DEPARTMENT OF ECONOMICS AND FINANCE
MASTER'S DEGREE IN ECONOMICS AND FINANCE
COURSE: EMPIRICAL FINANCE

Modeling Illiquidity Dynamics and Risk Premia

Prof. Paolo Santucci de Magistris
. . . .
SUPERVISOR

Prof. Emilio Barone
. . . .
CO-SUPERVISOR

Gian Marco Di Virgilio - 784651
. . . .
CANDIDATE

Academic Year 2024–2025

ABSTRACT

This thesis investigates whether market illiquidity shows a risk premium in U.S. equities. The analysis is grounded in the market microstructure view of illiquidity as *price impact*, i.e., the sensitivity of prices to order flow (Kyle, 1985), and builds on the practical low-frequency proxy proposed by Amihud (2002). The central empirical question is whether exposure to illiquidity-related factors is priced in the cross-section of stock returns, after controlling for standard risk factors.

A key contribution of the thesis is to connect Amihud-style illiquidity to a realized-variation framework and to develop scalable proxies suitable for large cross-sections. In particular, three families of measures are computed and compared: (i) the classic daily Amihud (Amihud, 2002); (ii) a realized Amihud based on intraday realized power variation when intraday data are available (Lacava et al., 2023; Barndorff-Nielsen and Shephard, 2002); and (iii) a range-based Amihud proxy where the volatility input is constructed from the daily high–low range using a Parkinson-type estimator (Parkinson, 1980). This design makes it possible to study the trade-off between informational content and data requirements, and to assess whether daily proxies can approximate realized illiquidity dynamics at scale.

To model the time-series behavior of illiquidity measures, which are non-negative and exhibit persistence and clustering, the thesis adopts the Multiplicative Error Model (MEM) framework (Engle, 2002; Engle and Gallo, 2006), following the cross-sectional implementation in Brownlees et al. (2011). The MEM decomposition delivers an economically interpretable separation between *expected illiquidity* (the predictable component) and *unexpected illiquidity* (the innovation), which is used to build illiquidity factors under tradable constructions.

Asset-pricing implications are examined through cross-sectional regressions in the spirit of Fama and MacBeth (1973). The thesis tests whether *expected* and *unexpected illiquidity* factors predict returns individually and jointly, and whether evidence of pricing remains after controlling for standard equity factors (e.g., FF3 (Fama and French, 1993)). The analysis is performed at the monthly frequency by aggregating daily illiquidity measures and stock returns, which makes it possible to evaluate how illiquidity pricing depends on the investment horizon.

CONTENTS

CHAPTER 1

INTRODUCTION

Liquidity is a central attribute of equity markets, yet it is also one of their most elusive. In everyday terms, a market is liquid when investors can trade quickly, in size, and without moving prices too much. In empirical work, however, this intuitive notion immediately fragments into several dimensions: trading costs, immediacy, depth, resiliency, and the sensitivity of prices to order flow. These dimensions matter for portfolio allocation, for risk management, and for the transmission of shocks, but they do not map one-to-one into a single observable statistic. A key difficulty is that liquidity is not a primitive object: it depends on the trading strategy, on the horizon, and on the state of the market. As a result, many empirical proxies capture mixtures of liquidity and other forces (news, volatility, and microstructure noise), and comparisons across assets become fragile when the measurement itself is unstable.

This thesis adopts a precise microstructure object: *price impact*. In the price-impact view, illiquidity reflects how strongly prices respond to trading pressure, i.e., the amount of price movement required to accommodate a given quantity of order flow (Kyle, 1985). Price impact is economically meaningful because it directly links liquidity to trading costs and to the ability of the market to absorb risk transfers. At the same time, price impact is not directly observable in standard daily datasets. For broad cross-sections and long samples, empirical work therefore relies on low-frequency proxies that summarize how much prices tend to move when trading volume is low. A classic and widely used proxy is the daily Amihud measure, which scales the absolute daily return by the dollar volume traded (Amihud, 2002). Interpreted through the price-impact lens, this ratio approximates the average price response per unit of traded value, making it a convenient bridge between market microstructure and asset-pricing applications.

Convenience, however, comes with a measurement problem. Daily illiquidity proxies are often noisy for reasons that are structural, not accidental. Close-to-close returns embed both intraday and overnight information, while the relationship between return magnitude and contemporaneous volume can be distorted by market-wide volatility, by discrete price formation, and by short-lived frictions such as bid–ask bounce. For illiquid stocks, irregular trading and occasional stale prices can further contaminate daily returns, producing mechanically large ratios that are not necessarily informative about persistent price impact. As a consequence, daily proxies may display strong persistence partly because of the latent liquidity state, but partly because of noise that is itself serially correlated through data artifacts and trading patterns. If the goal is to study liquidity *dynamics* and to test whether illiquidity is priced in the cross-section, these issues matter: a noisy signal can create spurious predictability, blur the distinction between persistent components and shocks, and ultimately weaken the interpretability of risk-premium tests.

A natural way to sharpen the measurement is to adopt the logic of realized variation. In volatility measurement, Barndorff-Nielsen and Shephard (2002) show that realized variance and related estimators aggregate high-frequency intraday returns to approximate an underlying integrated process more accurately than a single daily return. The same logic applies to illiquidity when illiquidity is framed as a volatility-per-volume object: instead of relying on one daily price change, one can aggregate intraday price movements over the trading day and then scale by traded volume. Building on this idea, Lacava et al. (2023) develop a realized Amihud measure in which the numerator is a realized power-variation component computed from intraday returns, providing a cleaner benchmark for integrated illiquidity. In principle, this approach addresses the core limitation of close-to-close measures: more intraday information reduces the role of idiosyncratic noise in the numerator and yields a signal that is closer to the latent price-impact component.

The obstacle is scalability. Intraday data are costly, heavy to process, and rarely available in a uniform way for thousands of stocks over long horizons. This thesis therefore takes a pragmatic route: it treats realized illiquidity as the *clean benchmark* when intraday data are available, and it develops a scalable daily alternative that retains as much of the realized-variation logic as possible. Specifically, I compute and compare three families of Amihud-style measures: (i) the classic daily Amihud proxy based on close-to-close returns (Amihud, 2002); (ii) a realized Amihud measure based on intraday realized power variation (Barndorff-Nielsen and Shephard, 2002; Lacava et al., 2023); and (iii) a range-based Amihud proxy in which volatility is proxied by the daily high–low range through a Parkinson-type estimator (Parkinson, 1980). The range-based design is motivated by a simple trade-off. Relative to realized measures, it is inevitably less informative because it compresses intraday prices into two points. Relative to close-to-close daily returns, however, the high–low range is designed to reflect the intraday path more directly, mitigating the reliance on a single closing price change. Operationally, this makes the range-based proxy attractive: it requires only daily high and low prices (available at scale), yet it is constructed to behave more like a realized object.

Measurement is only one side of the problem. Illiquidity is a non-negative variable that typically exhibits persistence and clustering, so its dynamics must be modeled in a way that respects these features. To this end, I adopt the Multiplicative Error Model (MEM) framework proposed by Engle (2002) and further developed in related work (Engle and Gallo, 2006). MEMs are designed for non-negative processes and provide a natural decomposition into a conditional mean component and a multiplicative innovation. Following the cross-sectional implementation proposed in Brownlees et al. (2011), I estimate MEM specifications stock by stock over a large panel. This yields an economically interpretable split between *expected illiquidity* (a slow-moving conditional mean that summarizes the predictable state) and *unexpected illiquidity* (the short-run disturbance around that state). In the context of illiquidity, this decomposition is especially useful: it separates persistent conditions that can be thought of as a liquidity environment from transient shocks that may reflect order-flow imbalances, temporary dislocations, or brief deteriorations in market depth.

The empirical objective of the thesis is to evaluate whether market illiquidity carries a risk premium in U.S. equities. The central question is not whether illiquid stocks earn higher average returns in a descriptive sense, but whether exposure to illiquidity-related *tradable* factors is priced in the cross-section after controlling for standard equity risk factors. To connect the illiquidity signals to asset-pricing tests, I map the estimated *expected* and *unexpected* components into monthly sorting variables and construct long–short factor-mimicking portfolios. Each month, stocks are sorted into equal-weighted deciles based on the relevant illiquidity signal, and a self-financing spread factor is formed as the high-minus-low return (decile 10 minus decile 1). This procedure yields four tradable illiquidity factors: *expected* and *unexpected illiquidity* under the daily proxy, and *expected* and *unexpected illiquidity* under the range-based proxy. The factors are evaluated using standard portfolio-sort diagnostics and time-series regressions against benchmark models, and they are then tested in cross-sectional pricing exercises using the two-step approach of Fama and MacBeth (1973). Benchmark controls include the market factor and common equity factors such as size, value, and momentum (Sharpe, 1964; Fama and French, 1993; Carhart, 1997).

The data construction is designed to combine broad coverage with transparent feasibility constraints. The equity universe is defined using the *CRSP U.S. Total Market Index* constituents as of September 30, 2025, which delivers a wide cross-section intended to represent the investable U.S. equity market[1]. Daily prices and volume are obtained from a publicly available online financial data provider through an automated download procedure, and the sample is filtered through data-quality checks to reduce the influence of missing prices, non-positive volume, and degenerate daily ranges. Daily illiquidity measures are then aggregated to the monthly frequency, and the portfolio evidence focuses on April 1998 to December 2025, using monthly excess returns. The emphasis on monthly testing reflects the goal of assessing pricing implications at a standard investment horizon while keeping the construction of daily signals and dynamics explicit.

In terms of contributions, the thesis first reinterprets Amihud-style illiquidity through a realized-variation perspective, treating intraday-based realized illiquidity as a benchmark and clarifying why close-to-close daily proxies can be intrinsically noisy (Amihud, 2002; Barndorff-Nielsen and Shephard, 2002; Lacava et al., 2023). Second, it develops and evaluates a scalable range-based Amihud proxy that relies on the daily high–low range, designed as a practical realized-type signal when intraday data are unavailable at scale (Parkinson, 1980; Lacava et al., 2023). Third, it models illiquidity dynamics in a large cross-section using MEMs, producing an expected/unexpected decomposition that is then used to construct tradable illiquidity factors and to assess their asset-pricing implications (Engle, 2002; Engle and Gallo, 2006; Brownlees et al., 2011; Fama and MacBeth, 1973).

The empirical results can be summarized without anticipating the full set of tables and diagnostics. On measurement, intraday-based realized illiquidity behaves as a superior benchmark, and the range-based proxy tends to track it more closely than the close-to-close daily proxy, consistent with the idea that the high–low range captures the intraday path more effectively. On dynamics, the MEM framework produces a coherent

[1] CRSP publishes official constituent lists at a quarterly frequency. Documentation on index levels and constituents is available at `https://www.crsp.org/indexes/levels-constituents/`.

separation between a slow-moving illiquidity state and short-run innovations; importantly, cleaner proxies reduce purely mechanical persistence and yield innovations that are more sharply defined. On pricing, portfolio sorts exhibit economically meaningful patterns, and long–short spreads can display positive abnormal performance relative to standard benchmarks. Yet, when the analysis shifts to cross-sectional pricing tests with standard factor controls, the evidence for a stable illiquidity price of risk is weak: any baseline support for *expected illiquidity* is at best marginal and tends to fade once common equity factors are included, while *unexpected illiquidity* does not deliver robust pricing under standard controls. Finally, the factor-exposure diagnostics suggest that expected-illiquidity spreads are mainly associated with size, value, and reversal-related behavior (often proxied by negative momentum), whereas unexpected-illiquidity spreads load more clearly on reversal, with limited market exposure.

