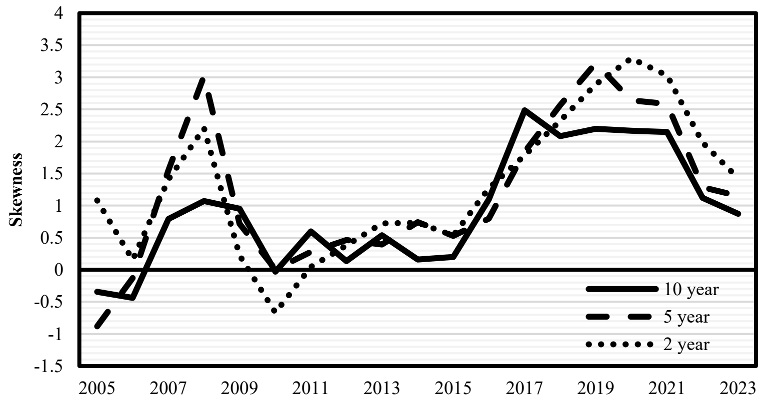

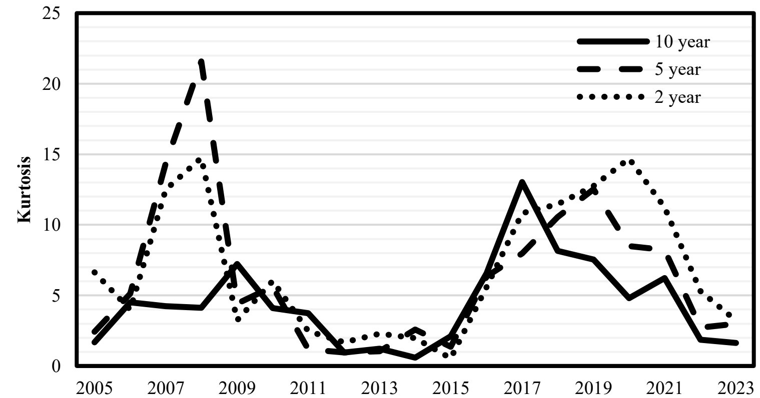

Panel C: Europe

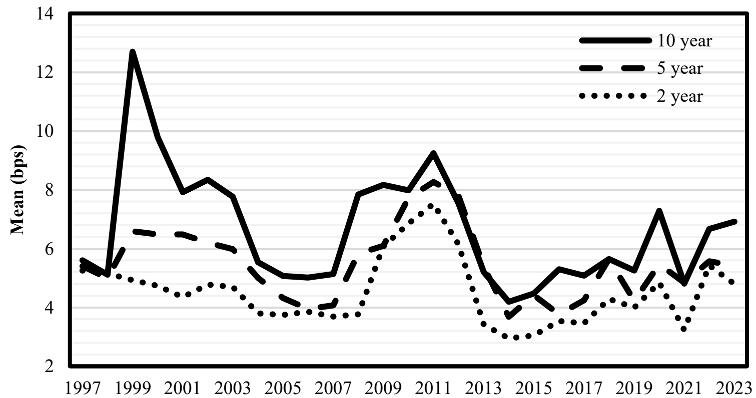

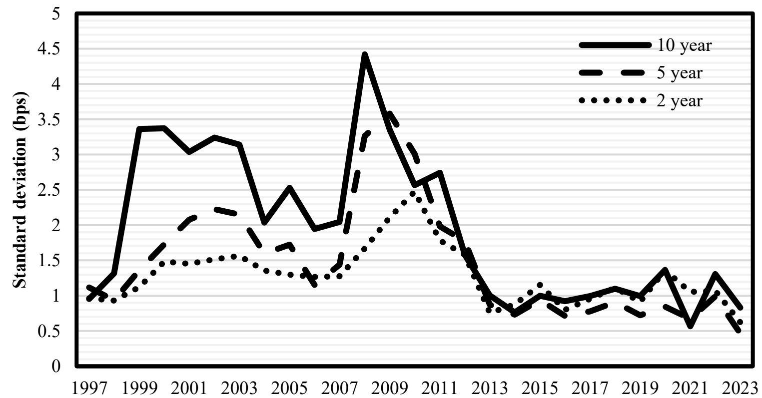

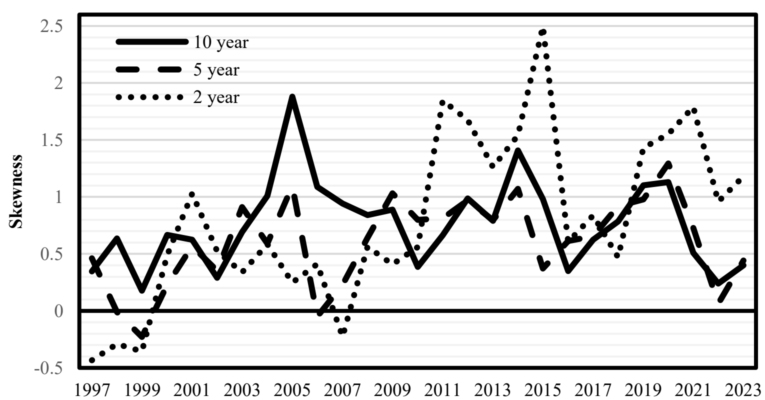

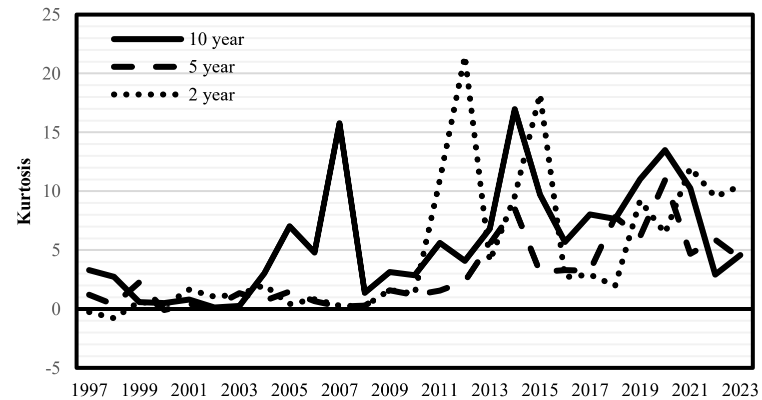

Figure 4. Cost differences between impatient and patient trades across skewness regimes
This figure presents the percentage difference in trading costs (i.e., prevailing bid-ask spread) between impatient and patient trades across three different bid-ask spread skewness regimes: low (0.82), medium (1.42), and high (2.7). The results are derived from the simulation described in Section 5. The vertical axis reports the difference between the average cost of impatient and patient trades as a percentage, while the horizontal axis represents the skewness regimes. Error bars indicate 95% confidence intervals.