The remainder of the thesis is organized as follows. Chapter 2 reviews the literature, moving from market microstructure foundations of price impact to low-frequency illiquidity proxies and to realized-variation approaches, and it situates the study within the broader evidence on illiquidity in asset pricing. Chapter 3 describes the data construction, the illiquidity measures, the MEM-based modeling strategy, and the portfolio and cross-sectional testing framework. Chapter 4 reports the empirical results, covering measurement comparisons, dynamic properties of illiquidity, factor construction, and pricing evidence. Chapter 5 concludes and discusses limitations, interpretation, and directions for further work.

CHAPTER 2

LITERATURE REVIEW

This chapter reviews the literature on market liquidity and illiquidity, with a particular focus on the notion of price impact, its empirical measurement, and its role in asset pricing. The discussion proceeds in three steps. First, I outline the market microstructure foundations of illiquidity, centering on Kyle's lambda (Kyle, 1985) as the canonical measure of price impact. Second, I review the evolution of empirical illiquidity proxies that are feasible in large samples, highlighting the role of Amihud-type measures and their recent extensions. Finally, I connect illiquidity dynamics to asset pricing, emphasizing the distinction between *expected* and *unexpected illiquidity* and motivating the empirical strategy adopted in this thesis.

2.1 From market microstructure to low-frequency illiquidity proxies

In market microstructure, liquidity is not a vague attribute of a security, but the outcome of a trading technology that maps order flow into transaction prices. The same asset can be liquid or illiquid depending on how quickly one can trade, how much one can trade without moving the price, and how fast temporary price pressure dissolves. The classical taxonomy emphasizes tightness (the cost of immediacy, often proxied by the bid–ask spread), depth (the quantity that can be absorbed at near-constant prices), and resiliency (the speed at which prices revert after liquidity shocks). Kyle (1985) makes this multidimensional view explicit within an equilibrium model, linking depth and resiliency to the strategic interaction between informed trading and noise trading. The central implication is that illiquidity can be understood as a price-impact mechanism: trades move prices because liquidity suppliers require compensation for bearing inventory and adverse-selection risk, and because order flow may reveal information about fundamentals.

This price-impact view is particularly sharp in Kyle (1985). In his framework, market makers set prices to satisfy a market-efficiency condition given aggregate order flow, while an informed trader optimally chooses trading intensity anticipating the induced price response. A single parameter summarizes the slope of the residual supply curve faced by the trader: the liquidity (price impact) parameter, whose reciprocal corresponds to market depth. In the linear equilibrium, depth is literally "the order flow necessary to induce prices to rise or fall" by a given amount, making price impact an operational object rather than a metaphor. Kyle (1985) also highlights that the same market can look liquid along one dimension (tightness) and illiquid along another (depth), which matters for empirical work because different proxies capture different margins of trading costs. The microstructure message is that illiquidity is inherently linked to the joint behavior of prices and trading activity, and that it should be allowed to vary over time with information arrivals and with the state of liquidity supply.

A complementary foundation comes from asymmetric-information models of dealer markets. In Glosten and Milgrom (1985), competitive market makers update beliefs about asset value from the direction of trades and set bid and ask quotes accordingly. The bid–ask spread emerges endogenously as compensation for trading against potentially informed counterparties, so that trading costs rise when adverse selection is more severe. This logic rationalizes the empirical observation that spreads widen when information asymmetry intensifies, and it also clarifies why transaction prices embed both compensation for liquidity provision and information revelation. In such settings, the act of trading generates a gap between transaction prices and fundamental value, and that deviation depends on order flow. This discrepancy is what makes the measurement of "true" liquidity challenging in practice: the object of interest is a conditional mapping from order flow to price changes, but researchers rarely observe the full order book, the identity of informed traders, or the dealer's inventory constraints.

The measurement problem becomes even more explicit in survey treatments. Madhavan (2000) stresses that liquidity is multi-faceted and that empirical measures reflect different microstructure channels, including order processing, inventory effects, and asymmetric information. He also notes a key asset-pricing link: if trading costs or price impact are state-dependent, then investors who anticipate future trading needs may require compensation for holding assets that become costly to trade precisely when liquidity is most valuable. In other words, illiquidity is not only a level effect (a contemporaneous cost) but can also be a risk attribute if it covaries with market-wide conditions. This perspective helps justify why a thesis that begins with microstructure naturally transitions toward asset pricing: once liquidity varies through time and co-moves across securities, it can enter expected returns either directly (as a carrying cost) or indirectly (as a systematic risk exposure).

The dependence of liquidity on information flow also motivates why volume and price variability often move together. Early empirical work such as Tauchen and Pitts (1983) formalizes the idea that both volatility and trading volume can be driven by a common latent arrival of information, implying that measures combining price variation and trading activity can be informative about trading frictions. While the volatility–volume link is not, by itself, a liquidity theory, it provides a practical cue for proxy design: quantities that scale price movements by the amount of trading may approximate the price response per unit of activity, which is precisely the primitive that price-impact models emphasize. This is also why a large family of liquidity measures in empirical finance uses both prices and volume, even when the underlying microstructure object (order flow at high frequency) cannot be observed directly for broad samples and long spans.

The practical constraint is that direct price-impact estimation is data intensive. Standard market-microstructure measures (effective spreads, realized spreads, and order-flow regressions in the spirit of Kyle's lambda) require intraday quotes and trades, and typically also careful handling of microstructure noise and trade classification. These requirements are feasible for a limited set of liquid stocks and relatively recent periods, but they are prohibitive when the goal is to study long histories or thousands of securities. For large cross-sections, researchers therefore rely on low-frequency proxies that preserve the intuition of price impact while using daily inputs (typically prices and volumes) that are widely available and comparable across long samples. The literature developed a broad menu of such proxies: turnover-based measures, zero-return or non-trading measures, spread approximations from daily highs and lows, and various price-movement-per-volume ratios[2].

Within this menu, the defining contribution of Amihud (2002) is to propose a scalable illiquidity proxy that can be computed from standard daily data while remaining close to the price-impact intuition: over a given period, the statistic summarizes how large absolute price changes tend to be relative to the amount of dollar trading activity, so that high values arise when sizable price moves occur on thin dollar volume, a pattern

[2]See Goyenko et al. (2009) for a systematic "horserace" of widely used low-frequency (daily-based) liquidity proxies against high-frequency benchmarks (effective/realized spreads and price-impact measures), and evidence on which proxies better track different liquidity dimensions.

naturally associated with low depth and high price impact in the Kyle sense; Amihud (2002) motivates the measure in these terms and links it to earlier empirical price-change-per-volume constructions in the microstructure tradition.

Two aspects of the Amihud construction explain its dominance in empirical work. First, it is nonparametric and uses only daily data, making it robust to model misspecification and applicable to very large panels. Second, it implicitly incorporates the microstructure idea that liquidity is about *quantity* and *price response* jointly: a day with a large absolute return is not necessarily illiquid if it coincides with extremely high dollar volume, whereas the same price move with thin volume is precisely the situation a price-impact perspective would label as illiquid. This "return scaled by activity" logic is also consistent with the view that order-flow information and trading frictions are intertwined, and it rationalizes why the measure is often used even outside equities, where granular order data are unavailable.

Beyond proposing a proxy, Amihud (2002) provides direct evidence that illiquidity is priced. His empirical strategy is not to introduce a fully articulated traded factor in the modern "factor-mimicking portfolio" sense, but to document an illiquidity premium through predictive and cross-sectional relationships between illiquidity and returns. In the cross-section, he constructs annual illiquidity measures and estimates regressions of stock returns on lagged illiquidity while controlling for standard characteristics and risk proxies. The key finding is that higher illiquidity predicts higher expected returns, consistent with investors requiring compensation for holding assets that are costly to trade. The analysis also incorporates market-wide illiquidity, connecting the microstructure notion of time-varying liquidity to an asset-pricing channel in which common liquidity conditions matter for expected returns.

A particularly relevant element for a thesis focused on decompositions into predictable and shock components is that Amihud (2002) also introduces a market illiquidity dynamics argument. He models the evolution of aggregate illiquidity using a simple forecasting equation in logs and defines the innovation term as an *unexpected illiquidity* component, denoted by an innovation ϵ_t in the market-illiquidity process. The economic interpretation matches the microstructure narrative: *expected illiquidity* reflects persistent market conditions that investors can anticipate, while *unexpected illiquidity* captures liquidity shocks that move transaction costs and price impact contemporaneously. In his time-series specifications, expected market illiquidity is associated with higher expected stock excess returns, whereas unexpected increases in illiquidity coincide with lower contemporaneous returns, consistent with a negative price reaction to adverse liquidity shocks. This distinction (*expected* versus *unexpected illiquidity*) does not rely on high-frequency microstructure estimation, yet it already embodies the state/shock logic that later motivates volatility-style modeling of positive and persistent liquidity measures.

Finally, the Amihud proxy naturally sits at the boundary between microstructure theory and large-sample empirical finance. It is "microstructure-consistent" in the sense that it tracks a price response per unit of trading activity, echoing the mapping that price-impact models formalize. At the same time, it is an intentionally coarse measurement choice: the numerator is based on close-to-close returns, compressing within-day dynamics into a single observation. This aggregation is precisely what makes the proxy scalable, but it also opens the door to measurement noise when intraday variation is substantial yet cancels out at the close, or when volume and volatility within the day carry information about trading frictions that the daily close does not reveal. These considerations motivate the subsequent move (central to the empirical design of this thesis) toward richer "within-day" information (range-based and realized constructions) while keeping the organizing principle anchored in price impact[3].

2.2 Realized and Range-based Illiquidity Measures

A central challenge in empirical liquidity analysis is the high measurement error plaguing traditional daily illiquidity proxies. The widely-used Amihud (2002) ratio of daily absolute return to dollar volume, for instance, relies on a single price change per day and aggregate volume, making it a noisy indicator of true price impact. Inspired by the volatility literature's advances in utilizing high-frequency data, recent research has sought to sharpen illiquidity measurement by incorporating intraday information. In the volatility context, Barndorff-Nielsen and Shephard (2002) showed that summing sufficiently frequent return observations yields a consistent and efficient estimator of the underlying integrated volatility. By analogy, liquidity researchers conjectured that ag-

[3]Daily close-to-close liquidity proxies can be noisy and only imperfectly track high-frequency price-impact benchmarks, which can weaken inference via attenuation from measurement error. See Goyenko et al. (2009) for a systematic comparison of low-frequency liquidity proxies to high-frequency benchmarks, and Lacava et al. (2023) for evidence that realized and high–low (range-based) Amihud-type measures substantially reduce noise relative to the classic daily proxy.

gregating intraday price moves and volumes could similarly improve the signal-to-noise ratio of illiquidity measures. This motivation culminated in the development of realized illiquidity metrics that build on high-frequency data to capture the cumulative price impact over a day more precisely than any single closing-price proxy.

The realized illiquidity framework extends the Amihud measure to intraday data, in effect blending the econometric theory of realized volatility with classic liquidity concepts. Lacava et al. (2023) formalize this approach by defining the realized Amihud as the ratio of daily realized return variation to total trading volume. In practice, the numerator can be computed as a non-parametric volatility estimator (such as realized variance or realized power variation à la Barndorff-Nielsen and Shephard 2003) using high-frequency returns, while the denominator is the day's volume. The resulting ratio estimates the integrated illiquidity of the day, meaning the aggregate price impact per unit volume over the trading session. Theoretical analysis confirms that realized Amihud converges to the true illiquidity in a continuous-time limit (with infinitely many price observations per day) and remains valid in the presence of stochastic volatility, jumps, and microstructure noise. Empirically, the benefits of using intraday data are striking: realized illiquidity series track the same general trends as the daily Amihud but with markedly lower high-frequency noise. In other words, much of the day-to-day noise in the classic illiquidity proxy are attributable to measurement error that is diversifiable when one aggregates many intraday price changes. The realized measure thus reveals cleaner dynamics, exhibiting stronger persistence and more pronounced patterns such as clustering of high-illiquidity periods and gradual mean reversion. It also uncovers intuitive relations like a leverage effect in liquidity: illiquidity tends to spike more following market drops than after rises, consistent with liquidity drying up during downturns. Notably, Lacava et al. (2023) find that realized illiquidity even has predictive power for short-term returns, with sudden liquidity shocks forecasting negatively the next-day performance of the asset. Such evidence underscores that using high-frequency information materially improves illiquidity measurement – an observation further confirmed by independent applications of realized Amihud in other markets.