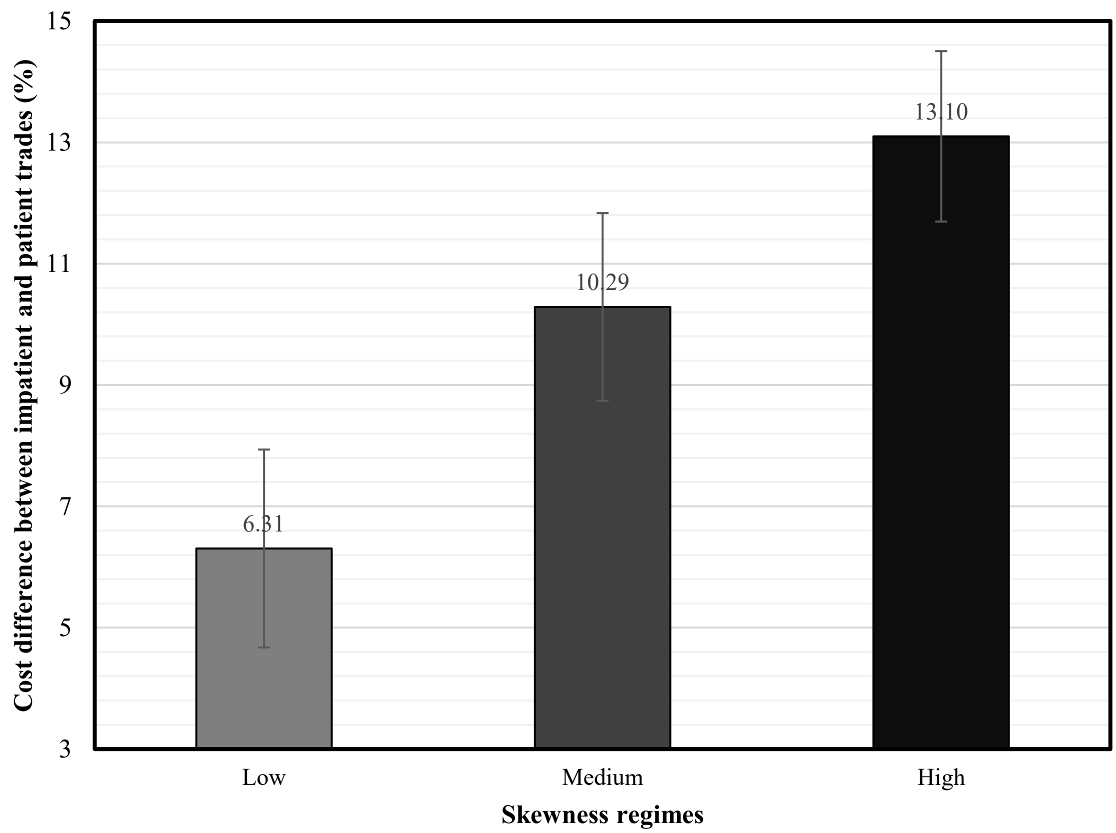

Figure 5. Trading cost against the probability of impatience for different skewness regimes
This figure presents the percentage change in trading costs relative to the benchmark probability of impatience $\mu = 0.1$ as a function of the probability of impatience (μ). The results are based on the simulation described in Section 5. The vertical axis shows the percentage change in trading costs (i.e., prevailing bid-ask spread) and the horizontal axis shows the probability of impatience ranging from 0.1 to 0.9. Three lines represent the low, medium, and high bid-ask spread skewness regimes. Error bands indicate 95% confidence intervals.

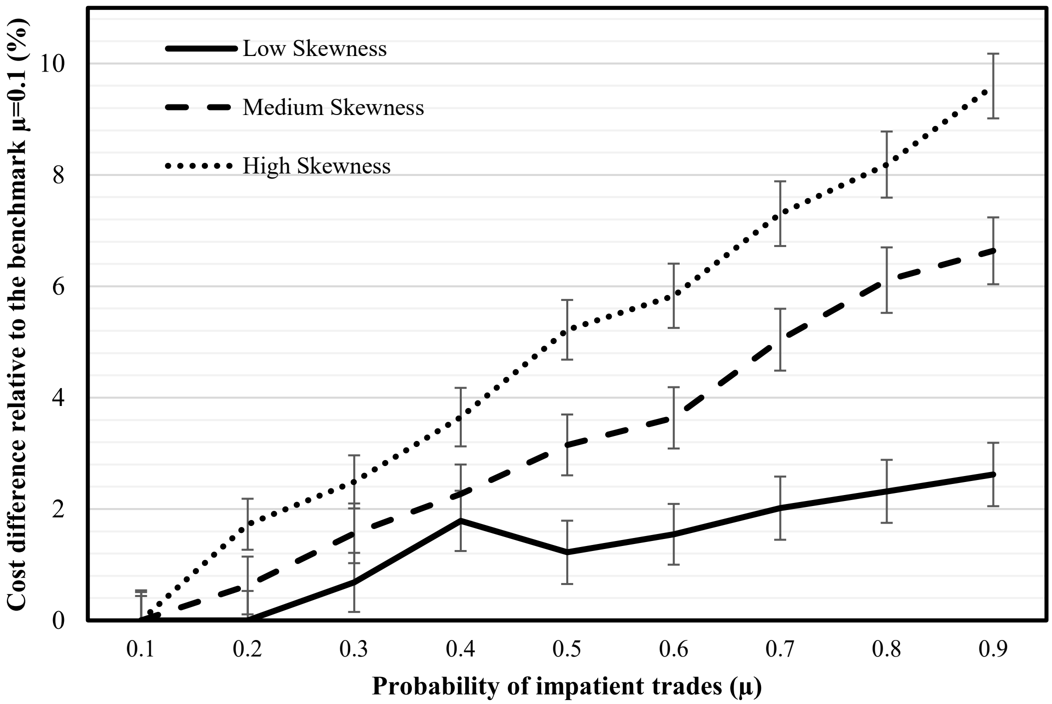

Figure 6. Alternative liquidity measures

This figure presents the mean, standard deviation, skewness, and kurtosis of alternative liquidity measures: effective spread, realized spread, and price impact in the US stock markets over time. For each stock-day, we calculate the moments of the liquidity measure and average them annually. The mean and standard deviation are reported in basis points.