While fully intraday-based measures are ideal, they require comprehensive high-frequency data that may not be available for all assets or sample periods. A practical compromise is to exploit daily range information – the high and low prices – to proxy intraday volatility. Range-based estimators, pioneered by Parkinson (1980), use the intuition that the distance between the day's extreme prices contains far more information about volatility than the mere closing price change. Indeed, Parkinson's classic result showed that the scaled high–low range is an unbiased volatility estimator with up to five times the efficiency of the close-to-close variance estimator. Following this insight, one can define a high–low Amihud illiquidity measure as the daily high–low price range divided by the daily volume. This range-based illiquidity metric leverages intraday price dispersion to approximate the day's aggregate price impact in the absence of full tick-by-tick data. Monte Carlo simulations in Lacava et al. (2023) show that using the daily high–low range instead of close-to-close returns delivers a much cleaner low-frequency proxy for the *integrated illiquidity*[4]. When intraday data are available, their realized construction remains the natural benchmark. When intraday data are not available, the range-based proxy is a strong second-best and a clear improvement over traditional daily measures.

The advent of realized and range-based illiquidity measures not only reduces measurement error but also yields time series with rich dynamics akin to those of volatility. Realized illiquidity typically exhibits strong autocorrelation and persistence, with bursts of illiquidity clustering in time and decaying only gradually. These properties point to underlying heterogeneous components driving liquidity over short and long horizons (in spirit of long-memory or HAR-type behavior; see Corsi 2009) and suggest that advanced time-series models are required to capture its evolution[5]. Another important characteristic is the non-negativity of illiquidity measures by construction – whether classic Amihud, realized, or high–low, they are strictly positive quantities. Consequently, modeling approaches must respect this positivity while accommodating the volatility-like features of persistence, clustering, and occasional spikes (jumps) in the series. A natural choice is the class of Multiplicative Error Models (MEM) introduced by Engle (2002) for exactly such non-negative processes. In an MEM specification, the variable of interest $\mathcal{A}_t$ (here a daily illiquidity measure) is decomposed as $\mathcal{A}_t = \mu_t \epsilon_t$, where $\mu_t = \mathbb{E}[\mathcal{A}_t \mid \mathcal{F}_{t-1}]$ is

[4]In Lacava et al. (2023), Monte Carlo evidence indicates that the high–low Amihud has a small negative bias (about 5%) and is about three times more efficient than the classic daily Amihud, making it a practical proxy for integrated illiquidity when intraday prices are unavailable.

[5]Empirically, illiquidity measures based on daily data tend to be highly persistent and clustered over time, and their economic impact is heterogeneous across firms. In particular, Amihud (2002) shows that *expected market illiquidity* varies through time and is associated with higher expected excess returns, with stronger effects for smaller stocks, consistent with the idea that liquidity contains both slow-moving components and short-run fluctuations.

a positive conditional expectation (or "scale factor") and ϵ_t is a mean-one innovation term with support on the positive real line. This multiplicative structure ensures x_t stays non-negative and lets one capture "conditional heteroskedasticity" in illiquidity without forcing a log-transformation of the data (Engle and Gallo, 2006). In effect, μ_t plays a role analogous to the conditional mean in an AR(1) or the conditional variance in a GARCH model, governing the persistent component of illiquidity, while ϵ_t captures random shocks. The MEM approach has proven effective in modeling financial time series like trade durations, volumes, and realized volatilities that share similar distributional features with illiquidity (Engle and Gallo, 2006; Brownlees and Gallo, 2010). By specifying dynamics for μ_t (e.g. an autoregressive structure that may include short- and long-term components), one can flexibly model the clustering of liquidity dry-ups and the slow reversion to "normal" liquidity levels, as well as potential asymmetries (for instance, allowing different responses to positive or negative return shocks). Extensions of the MEM have further enhanced its ability to mirror empirical patterns: Brownlees and Gallo (2010) incorporate smoothly time-varying long-run components to capture structural changes in volatility levels, and Caporin et al. (2017) introduce an MEM with jumps that can accommodate abrupt extreme moves in realized volatility (or illiquidity). These developments indicate that MEMs are well-suited to handle both the protracted persistence and the sudden spikes characteristic of realized illiquidity series.

Crucially, the MEM formulation yields a clear decomposition of illiquidity into expected and unexpected components at each point in time. The conditional mean μ_t can be interpreted as the *expected illiquidity* given all information up to $t-1$, while the innovation ϵ_t (with mean 1) represents the illiquidity shock on day t. This separation is especially useful when linking illiquidity to asset prices. For example, one can study whether it is the predictable variation in liquidity or the unpredictable shocks that matter more for returns. In their analysis of realized illiquidity, Lacava et al. (2023) find that after controlling for *expected illiquidity*, it is the *unexpected liquidity* shocks that bear a significant relation to subsequent returns. From an asset pricing perspective, this suggests that investors demand compensation primarily for innovations in illiquidity rather than for high illiquidity per se, consistent with the idea that unanticipated liquidity dry-ups are especially painful[6]. Motivated by this insight, the remainder of this thesis builds on

[6]A liquidity premium is typically interpreted as compensation for *liquidity risk*, i.e., for assets whose

these improved illiquidity measures and on the MEM-based decomposition to construct robust expected and unexpected components. The realized and range-based proxies developed here, together with MEM-implied conditional means and innovations, provide a disciplined toolkit to separate persistent liquidity conditions from short-run shocks. This framework sets the stage for the empirical analysis that follows, where the signals are aggregated into tradable long–short constructions and evaluated with standard econometric diagnostics and robustness checks.

2.3 Realized-variation illiquidity and scalable daily alternatives

In the realized volatility literature, intraday returns provide a direct route to measuring ex-post variability through realized variation and its generalizations. Under standard conditions, sums of powers of intraday returns converge to integrated variation objects, and realized measures can therefore be used to build informative volatility statistics without specifying a parametric diffusion model. This logic extends naturally to illiquidity measures built as ratios in which the numerator captures price variability and the denominator captures trading activity. Building on this insight, Lacava et al. (2023) propose a realized version of Amihud illiquidity by replacing the absolute daily return with realized absolute variation, thereby obtaining an intraday-based illiquidity statistic that retains the same basic structure of the classic proxy while exploiting richer within-day information. The construction connects illiquidity measurement to the realized power variation framework (Barndorff-Nielsen and Shephard, 2002) and helps clarify the role of sampling frequency as a practical design choice.

A recurring consideration in realized-measure construction is the trade-off between information efficiency and market microstructure noise. Sampling too finely can amplify noise and bid-ask bounce, while sampling with overly long intervals discards relevant information; moderate frequencies such as 5-minute returns are commonly used as a pragmatic compromise in empirical work (Hansen and Lunde, 2006). These considerations are directly relevant for realized-illiquidity numerators, since the numerator inherits the

returns covary with adverse *innovations* in aggregate liquidity. In this spirit, Pástor and Stambaugh (2003) show that a market-wide liquidity factor is priced in the cross-section, while Acharya and Pedersen (2005) formalize a liquidity-adjusted CAPM where expected returns depend on both liquidity levels and exposure to liquidity *shocks*. Consistently, Amihud (2002) documents that *expected market illiquidity* raises required returns, whereas *unexpected illiquidity* is associated with contemporaneous return drops, highlighting the economic relevance of liquidity innovations.

same measurement issues as realized volatility statistics computed from intraday prices.

Despite their appeal, intraday-based measures face a fundamental scalability constraint in broad cross-sections: intraday coverage is often uneven across stocks and time, historical depth is limited for many assets, and data costs and processing burdens can be substantial. This constraint motivates the use of daily alternatives that preserve the spirit of realized variability while relying only on daily inputs. A prominent route is to use daily range-based estimators. Parkinson-type range estimators (Parkinson, 1980) exploit the daily high and low to produce a volatility-type statistic that is typically more informative than the close-to-close absolute return, precisely because it embeds within-day extrema. When combined with dollar volume in a ratio form, a range-based Amihud proxy provides a daily, scalable approximation to realized-variation illiquidity, enabling large-panel analyses without requiring intraday data for every asset and every day.

These measurement considerations also have implications for modeling. Illiquidity series constructed from realized-variation principles, or from range-based proxies that mimic them, tend to display persistent dynamics and clustering patterns that are reminiscent of volatility clustering. Lacava et al. (2023) document that realized Amihud measures are markedly more persistent than the classic daily proxy, suggesting that improved measurement reveals a smoother latent illiquidity state. This motivates dynamic specifications that can capture persistence while respecting non-negativity.

2.4 Illiquidity in Asset Pricing

The microstructure foundations of illiquidity as trading cost and price impact naturally lead to an asset-pricing question: if trading frictions are costly and state-dependent, should investors require compensation for holding illiquid assets or assets exposed to systematic liquidity risk? The empirical asset-pricing literature has addressed this question along two complementary dimensions. The first concerns the *level effect* of illiquidity, that is, whether securities with higher average trading costs earn higher expected returns. The second concerns *liquidity risk*, namely whether exposure to fluctuations in market-wide liquidity commands a risk premium in equilibrium.

A direct test of the level effect is provided by Amihud (2002), who shows that stocks with higher illiquidity earn higher average returns, controlling for size and other standard characteristics. The economic intuition is straightforward: investors anticipate incurring higher trading costs over the life of an investment and therefore require a higher expected return as compensation. Importantly, Amihud (2002) also emphasizes that *expected market illiquidity* varies over time and affects aggregate required returns, suggesting that liquidity operates both at the cross-sectional and at the intertemporal level. Subsequent evidence confirms that illiquidity is priced not only in U.S. equities but also across international markets and asset classes[7].

While the level effect captures compensation for holding costly-to-trade assets, a deeper asset-pricing implication arises once liquidity is recognized as a state variable. In a conditional setting, what matters is not only how illiquid an asset is on average, but how its liquidity co-moves with aggregate market liquidity and returns. Pástor and Stambaugh (2003) provide a seminal contribution in this direction by constructing a traded liquidity factor based on the sensitivity of returns to order-flow–induced reversals. They show that stocks with higher loadings on aggregate liquidity shocks earn higher expected returns, even after controlling for the Fama–French factors. This finding shifts the focus from static trading costs to systematic liquidity risk.

The theoretical framework of Acharya and Pedersen (2005) formalizes this intuition within a liquidity-adjusted CAPM. In their model, expected returns depend on four distinct covariances: (i) the covariance between asset returns and market returns; (ii) the covariance between asset illiquidity and market illiquidity; (iii) the covariance between asset returns and market illiquidity; and (iv) the covariance between asset illiquidity and market returns. This structure implies that liquidity risk is multidimensional, and that both the level of illiquidity and its systematic interactions with aggregate states are priced in equilibrium. Empirically, they document evidence consistent with these predictions, reinforcing the interpretation of liquidity as a risk factor rather than a mere transaction cost.

These contributions place illiquidity within the broader factor-pricing tradition inaugurated by Fama and MacBeth (1973) and Fama and French (1993). In the Fama–MacBeth two-step framework, risk premia are identified by estimating asset-specific factor loadings in time-series regressions and then testing whether those loadings explain the cross-section of average returns. Liquidity-based asset-pricing tests typically follow this ap-

[7]Subsequent evidence confirms that liquidity risk is priced also outside the U.S., with significant premia documented across developed equity markets (e.g., Liang and Wei, 2012).

proach: first, construct either a liquidity factor or asset-level liquidity betas; second, estimate whether the cross-sectional price of liquidity risk is statistically different from zero. The Fama–French three-factor model (Fama and French, 1993) and its Carhart extension (Carhart, 1997) provide the standard benchmark against which liquidity premia are evaluated. A key empirical issue is whether liquidity-related factors retain explanatory power once market, size, value, and momentum exposures are controlled for.