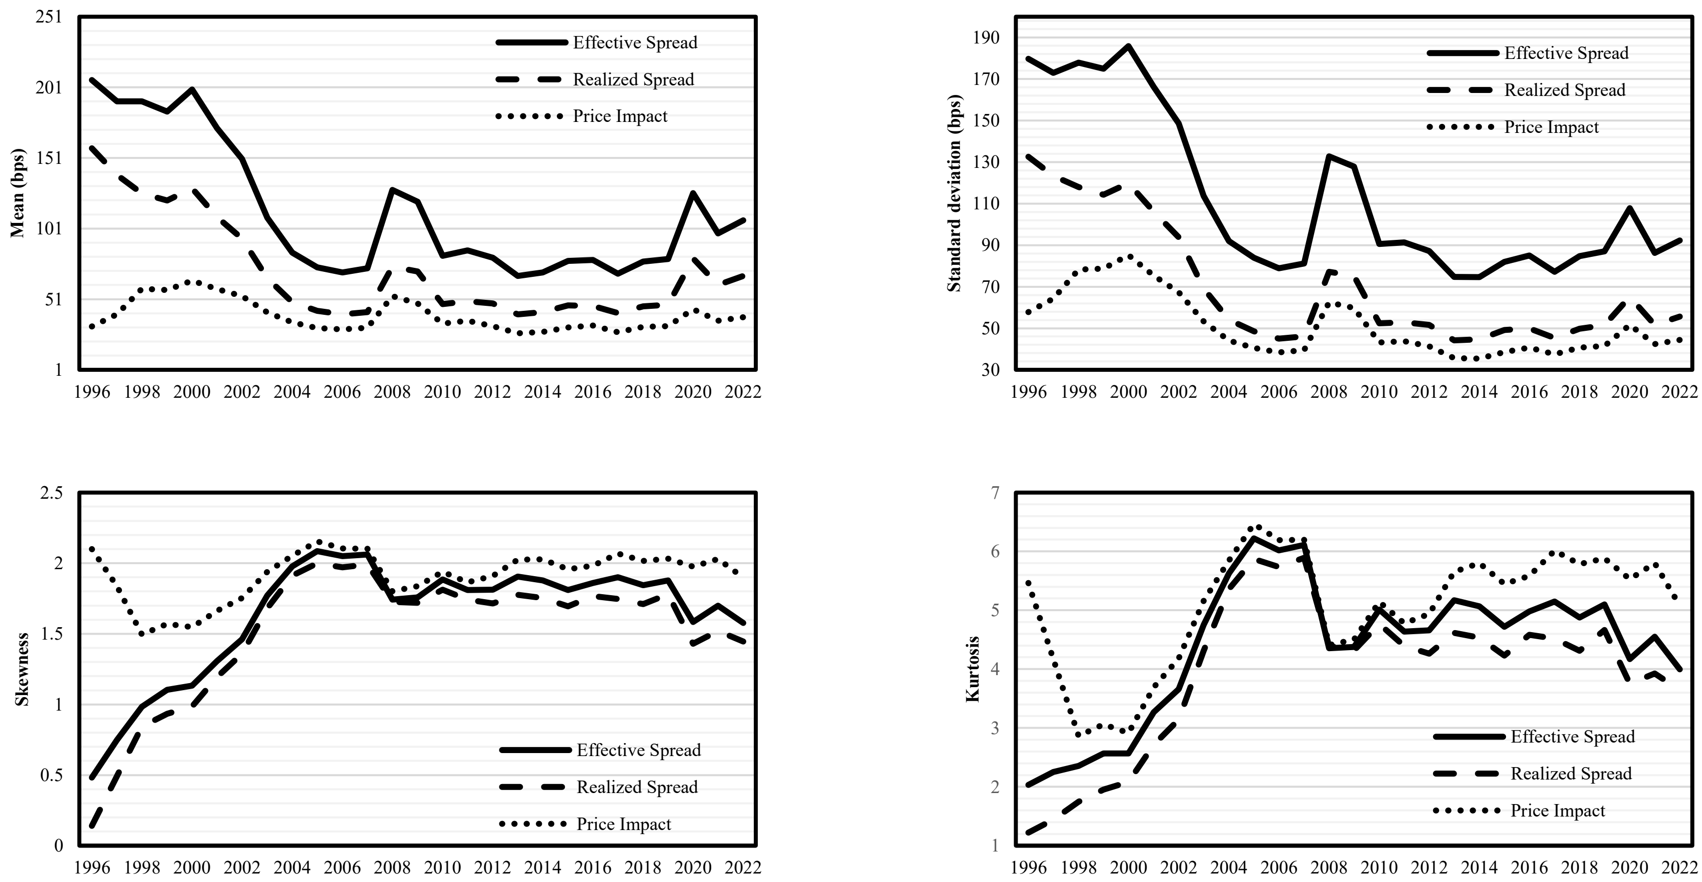

Table 1. Moments of the bid-ask spread distribution across asset classes and regions
This table reports the average values of the first four moments (mean, standard deviation, skewness, and kurtosis) of the bid-ask spread distribution across different asset classes (equities, bonds, and foreign exchange)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

This table presents the results of the trend regression analyses for the mean and skewness of the bid-ask spread. Specifically, we estimate the following model:

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

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

	$Mean_t$		$Skewness_t$	
	Coefficient (bps)	**t-stat**	**Coefficient (bps)**	**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
Medium				
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 in the mean of bid-ask spreads across different asset classes and regions

This table shows the breaks in the time series of the mean of bid ask spread identified using the Bai and Perron (1998) procedure discussed in Section 4.2. Upward (downward) breaks are represented with green (red) arrows. The test identifies the month of each break , but results are reported in six-month intervals for clarity.

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

Table 4. Breaks in the skewness of bid-ask spreads across different asset classes and regions

This table shows the breaks in the time series of the skewness of bid ask spread identified using the Bai and Perron (1998) procedure discussed in Section 4.2. Upward (downward) breaks are represented with green (red) arrows. The test identifies the month of each break , but results are reported in six-month intervals for clarity.

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

Table 5. Key events associated with breaks in the time series of mean and skewness of bid-ask spreads
This table describes key events that coincide with structural breaks in the time series of the mean or skewness of bid-ask spreads, as identified by the Bai-Perron method in Section 4.2. Events are grouped into three categories: macroeconomic events, market structure changes, and regulatory reforms. The first column lists the event name, the second provides its timing, and the third provides a brief description.

Event	Timing	Description
		Market structure changes
Decimalization in US equity markets	2000-2001	Decimalization refers to the change in US equity markets from quoting prices in fractions (e.g., 1/8, 1/16 of a dollar) to decimal units (e.g., \$0.01). The move, ultimately mandated by the US Securities and Exchange Commission (SEC), was intended to enhance price competition, reduce transaction costs, and improve transparency. The reduction in tick size to one cent narrowed bid-ask spreads across most equities and facilitated algorithmic trading strategies that benefit from finer price increments. See for example, https://www.sec.gov/files/decimalization-072012.pdf
Autoquote adoption	2003	Before Autoquote, NYSE specialists were required to manually disseminate the inside quote. The introduction of Autoquote in NYSE in 2003 replaced this process with a new automated quote. Hendershott, Jones and Menkveld (2011) document that the introduction of Autoquote on NYSE narrowed bid-ask spreads, and they attribute this to an increase in algorithmic trading.
API access in FX	2004 (banks); 2005 (non-banks)	FX venues like EBS enabled API connectivity, initially for banks, then for HFTs and principal trading firms (PTFs). API access allowed participants to send and receive quotes programmatically and with low latency. As a result, algorithmic trading activity increased in the market, particularly arbitrage strategies that exploited small price discrepancies across venues. Chaboud et al. (2014) document that following the API access, quote activity, trading volume, and order cancellation rates increased significantly. The introduction of APIs also changed the market power dynamics, with non-bank liquidity providers becoming increasingly dominant in interdealer trading.
Speed bumps in FX	2013-2014	In response to increasing concerns over latency arbitrage, several FX platforms introduced venue-level interventions, such as randomized order matching, throttled message processing, and trade batching around 2013 and 2014. These mechanisms, often referred to as "speed bumps", were designed to neutralize the speed advantage of HFTs by introducing small, controlled delays in order processing. For example, EBS introduced randomized pauses before matching orders, and other venues batched orders for simultaneous processing. These changes modestly increased average spreads but helped reduce predatory trading behavior. Budish et al. (2015) provide a theoretical rationale for batch auctions as a way to restore fairness in fragmented, speed-sensitive markets.