An important distinction in this literature concerns whether liquidity is treated as a *characteristic* or as a *risk factor*. Characteristic-sorted portfolio evidence, such as that in Amihud (2002), documents return spreads between high- and low-illiquidity stocks. However, such spreads may reflect correlated exposures to other firm attributes, particularly size and value. Risk-based tests ask whether a stock's exposure to liquidity innovations commands a premium, rather than whether the level of illiquidity does. This distinction becomes especially relevant when illiquidity is decomposed into predictable and unexpected components. Persistent illiquidity may overlap with slow-moving firm characteristics, while *unexpected liquidity* shocks are more naturally interpreted as systematic risk realizations.

Empirical findings across studies are not fully uniform. While Pástor and Stambaugh (2003) and Acharya and Pedersen (2005) document significant liquidity risk premia, other contributions show that the magnitude and robustness of liquidity pricing are sensitive to measurement choices, sample construction, and benchmark specifications. For instance, alternative liquidity proxies can yield materially different cross-sectional results, and the inclusion of momentum or profitability factors may attenuate estimated liquidity premia. More recent surveys and replication studies highlight that liquidity pricing evidence can weaken once methodological refinements are introduced. These tensions suggest that measurement precision and factor construction play a central role in identifying liquidity premia.

Within this context, the construction of *tradable* liquidity factors becomes crucial. A factor-mimicking portfolio, formed by going long in high-illiquidity assets and short in low-illiquidity assets, provides a self-financing representation of liquidity exposure analogous to the SMB or HML factors (Fama and French, 1993). Such tradable spreads allow researchers to test whether liquidity earns abnormal returns relative to standard benchmarks and whether individual stocks with higher betas to these spreads earn higher expected returns in Fama–MacBeth regressions. The economic interpretation of a positive price of liquidity risk is that assets more sensitive to adverse liquidity states (for example, those that become illiquid precisely when market liquidity deteriorates) must offer compensation in equilibrium.

The recent development of realized and model-based illiquidity measures introduces an additional layer to this debate. If traditional daily proxies suffer from measurement error, estimated liquidity betas may be attenuated, biasing cross-sectional tests toward insignificance. More precise measures of price impact, combined with dynamic models that separate *expected* from *unexpected illiquidity*, may therefore sharpen identification of liquidity risk. In particular, innovations extracted from a Multiplicative Error Model provide a direct empirical counterpart to liquidity shocks in theoretical models. Whether these shocks command a positive price of risk is ultimately an empirical question, to be addressed through portfolio sorts and Fama–MacBeth cross-sectional regressions using tradable illiquidity factors.

The asset-pricing literature thus provides both a conceptual framework and a set of empirical tools for evaluating illiquidity premia. It establishes that liquidity can matter through levels, through systematic covariances, and through state-dependent interactions with market returns. At the same time, it leaves open the quantitative magnitude and robustness of such premia, particularly when measurement is refined and standard factor controls are imposed. The empirical analysis that follows builds directly on this tradition by constructing tradable factors based on *expected* and *unexpected illiquidity* and testing their pricing implications within the Fama–MacBeth framework.

CHAPTER 3

DATA AND METHODOLOGY

This chapter describes the data sources, sample construction, and empirical design adopted in the thesis. I first outline how the U.S. equity universe is assembled and how daily price and volume series are collected and cleaned. I then define the illiquidity proxies used in the analysis and motivate the range-based construction as a volatility-type daily signal. Next, I introduce the Multiplicative Error Model framework employed to decompose illiquidity into a predictable component (*expected illiquidity*) and an innovation (*unexpected illiquidity*). Finally, I describe how these MEM-implied signals are aggregated to the monthly horizon, mapped into tradable decile-spread factors, and brought to standard asset-pricing tests.

3.1 Data collection

The starting point of the empirical analysis is the construction of a broad and representative universe of U.S. equities. To this end, I rely on the *CRSP U.S. Total Market Index*, which is designed to represent approximately 100% of the U.S. investable equity market. The index coverage makes it particularly suitable for a comprehensive analysis of market-wide liquidity dynamics, avoiding selection biases associated with narrow benchmark indices. The list of eligible securities is obtained from the CRSP *Quarterly Constituents List*[8]. Specifically, I use the snapshot of constituents as of September 30, 2025, which defines the cross-sectional universe underlying the empirical analysis. The list was processed to obtain a structured and consistent set of ticker symbols and firm identifiers, with a preliminary cleaning step addressing the presence of multiple share classes for the same firm (e.g., Class A and Class B shares). In such cases, a single representative security was retained. The selection criterion was based on data availability: among multiple classes, I retained the class with the longest historical price series available in Yahoo Finance[9]. This choice ensures internal consistency of ticker identifiers and facilitates the

[8] CRSP publishes official constituent lists at a quarterly frequency. Documentation on index levels and constituents is available at `https://www.crsp.org/indexes/levels-constituents/`.

[9] Yahoo Finance (`https://www.finance.yahoo.com`) is a publicly accessible data provider that aggregates historical price, volume, and corporate information for a wide range of financial instruments across global markets.

subsequent data collection process, which relies on the `yfinance` Python library. After harmonizing ticker symbols and removing duplicate share classes, the initial universe consisted of 3,487 individual stocks. While this broad cross-section is well suited to capturing aggregate liquidity dynamics, not all securities provide data of sufficient quality and length for reliable econometric modeling. The initial universe was therefore further refined to enhance the robustness of the empirical analysis and to ensure that the dynamic specifications could be estimated on sufficiently informative time series. To this end, two additional filters were applied.

First, I excluded very small firms by imposing a minimum market capitalization threshold of 100 million U.S. dollars, measured at the time of data download. This restriction is not driven by universe size considerations per se, but by econometric concerns: extremely small and illiquid stocks are more prone to data errors, extreme price movements, and irregular trading patterns, which may distort both illiquidity measures and model estimation.

Second, I imposed a restriction on the minimum length of the available price history by requiring that each stock had begun trading no later than January 1, 2021. Given the end date of the dataset (January 21, 2026), this ensures a minimum time span of approximately five years of daily observations for each stock[10].

Importantly, no restrictions were imposed on stocks entering the market after December 31, 1997, provided that they satisfied the minimum data availability condition described above. After applying these filters, the sample was reduced to 2,565 stocks.

Having defined the equity universe and the associated selection criteria, the next step consists in collecting consistent daily price and volume data for each selected stock. For each stock in the filtered universe, I collected daily price and volume data from Yahoo Finance using the `yfinance` Python library. The dataset includes adjusted daily Open, High, Low, Close prices, as well as trading volume expressed in number of shares. All price series are fully adjusted for corporate actions, including stock splits and dividend payments. Importantly, adjustments were applied consistently not only to closing prices,

[10]This restriction ensures at least five years of daily observations for each stock, which helps mitigate small-sample distortions in the estimation of dynamic econometric specifications and improves the stability of persistence estimates. In closely related volatility-type models, Monte Carlo evidence shows that small samples can generate sizeable bias and convergence problems, and that several hundred observations may be required to obtain reliable parameter estimates (e.g., at least 500 observations for standard GARCH(1,1) estimation under common non-negativity constraints; see Hwang and Valls Pereira (2006)).

but also to high and low prices, ensuring coherence in the construction of range-based measures of volatility and illiquidity.

The data span the period from December 31, 1997 to January 21, 2026, depending on individual stock availability. To accurately assess data completeness and detect missing observations, I relied on the official NYSE trading calendar. This allows for a precise identification of non-trading days and avoids conflating genuine market closures with missing or corrupted observations.

3.2 Data Quality Control

Beyond standard filters, a detailed data quality assessment was conducted for each stock. For every time series, I computed a set of diagnostic indicators, including: the total number of observations, the number of missing trading days relative to the official calendar, the number of days with non-positive trading volume, the number of days with zero or negative price ranges (High minus Low), the number of missing or non-numeric values, and the maximum observed positive and negative daily returns.

Stocks exhibiting negative prices were excluded from the dataset. In addition, I applied a data quality filter based on the fraction of unusable observations within each time series, defined as the ratio between the number of unusable daily observations and the total number of available trading days for each stock. Unusable observations include days with missing prices, non-positive trading volume, or degenerate price ranges.

A tolerance threshold of 0.5% was imposed on this fraction. This restriction ensures that only stocks with a negligible share of unusable observations are retained, thereby preserving the reliability of the daily illiquidity measures. Stocks exceeding this tolerance level were excluded from the analysis. In borderline cases, time series were manually inspected, and economically implausible or unreliable stocks were removed.

After applying all data quality filters, the final dataset consists of 1,893 stocks with clean daily observations for adjusted Close, High, Low prices and trading volume.

3.3 Illiquidity Measures

This section defines the illiquidity proxies used in the empirical analysis. The economic interpretation of these measures and their connection to the market microstructure notion of price impact are discussed in Chapter 2; here the focus is on notation, construction, and comparability across alternative proxies.

Let $P_{i,t}$ denote the adjusted close price of stock i on day t, and let $\text{Volume}_{i,t}$ denote daily trading volume in shares. Daily dollar volume is computed as $v_{i,t} = P_{i,t} \times \text{Volume}_{i,t}$. Daily simple returns are defined as $R_{i,t} = P_{i,t}/P_{i,t-1} - 1$. For comparisons across illiquidity measures, I also work with log returns $r_{i,t} = \ln(P_{i,t}/P_{i,t-1}) = \ln(1 + R_{i,t})$, which are convenient for time aggregation and scale comparability[11].

Following Amihud (2002), for stock i in year y the annual Amihud illiquidity measure is

$$ILLIQ_{i,y} = \frac{1}{D_{i,y}} \sum_{d=1}^{D_{i,y}} \frac{|R_{i,y,d}|}{v_{i,y,d}}, \tag{1}$$

where $D_{i,y}$ is the number of trading days with available observations in year y, $R_{i,y,d}$ is the daily return, and $v_{i,y,d}$ is daily dollar volume. In my empirical implementation at the daily frequency, the corresponding daily series is the special case of (1) with $D_{i,y} = 1$:

$$\mathcal{A}^D_{i,t} = \frac{|R_{i,t}|}{v_{i,t}}. \tag{2}$$

A practical drawback of (2) is that the numerator uses the close-to-close absolute return. When the close does not change between $t-1$ and t, one has $R_{i,t} = 0$ and the measure mechanically returns zero even if within-day price variation and trading frictions were present. This feature motivates alternative numerators that exploit within-day information when such information is available, or that approximate it using daily data.

When intraday prices are available, I compute the realized Amihud measure proposed by Lacava et al. (2023). The measure replaces the close-to-close absolute return with realized absolute variation (realized power variation of order one) computed from intraday log returns:

$$\mathcal{A}_{i,t} = \frac{\text{RPV}_{i,t}}{v_{i,t}}, \qquad \text{RPV}_{i,t} = \sum_{j=1}^{I_{i,t}} \left| r_{i,t,j} \right|,$$

where $r_{i,t,j}$ denotes the j-th intraday log return for stock i on day t, with the index j running over equally spaced intraday intervals within the trading day. The total number of intraday observations is $I_{i,t}$, which depends on the chosen intraday sampling frequency, and $v_{i,t}$ denotes daily dollar trading volume[12]. In the present setting, intraday data are

[11] In the cross-sectional asset pricing tests (Fama–MacBeth), I employ simple returns, consistently with the standard implementation based on monthly cross-sections of stock returns (Fama and MacBeth, 1973; Fama and French, 1993).

[12] In the empirical implementation, intraday prices are sampled at a 5-minute frequency. Starting from

available only for a restricted subset of stocks and sample periods; accordingly, realized Amihud is used primarily for within-stock comparisons and validation exercises, while the main panel analysis relies on daily proxies that are available for the full cross-section.

To retain a volatility-type numerator while relying only on daily information, I construct a range-based proxy in which the numerator is built from a Parkinson-type range estimator (Parkinson, 1980). Let $H_{i,t}$ and $L_{i,t}$ denote the (adjusted) daily high and low prices, and define log-high and log-low as $p^H_{i,t} = \ln(H_{i,t})$ and $p^L_{i,t} = \ln(L_{i,t})$. The daily range statistic is

$$Range_{i,t} = \sqrt{\frac{1}{4\log(2)} \left(p^H_{i,t} - p^L_{i,t}\right)^2},$$

and the corresponding range-based Amihud proxy is $\mathcal{A}^{HL}_{i,t} = \frac{Range_{i,t}}{v_{i,t}}$. The daily and range-based measures, (2) and $\mathcal{A}^{HL}_{i,t}$, are available for the full panel of stocks over the daily sample. The realized measure $\mathcal{A}_{i,t}$ is computed only when intraday coverage is available and sufficiently reliable; when present, it serves as a benchmark to assess how closely daily proxies track intraday-based variation.

To provide a concrete mapping between the three constructions, Figure 1 reports normalized time series for a representative large-cap stock with intraday coverage (AAPL[13]), while Table 6 complements the visual evidence by summarizing basic distributional moments at daily, weekly, and monthly horizons. Taken together, these diagnostics highlight two empirical regularities that are central for the subsequent modeling choice.

First, the classic daily Amihud is substantially noisier at the daily horizon than both the realized and the range-based measures. This is evident in the time series comparison, where the daily proxy exhibits frequent isolated spikes, and is confirmed by the higher dispersion and coefficient of variation reported in the Table 6. Second, time aggregation systematically reduces dispersion for all three measures. Importantly, when aggregation is performed by averaging daily observations within the period, the reduction in variability is broadly comparable across constructions, suggesting that the underlying illiquidity signal is preserved across horizons.

minute-level prices, the first 15 minutes of trading are discarded to mitigate opening effects, and the remaining observations are aggregated into non-overlapping 5-minute intervals. The use of moderate intraday sampling frequencies such as 5-minute returns is standard in realized volatility work as a trade-off between information efficiency and market microstructure noise; see Hansen and Lunde (2006).

[13]The same comparison is performed for MSFT and IBM. The resulting dynamics are qualitatively very similar and are therefore reported separately in the Figures and Tables section (Figure 4).

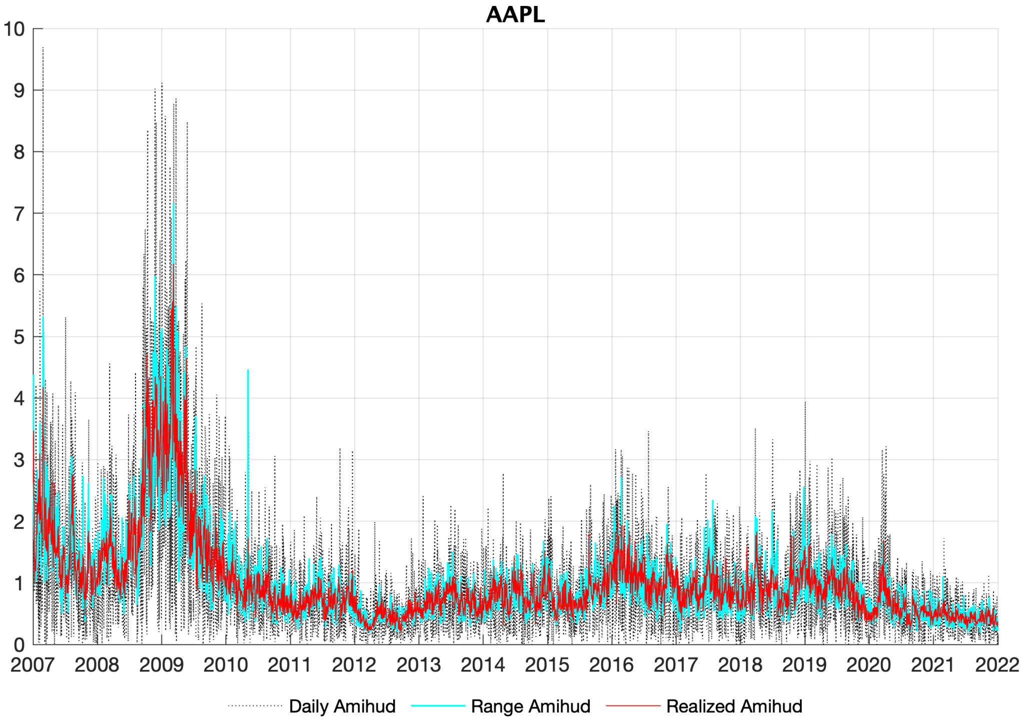

Figure 1: Comparison between the classic daily Amihud illiquidity measure (black dotted line), the range-based Amihud (cyan solid line), and the realized Amihud (red solid line) for AAPL. The sample period spans from January 3, 2007 to December 31, 2021. All illiquidity proxies are normalized by their respective sample means. For each proxy $\mathcal{A}_t^{(j)}$, with $j \in \{\text{daily}, \text{range}, \text{realized}\}$, the normalized series is defined as $\tilde{\mathcal{A}}_t^{(j)} = \mathcal{A}_t^{(j)} / \left(\frac{1}{T}\sum_{t=1}^{T}\mathcal{A}_t^{(j)}\right)$, where T denotes the available sample size for the given stock. This normalization preserves the relative time-series dynamics while rescaling each proxy to have unit mean, facilitating comparability across measures constructed on different scales.

These patterns reinforce the interpretation emerging from Figure 1: measures that embed richer within-day information either directly through intraday returns or indirectly through the daily price range deliver a smoother and more persistent illiquidity signal at the daily frequency, while remaining coherent under temporal aggregation. This evidence motivates treating range-based illiquidity as a realized-type daily measure and supports the use of volatility-style dynamic models to capture its evolution over time.

3.4 Illiquidity Dynamics and MEM

Illiquidity proxies are non-negative by construction and exhibit pronounced persistence over time. Figure 2 documents this feature through autocorrelation functions computed for the daily, range-based, and realized Amihud measures for AAPL[14]. The three panels

[14]The same autocorrelation analysis was conducted for other large-cap stocks with reliable intraday coverage, including MSFT and IBM. Since the qualitative patterns are virtually identical across firms, these

reveal a clear ranking in persistence across constructions. The classic daily Amihud displays relatively low and fairly flat autocorrelations, consistent with substantial measurement noise at the daily horizon. By contrast, the range-based proxy exhibits markedly higher short and medium lag dependence, indicating that incorporating intraday price dispersion through the high–low range substantially smooths the illiquidity signal. The realized Amihud shows the strongest persistence, with autocorrelations close to unity at short lags and a very slow decay over time, reflecting the ability of intraday information to capture a highly persistent latent illiquidity state.

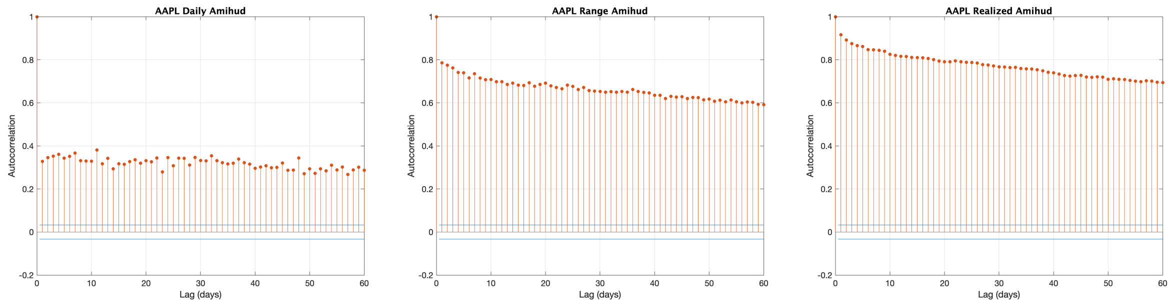

Figure 2: Empirical autocorrelation functions (ACF) of illiquidity measures for AAPL. The left panel reports the classic daily Amihud proxy, the middle panel the range-based Amihud, and the right panel the realized Amihud. The sample period spans from January 3, 2007 to December 31, 2021. The horizontal axis reports lags in trading days (up to 60). The blue bands denote approximate 95% confidence bounds under the white-noise null.

This pattern mirrors the well-documented clustering behavior of realized volatility and is fully consistent with the evidence reported in Lacava et al. (2023). Taken together, these diagnostics suggest that richer within-day information either explicitly through intraday returns or implicitly through range-based measures extracts a smoother and more persistent illiquidity component than the close-to-close proxy. These empirical regularities motivate the adoption of a volatility-style dynamic specification that simultaneously (i) respects the non-negativity of illiquidity measures and (ii) accommodates a slowly evolving conditional component, as formalized by the Multiplicative Error Model framework introduced in the next section.

Following the exposition in Brownlees et al. (2011), I assume that the illiquidity process $\{\mathcal{A}_{i,t}\}$ follows a Multiplicative Error Model (MEM), originally introduced by Engle (2002) and further discussed in Engle and Gallo (2006). Suppressing the stock index i

additional ACF plots are reported separately in the Figures and Tables section (Figure 5).

when no ambiguity arises, the MEM representation is

$$\mathcal{A}_t = \mu_t \, \epsilon_t, \tag{3}$$

where, conditionally on $\mathcal{F}_{t-1}$, μ_t is a strictly positive predictable component and ϵ_t is a positive innovation. The innovation is assumed to satisfy $\mathbb{E}[\epsilon_t \mid \mathcal{F}_{t-1}] = 1$ and $\mathrm{Var}(\epsilon_t \mid \mathcal{F}_{t-1}) = \sigma^2$, which under (3) implies $\mathbb{E}[\mathcal{A}_t \mid \mathcal{F}_{t-1}] = \mu_t$ and $\mathrm{Var}(\mathcal{A}_t \mid \mathcal{F}_{t-1}) = \sigma^2\mu_t^2$.

When $\mathcal{F}_{t-1}$ contains only past values of the series, I adopt the baseline MEM(1,1) recursion

$$\mu_t = \omega + \alpha\mathcal{A}_{t-1} + \beta\mu_{t-1}, \tag{4}$$

with $\omega > 0$, $\alpha \geq 0$, and $\beta \geq 0$ to preserve positivity. This specification mirrors the familiar generalized ARCH dynamics (Bollerslev, 1986), with α capturing the immediate impact of the most recent realization and β capturing the inertia of the predictable component. If the process is mean-stationary, unconditional expectations yield $\mathbb{E}[\mathcal{A}_t] = \mathbb{E}[\mu_t] \equiv \mu$ and the standard relationship $\omega = (1 - \alpha - \beta)\mu$, valid when $\alpha + \beta < 1$.

To obtain a fully parametric likelihood, I specify a Gamma innovation with unit mean using the shape–scale parametrization

$$\epsilon_t \mid \mathcal{F}_{t-1} \sim \Gamma\left(\theta, \frac{1}{\theta}\right),$$

which implies $\mathbb{E}[\epsilon_t \mid \mathcal{F}_{t-1}] = 1$ and $\mathrm{Var}(\epsilon_t \mid \mathcal{F}_{t-1}) = 1/\theta$, hence $\sigma^2 = 1/\theta$. Parameters are estimated by maximizing the corresponding conditional log-likelihood under the positivity constraints on (ω, α, β) and $\theta > 0$. For each stock and each illiquidity proxy (daily, range-based, and realized when available), the likelihood is constructed using all valid daily observations after the data-quality filters described in the previous subsection; missing days are treated as missing observations and are excluded from the likelihood summation. The MEM representation provides a direct decomposition of observed illiquidity into a predictable component and an innovation. In the empirical analysis, the fitted conditional mean $\widehat{\mu}_{i,t}$ is interpreted as *expected illiquidity* for stock i at time t. The standardized innovation is obtained as

$$\widehat{\epsilon}_{i,t} = \frac{\mathcal{A}_{i,t}}{\widehat{\mu}_{i,t}},$$

and captures *unexpected illiquidity* shocks, i.e., deviations from the predictable illiquidity state. This decomposition is then used to construct illiquidity-based signals and factors, and to assess their pricing implications in the cross-section.