Euro introduction	1999 (electronic); 2002 (cash)	The euro was launched on 1 January 1999. For the first three years it was an 'invisible' currency, only used for accounting purposes and electronic payments. Coins and banknotes were launched on 1 January 2002, and in 12 EU countries the biggest cash changeover in history took place. See, for example, https://european-union.europa.eu/institutions-law-budget/euro/history-and-purpose_en.
	Regulatory reforms	
Reg NMS (US)	Approved 2005; Implemented 2007	Regulation National Market System (Reg NMS), implemented in 2007, aimed to improve execution quality and price competition in US equity markets. Its core rule was the Order Protection Rule, which required brokers to route orders to the venue with the best available price. While this improved average liquidity, it also triggered market fragmentation by incentivizing new venues and complex routing practices. Execution became harder to monitor, especially for investors without fast data access. See, https://www.sec.gov/rules-regulations/2005/06/regulation-nms.
MiFID I (EU)	2007	The Markets in Financial Instruments Directive (MiFID I), implemented in the European Union in November 2007, was the first major attempt to harmonize trading rules across EU member states. It introduced the concept of a "single European market for financial services" and allowed for the creation of multilateral trading facilities (MTFs). While MiFID I increased transparency through post-trade reporting obligations, it also triggered significant fragmentation of order flow across lit and dark venues. See, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32004L0039.
MiFID II (EU)	2018	The revised Markets in Financial Instruments Directive (MiFID II), implemented in January 2018, was designed to address many of the shortcomings of MiFID I, particularly around transparency, dark trading, and investor protection. MiFID II introduced volume caps on dark pool trading, mandated pre- and post-trade transparency for a wider range of instruments (including bonds and derivatives), and created new categories of venues such as Organized Trading Facilities (OTFs). These reforms aimed to bring more trading onto transparent venues and reduce information asymmetries. However, they also introduced operational complexity and compliance burdens. See, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32014L0065.
Financial Instruments and Exchange Act (Japan)	2006	Japan's Financial Instruments and Exchange Act (FIEA) was passed in 2006 (took effect in 2007) in response to growing market complexity, increased cross-border activity, and the need for stronger investor protection in Japan. It aimed to modernize Japan's capital markets by creating a unified, cross-sectional legal framework that covers a wide range of financial instruments and services. FIEA rests on four main pillars: (i) establishing a consistent legal structure for investor protection across financial products; (ii) enhancing corporate disclosure requirements; (iii) strengthening the governance and oversight of self-regulatory organizations like exchanges; and (iv) introducing