3.5 Testing for an Illiquidity Premium

To test whether illiquidity risk is priced in the cross-section of stock returns, I construct tradable illiquidity factors and estimate risk premia using the two-step Fama–MacBeth procedure (Fama and MacBeth, 1973; Fama and French, 1993). All asset-pricing tests are implemented at the *monthly* horizon and are repeated separately for (i) the expected and unexpected components of illiquidity extracted from the MEM and (ii) two underlying illiquidity proxies (daily Amihud and range-based Amihud).

Illiquidity signals are measured at the daily frequency through the MEM-implied quantities $\widehat{\mu}_{i,t}$ and $\widehat{\epsilon}_{i,t}$, but the return horizon used for factor construction and cross-sectional tests is monthly. Let τ index consecutive non-overlapping calendar months. Denote by N_τ the number of trading days contained in month τ, which may vary across months due to the trading calendar. I define month-τ illiquidity signals by averaging the daily MEM-implied quantities over all trading days in τ:

$$\overline{\mu}_{i,\tau} = \frac{1}{N_\tau} \sum_{t \in \tau} \widehat{\mu}_{i,t}, \qquad \overline{\epsilon}_{i,\tau} = \frac{1}{N_\tau} \sum_{t \in \tau} \widehat{\epsilon}_{i,t}.$$

Monthly returns are computed over the same calendar months from end-of-month adjusted prices. This design preserves the informational content of the daily illiquidity signal extracted from the MEM while aligning the signal horizon with the return horizon used in the asset-pricing tests. The aggregation principle parallels the treatment of realized measures in Andersen et al. (2003), who emphasize that realized quantities can be coherently summarized over longer horizons without requiring a re-specification of the underlying dynamics[15]. At the monthly horizon, I form decile portfolios based on the *lagged* MEM-implied illiquidity signals. Specifically, for a given month τ, stocks are sorted into ten groups according to either $\overline{\mu}_{i,\tau-1}$ (*expected illiquidity*) or $\overline{\epsilon}_{i,\tau-1}$ (*unexpected*

[15] In the realized volatility setting, temporal aggregation is justified by continuous-time quadratic variation arguments (Andersen et al., 2003). In this thesis, intraday-based realized illiquidity is not available on a broad panel; the range-based proxy is therefore treated as a realized-type daily signal and is assumed to inherit similar aggregation properties in practice.

illiquidity). Let $R_\tau^{(k)}(\overline{\mu}_{i,\tau-1})$ denote the *equal-weighted*[16] simple return over month τ of the stocks that fall in decile k when sorting on $\overline{\mu}_{i,\tau-1}$. Analogously, $R_\tau^{(k)}(\overline{\epsilon}_{i,\tau-1})$ denotes the equal-weighted return of decile k when sorting on $\overline{\epsilon}_{i,\tau-1}$. In particular, $R_\tau^{(10)}(\cdot)$ (*respectively* $R_\tau^{(1)}(\cdot)$) refers to the decile containing stocks with the highest (*lowest*) signal realizations. The corresponding tradable illiquidity factors are then defined as high-minus-low spreads:

$$\mathrm{EI}_\tau \equiv R_\tau^{(10)}(\overline{\mu}_{i,\tau-1}) - R_\tau^{(1)}(\overline{\mu}_{i,\tau-1}), \qquad \mathrm{UI}_\tau \equiv R_\tau^{(10)}(\overline{\epsilon}_{i,\tau-1}) - R_\tau^{(1)}(\overline{\epsilon}_{i,\tau-1}), \tag{5}$$

where the labels EI and UI denote factors built from the *expected* and *unexpected illiquidity* signals, respectively. The one-period lag in the sorting signal enforces a predictive design: information in month $\tau - 1$ is used to form portfolios whose returns are realized over month τ. The same construction is repeated for both underlying illiquidity proxies, to make this explicit, I denote by EI_τ^{HL} and UI_τ^{HL} the factors obtained from the range-based (high–low) proxy, and by EI_τ^{D} and UI_τ^{D} the corresponding factors obtained from the daily-based proxy. In each case, the sorting signals $\overline{\mu}_{i,\tau-1}$ and $\overline{\epsilon}_{i,\tau-1}$ are constructed by aggregating the daily MEM-implied quantities over the trading days contained in period $\tau - 1$. Having defined the tradable illiquidity factors, I evaluate their asset-pricing implications in two ways. I first use time-series regressions of the decile portfolios and of the high-minus-low factors on standard benchmarks (CAPM, Fama–French three factors, and the Carhart four-factor model) to summarize abnormal performance (alphas) and factor sensitivities (betas). I then turn to cross-sectional tests based on the Fama–MacBeth methodology (Fama and MacBeth, 1973).

Formally, for each sorting signal and for each underlying illiquidity proxy, I run standard time-series pricing regressions on the returns of the illiquidity-sorted decile portfolios and on the corresponding self-financing high-minus-low spread. For each portfolio k in the decile set (with k indicating the decile rank) and for the corresponding spread portfolio constructed from the top and bottom deciles, I estimate the standard linear factor

[16] Decile-portfolio returns, denoted generically by $R_\tau^{(k)}(\cdot)$, are computed as simple equal-weighted averages of constituent stock returns within each group. A value-weighted (market-cap-weighted) construction is often preferred in the asset-pricing literature (see, e.g., Fama and French (1993)), but it requires time-varying firm characteristics measured at the same horizon and updated each period—most importantly market capitalization and, for characteristic-based refinements, accounting variables such as book equity. These inputs are typically obtained from CRSP/Compustat and related academic datasets, but are not available in the present setting at the required frequency and coverage. Given the objective of this thesis and the available data, I therefore implement equal-weighted decile sorts.

pricing regression of excess returns on the benchmark factors.

$$R_{\tau}^{(k)}(\cdot) - R_{f,\tau} = \alpha^{(k)} + \boldsymbol{\beta}^{(k)\top}\mathbf{F}_{\tau} + \varepsilon_{\tau}^{(k)}, \tag{6}$$

where $\mathbf{F}_{\tau}$ denotes the benchmark factor vector, set to {MKT−RF} under the CAPM (Sharpe, 1964), to {MKT−RF, SMB, HML} under FF3 (Fama and French, 1993), and augmented with MOM under the Carhart four-factor model (Carhart, 1997). The intercept $\alpha^{(k)}$ is Jensen's alpha (Jensen, 1968), capturing benchmark-adjusted performance, namely the portion of the portfolio's average excess return that cannot be attributed to its exposures to the benchmark factor set. The primary portfolio-level object of interest is the alpha of the long–short spread (the 10−1 portfolio), because it directly tests whether the tradable illiquidity strategy delivers a return premium beyond standard equity risk exposures. At the same time, the estimated factor loadings $\boldsymbol{\beta}^{(k)}$ provide a compact characterization of the systematic tilts induced by sorting on illiquidity (e.g., market, size, value, and momentum/reversal exposures). This helps assess whether any return gradient across deciles primarily reflects compensation for conventional risk factors or whether it contains an additional component potentially attributable to illiquidity risk. These time-series regressions therefore serve as portfolio-level diagnostics and performance attribution; to directly estimate whether exposure to illiquidity is rewarded in the cross-section of individual stocks,

I next implement Fama–MacBeth pricing tests. In the first step, I estimate factor betas for each stock i in the sample (1,893 U.S. equities) from the time-series regression

$$R_{i,\tau} - R_{f,\tau} = \alpha_i + \beta_i^{\text{Illiq}}\,\text{Illiq}_{\tau} + \boldsymbol{\beta}_i^{\top}\mathbf{F}_{\tau} + u_{i,\tau}, \tag{7}$$

where $R_{i,\tau} - R_{f,\tau}$ is the excess return[17] of stock i over period τ, Illiq_{τ} denotes the illiquidity factor under consideration (either EI_{τ}^{HL} or UI_{τ}^{HL} when using the range-based proxy $\mathcal{A}^{HL}$, and either EI_{τ}^{D} or UI_{τ}^{D} when using the daily-based proxy $\mathcal{A}^{D}$), and $\mathbf{F}_{\tau}$ is the benchmark factor vector defined above.

In the second step, I estimate the price of illiquidity risk by running cross-sectional

[17]In the regression analysis, individual stock returns are expressed in excess-return form whenever required by the pricing specification. This choice is *inconsequential* for the long–short factors, because subtracting the risk-free rate from both legs cancels mechanically: $\big(R_{\tau}^{(10)}(\cdot) - R_{f,\tau}\big) - \big(R_{\tau}^{(1)}(\cdot) - R_{f,\tau}\big) = R_{\tau}^{(10)}(\cdot) - R_{\tau}^{(1)}(\cdot)$. Accordingly, EI_{τ} and UI_{τ} can be interpreted as self-financing factor returns that are directly comparable to other traded factors.

regressions at each time τ, regressing realized excess returns on the estimated loadings from (7):

$$R_{i,\tau} - R_{f,\tau} = \gamma_{0,\tau} + \gamma_\tau^{\text{Illiq}} \widehat{\beta}_i^{\text{Illiq}} + \boldsymbol{\gamma}_\tau^\top \widehat{\boldsymbol{\beta}}_i + \eta_{i,\tau}, \tag{8}$$

where $\gamma_\tau^{\text{Illiq}}$ is the period-τ price of illiquidity risk, and $\widehat{\boldsymbol{\beta}}_i$ collects the estimated factor loadings associated with $\mathbf{F}_\tau$. The illiquidity risk premium is obtained as the time-series average $\overline{\gamma}^{\text{Illiq}} = \frac{1}{T} \sum_\tau \gamma_\tau^{\text{Illiq}}$. Inference is based on the time-series variability of $\gamma_\tau^{\text{Illiq}}$, using the usual Fama–MacBeth t-statistic:

$$t\left(\overline{\gamma}^{\text{Illiq}}\right) = \frac{\overline{\gamma}^{\text{Illiq}}}{\widehat{\sigma}\left(\gamma_\tau^{\text{Illiq}}\right) / \sqrt{T}}, \qquad \widehat{\sigma}^2\left(\gamma_\tau^{\text{Illiq}}\right) = \frac{1}{T-1} \sum_{\tau=1}^{T} \left(\gamma_\tau^{\text{Illiq}} - \overline{\gamma}^{\text{Illiq}}\right)^2.$$

The null hypothesis is $H_0 : \overline{\gamma}^{\text{Illiq}} = 0$ (no illiquidity premium), against the two-sided alternative $H_1 : \overline{\gamma}^{\text{Illiq}} \neq 0$. A statistically significant $\overline{\gamma}^{\text{Illiq}}$ provides evidence that exposure to the illiquidity factor is systematically rewarded in expected returns. Comparing results across EI and UI further allows assessing whether pricing differs between the predictable and shock components of illiquidity, while contrasting the HL- and D-based versions highlights the role of the underlying proxy used to construct the signal.

The next section presents the empirical evidence in the same order as the testing strategy, starting with the decile-sort portfolios and the associated spread factors. It then turns to the cross-sectional results from the Fama–MacBeth estimates of $\overline{\gamma}^{\text{Illiq}}$.

CHAPTER 4

EMPIRICAL RESULTS

4.1 Dynamic properties of illiquidity

In 3.4 I introduced the MEM framework for modeling illiquidity. I now discuss what the empirical estimates imply for illiquidity dynamics, focusing first on three highly liquid stocks (AAPL, MSFT, IBM) as benchmark cases. These stocks are chosen for their reliable intraday data coverage and representative trading activity, which allow a clean comparison across alternative illiquidity measures without cherry-picking results. Each stock's illiquidity is quantified in three ways: a Daily price-impact proxy (in the spirit of Amihud (2002)), a range-based proxy (using daily high–low price range), and a Realized measure constructed from intraday data. The MEM(1,1) model is then estimated for each series to capture the temporal dependence in illiquidity.

	MEM Estimation								
	Daily Amihud ($\mathcal{A}_t^D$)			Range Amihud ($\mathcal{A}_t^{HL}$)			Realized Amihud ($\mathcal{A}_t$)		
Parameters	AAPL	MSFT	IBM	AAPL	MSFT	IBM	AAPL	MSFT	IBM
ω	0.0000^c	0.0000^c	0.0000^a	0.0000^a	0.0000^b	0.0000^a	0.0000^a	0.0000^a	0.0000^a
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
α	0.0437^a	0.0555^a	0.0766^a	0.2003^a	0.1912^a	0.2576^a	0.3500^a	0.3847^a	0.4389^a
	(0.0055)	(0.0078)	(0.0094)	(0.0159)	(0.0164)	(0.0192)	(0.0199)	(0.0228)	(0.0260)
β	0.9538^a	0.9421^a	0.9133^a	0.7917^a	0.8038^a	0.7199^a	0.6387^a	0.6068^a	0.5243^a
	(0.0060)	(0.0085)	(0.0113)	(0.0168)	(0.0173)	(0.0222)	(0.0209)	(0.0239)	(0.0301)
θ	1.5314^a	1.4819^a	1.5358^a	10.5512^a	9.1136^a	8.7494^a	22.5212^a	18.3029^a	15.7061^a
	(0.0351)	(0.0321)	(0.0362)	(0.2568)	(0.2193)	(0.2068)	(0.5407)	(0.4723)	(0.4129)
LogLik	94788.86	91326.14	88654.16	98990.31	95076.05	91556.79	92901.83	88777.04	85071.04
	Ljung–Box statistics								
LB1	0.0899	0.0602	0.0006	0.0238	0.0029	0.0159	0.0000	0.0001	0.0017
LB5	0.0802	0.0205	0.0000	0.0262	0.0031	0.0287	0.0000	0.0000	0.0000
LB10	0.0989	0.0034	0.0005	0.0003	0.0016	0.0005	0.0000	0.0000	0.0000

Table 1: The table reports MEM(1,1) estimates for Daily, Range, and Realized Amihud illiquidity measures for AAPL, MSFT, and IBM. The estimation period spans from January 3, 2007 to December 31, 2021. Parameter estimates are reported with robust standard errors in parentheses. Statistical significance is denoted by superscripts a, b, and c, corresponding to $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively. The Ljung–Box statistics are reported as p-values for the null hypothesis of no residual autocorrelation up to lags 1, 5, and 10 (computed on ϵ_t). LogLik denotes the maximized conditional log-likelihood.

Table 1 reports the estimated MEM(1,1) parameters for AAPL, MSFT, and IBM under each illiquidity proxy. Overall, the MEM delivers well-behaved estimates across all three constructions: $\widehat{\alpha}$ and $\widehat{\beta}$ are highly statistically significant throughout, consistent with the strong persistence typically found in volatility-style dynamics (Engle and Bollerslev, 1986) and with the broader evidence on liquidity dynamics and commonality (e.g., Chordia et al., 2000; Pástor and Stambaugh, 2003). The only coefficient that shows comparatively weaker statistical support is the intercept $\widehat{\omega}$ for the Daily Amihud specification, which is significant only at the 10% level for AAPL and MSFT (while remaining strongly significant for IBM). This pattern is consistent with the idea that the close-to-close daily proxy contains more measurement noise, so that the model's long-run level parameter is identified less sharply than the dynamic feedback terms.

A clear pattern emerges when comparing estimates across the three measurement approaches. Moving from the $\mathcal{A}^D$ to the $\mathcal{A}^{HL}$ to the $\mathcal{A}$, the MEM parameters shift toward stronger short-run responsiveness and weaker inertia. Specifically, $\widehat{\alpha}$ which governs how strongly the conditional mean reacts to the most recent realization, increases as the proxy incorporates richer within-day information, while $\widehat{\beta}$ which captures the persistence of the conditional mean, declines.

For AAPL, Table 1 illustrates this reallocation sharply. Under the Daily proxy, the dynamics are dominated by inertia ($\widehat{\beta} = 0.9538$) with only a limited short-run response ($\widehat{\alpha} = 0.0437$). Under the Realized proxy, the model reacts much more to the latest observation ($\widehat{\alpha} = 0.3500$) and relies less on inertia ($\widehat{\beta} = 0.6387$). In both cases, $\widehat{\alpha} + \widehat{\beta}$ remains close to one, indicating that the conditional mean is highly persistent and adjusts only gradually. MSFT and IBM display the same qualitative pattern: Range and Realized proxies are associated with higher $\widehat{\alpha}$ and lower $\widehat{\beta}$ relative to the Daily proxy, suggesting that the result is not stock-specific. More generally, values of $\widehat{\alpha} + \widehat{\beta}$ close to unity across specifications imply that illiquidity shocks dissipate slowly, consistent with the persistence typically observed in other market frictions. (Engle, 1982; Bollerslev, 1986)

This systematic shift in $\widehat{\alpha}$ and $\widehat{\beta}$ can be given an economic interpretation. A proxy that incorporates more information (such as intraday data for realized volatility or the full daily price range) provides a less noisy signal of underlying illiquidity on each day. The MEM then places more weight on the latest observation (higher $\widehat{\alpha}$) and relies less on persistence from past values (lower $\widehat{\beta}$) to fit the observed variation. In contrast, a noisier proxy (for example, a daily price-impact measure based on only the closing price and daily volume) induces the model to smooth more heavily over time, effectively compensating for measurement error by way of a higher $\widehat{\beta}$. This finding aligns with the literature on realized volatility measures providing cleaner signals than daily return-based estimates. Just as using intraday data improves volatility estimates (Andersen et al., 2003; Barndorff-Nielsen and Shephard, 2002), using richer illiquidity proxies yields more responsive dynamics. Similarly, range-based estimators have long been recognized to contain more information about variability than single-point daily changes (Parkinson, 1980), which here translates into the model reacting more to each new range-based observation.

Another noteworthy result concerns the estimated innovation-shape parameter $\widehat{\theta}$. Under the unit-mean Gamma specification introduced in Section 3.4, $\widehat{\theta}$ directly controls the dispersion of the multiplicative innovation ϵ_t (Cipollini et al., 2009): higher values of $\widehat{\theta}$ correspond to lower conditional variance, since $\mathrm{Var}(\epsilon_t) = 1/\theta$ under the $\Gamma(\theta, 1/\theta)$ parametrization. In practical terms, $\widehat{\theta}$ provides a compact summary of how "tight" the unexpected component is around its unit mean once the predictable component μ_t has been filtered out.

The estimates in Table 1 show a clear monotonic pattern across illiquidity constructions. For AAPL, the Daily proxy implies $\widehat{\theta} = 1.5314$, which translates into $\widehat{\mathrm{Var}}(\epsilon_t) \approx 0.653$. Moving to the Range proxy, $\widehat{\theta}$ increases to 10.5512, implying a much smaller innovation variance, $\widehat{\mathrm{Var}}(\epsilon_t) \approx 0.095$. Finally, under the Realized proxy, $\widehat{\theta}$ rises further to 22.5212, corresponding to $\widehat{\mathrm{Var}}(\epsilon_t) \approx 0.044$. The same ranking appears for MSFT and IBM: $\widehat{\theta}$ is lowest for the Daily proxy, substantially higher for the range-based construction, and highest when intraday information is used.

This ordering is economically coherent with the earlier interpretation of proxy quality. When illiquidity is measured using only close-to-close variation, the observed series contains more mechanical discreteness and idiosyncratic spikes, so the model leaves a larger share of the residual variability in the innovation term. As the proxy incorporates richer within-day information (either indirectly through the high–low range or directly through intraday returns), the observed series becomes a cleaner signal of the underlying illiquidity state, and the MEM can absorb more of its variation into the predictable component, leaving a lower dispersion of multiplicative shocks. In this sense, the range-based proxy behaves much closer to the realized proxy than to the daily one: its $\widehat{\theta}$ (and therefore the implied innovation variance) sits between the two, but it is quantitatively far nearer to the realized case than to the daily proxy.

Residual autocorrelation diagnostics (Ljung–Box tests on $\widehat{\epsilon}_t$) are reported for completeness in the lower panel of Table 1; they do not affect the main conclusion that the relative positioning of Daily, Range, and Realized proxies is primarily reflected in the systematic shifts in $(\widehat{\alpha}, \widehat{\beta})$ and in the innovation dispersion captured by $\widehat{\theta}$.

4.2 Cross-sectional evidence on illiquidity dynamics

Having examined three benchmark stocks in detail, I now investigate whether the same dynamic patterns hold across a broad cross-section of stocks. Understanding the generality of these findings is important in light of evidence that liquidity has common systematic components across assets (Chordia et al., 2000; Pástor and Stambaugh, 2003; Acharya and Pedersen, 2005). For this part of the analysis, I estimate the MEM(1,1) model for each of the 1,893 stocks in the sample. To make this large-scale estimation feasible, and because intraday data for the realized measure are not available for all stocks, I focus on the Daily ($\mathcal{A}^D$) and Range ($\mathcal{A}^{HL}$) illiquidity proxies in the cross-sectional study. In other words, each stock's illiquidity time series (daily proxy and range-based proxy) is modeled separately to obtain a distribution of $\widehat{\omega}$, $\widehat{\alpha}$, $\widehat{\beta}$, and $\widehat{\theta}$ estimates across the entire universe of stocks.

Figure 3 summarizes the cross-sectional distributions of the estimated MEM parameters for the two proxies. The top row corresponds to the Range-based illiquidity measure and the bottom row to the Daily measure. Each panel shows the density of a given parameter estimate across the 1,893 stocks. Several clear differences emerge between the Range and Daily distributions, echoing the benchmark-stock results from Section 4.1. First, consider the panel corresponding to $\widehat{\alpha}$ in Figure 3. The distribution for the Range proxy is shifted noticeably to the right of the distribution for the Daily proxy. In fact, the median $\widehat{\alpha}$ under the Range measure exceeds that of the Daily measure by approximately 0.10, and the entire Range $\widehat{\alpha}$ distribution is skewed toward larger values. This indicates that, for most stocks, illiquidity reacts more strongly to recent innovations when it is measured by the high–low range. On the other hand, the $\widehat{\beta}$ estimates tend to be lower for the Range-based models than for the Daily-based models. The Daily $\widehat{\beta}$ distribution is tightly