		stricter countermeasures against unfair trading practices. These changes increased confidence in Japan's financial markets. See, for example, https://www.fsa.go.jp/en/policy/fiel/index.html.
FX Global Code	2016-2017	The FX Global Code is a set of principles developed by central banks and market participants to promote integrity and transparency in the foreign exchange market. The Code was a response to a series of misconduct scandals, including benchmark rate manipulation and front-running, that had eroded trust in FX trading. It outlines best practices across different areas. The Code is voluntary but has been widely adopted. See, https://www.bis.org/about/factmktc/fx_global_code.htm.
	Macroeconomic events	
Global Financial Crisis (GFC)	2007-2009	The Global Financial Crisis (GFC), which began in mid-2007 and escalated with the collapse of Lehman Brothers in 2008, was one of the most severe global market dislocations since the Great Depression. Initially triggered by a breakdown in the US subprime mortgage market and the widespread mispricing of credit risk, the crisis rapidly spread across asset classes and geographies, leading to widespread liquidity freezes, counterparty concerns, and extreme volatility. Central banks responded with emergency programs, but market functioning remained impaired for an extended period. See, https://www.federalreserve.gov/newsevents/speech/bernanke20090113a.htm.
"Whatever it takes" speech (ECB)	2012	In July 2012, then European Central Bank (ECB) President Mario Draghi delivered his now-famous "whatever it takes" speech at a time when the euro area was facing severe sovereign debt tensions. Speaking at a conference in London, Draghi said: "within our mandate, the ECB is ready to do whatever it takes to preserve the euro. And believe me, it will be enough." This marked a pivotal moment in the European sovereign debt crisis and helped to restore investor confidence. While the speech did not announce an immediate policy action, it was widely seen as a signal of the ECB's commitment to backstop the euro. See verbatim remarks made by Mario Draghi https://www.ecb.europa.eu/press/key/date/2012/html/sp120726.en.html.
Abenomics	2012	Abenomics refers to the set of economic policies introduced by Japanese Prime Minister Shinzō Abe in late 2012, aimed at reviving Japan's economy from decades of deflation and stagnation. In early 2013, Japan implemented a monetary regime shift: the Bank of Japan (BoJ) adopted a 2% inflation target and launched Quantitative and Qualitative Easing (QQE), committing to achieve the inflation goal by 2015. The government supported this through expansionary fiscal measures and a roadmap for structural reforms. This "three-arrow" strategy (monetary easing, fiscal stimulus, and growth-enhancing reforms) became collectively known as Abenomics. See, for example, https://www.elibrary.imf.org/display/book/9781498324687/ch002.xml.

Table 6. 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 intraday 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 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 Market

	$\boldsymbol{Mean_{i,m+1}}$	$\boldsymbol{Skewness_{i,m+1}}$
$Algo_{i,m}$	-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)
Fixed Effect	Stock	Stock
N obs.	811,638	807,941
R^2	7.2%	15.2%

Panel B: Foreign exchange		
	$\boldsymbol{Mean_{i,m+1}}$	$\boldsymbol{Skewness_{i,m+1}}$
$Algo_{i,m}$	-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)
Fixed Effect	FX pair	FX pair
N obs.	378	375
R^2	17.4%	36.1%

Panel C: Government bond market		
	$\boldsymbol{Mean_{i,m+1}}$	$\boldsymbol{Skewness_{i,m+1}}$
$Algo_{i,m}$	-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)
Fixed Effect	Bond	Bond
N obs.	2,061	2,059
R^2	8.1%	6.9%

Previous volumes in this series

1228 November 2024	Through stormy seas: how fragile is liquidity across asset classes and time?	Nihad Aliyev, Matteo Aquilina, Khaladdin Rzayev and Sonya Zhu
1227 November 2024	Decentralised dealers? Examining liquidity provision in decentralised exchanges	Matteo Aquilina, Sean Foley, Leonardo Gambacorta and William Krekel
1226 November 2024	How Does Fiscal Policy affect the Transmission of Monetary Policy into Cross-border Bank Lending? Cross-country Evidence	Swapan-Kumar Pradhan, Előd Takáts and Judit Temesvary
1225 October 2024	Trade fragmentation, inflationary pressures and monetary policy	Ludovica Ambrosino, Jenny Chan and Silvana Tenreyro
1224 October 2024	A theory of economic coercion and fragmentation	Chris Clayton, Matteo Maggiori and Jesse Schreger
1223 October 2024	Carbon prices and reforestation in tropical forests	José A Scheinkman
1222 October 2024	Artificial intelligence and big holdings data: opportunities for central banks	Xavier Gabaix, Ralph S J Koijen, Robert Richmond and Motohiro Yogo
1221 October 2024	Bank geographic diversification and funding stability	Sebastian Doerr
1220 October 2024	The road to net zero: a fund flow investigation	Louisa Chen and Koji Takahashi
1219 October 2024	Stablecoins, money market funds and monetary policy	Iñaki Aldasoro, Giulio Cornelli, Massimo Ferrari Minesso, Leonardo Gambacorta, Maurizio Michael Habib
1218 October 2024	Bank specialisation and corporate innovation	Hans Degryse, Olivier De Jonghe, Leonardo Gambacorta and Cédric Huylebroek
1217 October 2024	Global inflation, inflation expectations and central banks in emerging markets	Ana Aguilar, Rafael Guerra and Berenice Martinez
1216 October 2024	Trade credit and exchange rate risk pass-through	Bryan Hardy, Felipe Saffie, Ina Simonovska

All volumes are available on our website www.bis.org.