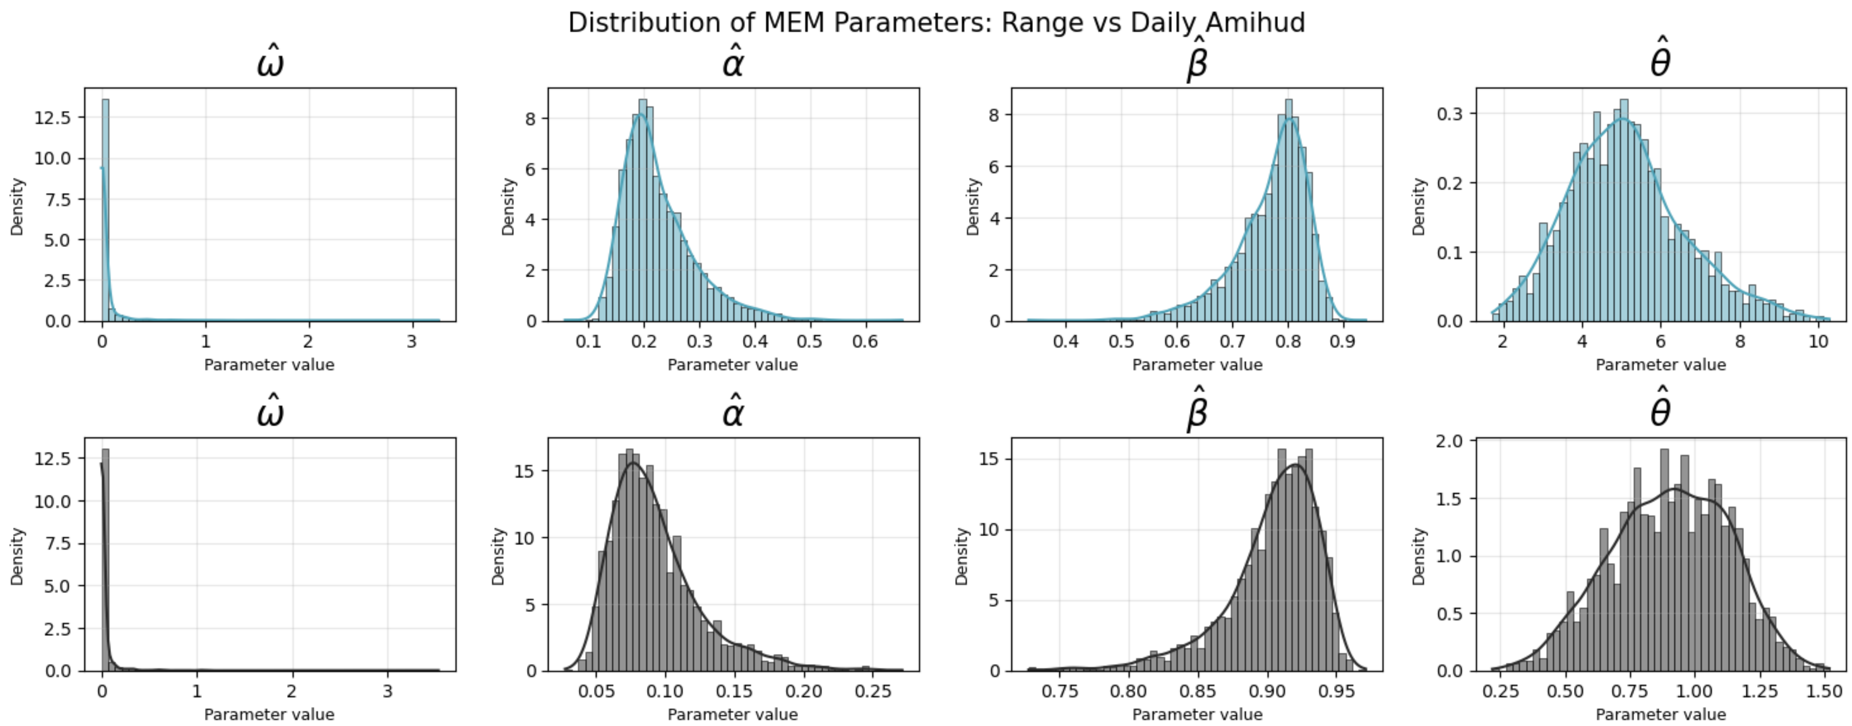

Figure 3: The figure reports kernel-smoothed histograms of the MEM parameters $\widehat{\omega}$, $\widehat{\alpha}$, $\widehat{\beta}$, and $\widehat{\theta}$ estimated at the daily frequency for each stock in a cross-section of 1,893 U.S. equities. The top row refers to the range-based Amihud specification, while the bottom row refers to the daily Amihud specification. The parameter $\widehat{\omega}$ is rescaled by a factor 10^9 to improve readability. The overall data coverage spans from December 31, 1997 to January 21, 2026, with effective sample periods determined by the stock-level selection criteria described in Section 3.

clustered towards the upper end (many stocks have $\widehat{\beta}$ around 0.9 or above), whereas the Range $\widehat{\beta}$ distribution is centered a bit lower (around 0.8) and is more dispersed. In particular, a significant subset of stocks show substantially lower persistence when illiquidity is measured via range, as seen by the left tail of the Range $\widehat{\beta}$ density extending to smaller values that are rare under the Daily measure. Despite these shifts, it remains true that $\widehat{\alpha} + \widehat{\beta}$ is high for virtually all stocks under both measures (typically on the order of 0.95 to 0.99) underscoring that the dominant feature of stock-level illiquidity is a high persistence or slow mean reversion. This pervasive persistence is reminiscent of other financial market variables, such as volatility, that exhibit long-lasting effects of shocks (Andersen et al., 2003) and also aligns with the notion that liquidity supply/demand shocks are persistent and sometimes common across the market[18].

The shape parameter $\widehat{\theta}$ distributions (right column of Figure 3) provide additional insight. Consistently with the single-stock analysis, $\widehat{\theta}$ tends to be larger when using the Range proxy (top-right panel) than with the Daily proxy (bottom-right panel) for the majority of stocks. The Range $\widehat{\theta}$ distribution is centered around higher values (e.g., a median around 5) compared to the Daily $\widehat{\theta}$ (with a median below 1). This suggests that the multiplicative errors in the MEM are less variable (relative to the mean) when illiquidity is measured by the high–low range. In practical terms, a larger $\widehat{\theta}$ implies that the idiosyncratic component of illiquidity (the ϵ_t in the MEM) is more tightly concentrated around 1, meaning fewer large shocks relative to the conditional expectation. The fact that this occurs with the Range measure indicates that using a more informative proxy filters out more of the noise, leaving a cleaner signal of the underlying liquidity condition. Stated differently, the model's uncertainty (unpredictable volatility of illiquidity) is lower when the input data contain more information about true liquidity fluctuations. This pattern is analogous to findings in volatility modeling, where realized measures based on high-frequency data are shown to reduce measurement error relative to daily return-based proxies[19].

Beyond the differences in central tendency between proxies, Figure 3 highlights con-

[18] Evidence of strong persistence and common components in stock-level liquidity is well documented in the market microstructure literature. In particular, Chordia et al. (2000) show that several liquidity measures display significant co-movements over time at both the market and industry level, even after controlling for firm-specific determinants, indicating that liquidity shocks contain an important persistent common component.

[19] High-frequency-based volatility measures are known to reduce measurement error relative to daily return-based proxies (Andersen et al., 2003).

siderable heterogeneity in illiquidity dynamics across stocks. The width of each distribution (and the presence of fat tails) implies that while most stocks cluster around a common behavior, there are many outliers. For instance, a few stocks have extremely high $\widehat{\alpha}$ estimates (far right tail of the $\widehat{\alpha}$ densities), indicating that for those stocks illiquidity responds almost instantaneously to new information (perhaps those are stocks with very high trading activity and transparent order books). Conversely, a few stocks show unusually low $\widehat{\alpha}$ and high $\widehat{\beta}$, suggesting nearly unit-root illiquidity processes where shocks dissipate only very slowly. There is also variation in $\widehat{\omega}$, despite its small absolute scale: some stocks exhibit a higher intercept (after scaling), which may reflect a higher long-run average level of illiquidity, potentially associated with smaller or less-traded companies[20]. This cross-sectional dispersion is economically meaningful; it likely reflects differences in firm characteristics such as trading frequency, volatility, and investor base, all of which can influence liquidity dynamics. For example, larger firms with active markets may exhibit more stable liquidity (lower $\widehat{\beta}$) and faster reactions to new information (higher $\widehat{\alpha}$), whereas small-cap stocks may display more persistent illiquidity due to sporadic trading. A thorough investigation of these determinants is beyond the scope of this subsection, but the patterns in the figure lay the groundwork for linking estimated MEM parameters to observable firm attributes in subsequent analysis.

Despite the substantial dispersion in parameter estimates, the qualitative differences between the Daily and Range proxies are remarkably stable across the cross-section. The shift toward higher $\widehat{\alpha}$ and lower $\widehat{\beta}$ observed for the Range measure, already documented for the benchmark stocks in Section 4.1, is therefore not confined to a small set of large-cap names but represents a pervasive feature of the data. This robustness is consistent with a common underlying mechanism: as measurement error in illiquidity is reduced through the use of more information-rich proxies, liquidity dynamics become more responsive to new information and slightly less persistent. The uniformity of this pattern across stocks aligns with theories of liquidity commonality, according to which assets are jointly affected by market-wide liquidity conditions (Chordia et al., 2000) or funding liquidity constraints (Acharya and Pedersen, 2005). Taken together, these results sug-

[20]Cross-sectional differences in stock illiquidity are systematically related to firm characteristics such as size and trading activity, with smaller and less frequently traded stocks exhibiting higher illiquidity. See Amihud (2002) for evidence that illiquidity is strongly negatively related to firm size and affects small stocks more strongly.

gest that the MEM(1,1) framework captures the essential features of illiquidity dynamics across a broad set of stocks and provides a suitable baseline for examining how firm-level characteristics shape liquidity behavior in subsequent analysis.

4.3 Illiquidity Premium

Section 4.2 showed that the MEM yields stable and economically meaningful estimates of illiquidity dynamics across a broad cross-section of stocks. I now turn from dynamics to asset-pricing implications. Using the model-based decomposition into a predictable component (*expected illiquidity*) and an innovation component (*unexpected illiquidity*), I construct monthly decile-sorted portfolios and the associated high-minus-low (10–1) traded spreads. This long–short construction follows the standard portfolio-sorting approach in empirical asset pricing, where extreme-sorted spreads are used to summarize the return differential associated with a characteristic while maintaining a diversified, self-financing trading strategy (e.g., Fama and French, 1993; Carhart, 1997; Amihud, 2002).

Table 2 reports descriptive statistics for the illiquidity-sorted decile portfolios and the corresponding 10–1 spreads over April 1998 to December 2025. Clear differences emerge between factors built from the expected component of illiquidity and those built from the unexpected component. For *expected illiquidity*, the range-based proxy delivers an essentially monotone profile: average returns rise from about 0.62% per month in decile 1 to 2.49% in decile 10. The daily-based proxy shows the same upward gradient overall (roughly 0.65% to 2.45%), although small reversals appear in a few middle deciles. Dispersion increases in parallel under both proxies, with standard deviations moving from about 4.8% in the lowest-signal portfolio to about 8.0% in the highest-signal portfolio. The long–short spreads are economically large already at the unconditional level: the expected 10–1 factors average 1.88% per month (range-based) and 1.80% (daily-based), with Sharpe ratios around 0.34 and 0.33. Taken together, these moments indicate that the MEM-implied conditional level of illiquidity generates a stable cross-sectional ranking that translates into a sizeable return differential.

Once the sorting signal is the *unexpected* component, the decile pattern becomes visibly weaker and the spreads shrink. A mechanical feature is important in the daily-based case: the low-signal portfolio does not act as a consistently low-return leg, since decile 1 earns about 1.11% per month on average. This compresses the high-minus-low difference and results in a mean 10–1 return of 0.48% per month and a Sharpe ratio of about 0.12. The range-based proxy is less affected, but the spread remains moderate at about 0.74% per month (Sharpe around 0.21). These outcomes are consistent with the interpretation that innovations are less persistent than conditional levels, so rankings move more from month to month and the long–short construction becomes more sensitive to the behavior of the low-signal leg.

Distributional measures reinforce this contrast, since for the expected sorts skewness moves from slightly negative in the low-signal portfolios to clearly positive in the high-signal portfolios and the 10–1 spreads are strongly positively skewed (about 0.92 for the range-based proxy and about 0.88 for the daily-based proxy). Kurtosis for the expected spreads is *lower* than for the individual deciles (4.51 for range and 4.32 for daily), which is consistent with some tail diversification in the long–short return. The min–max ranges point in the same direction: while the expected decile portfolios can experience minima ranging roughly from −17% to about −28%, the expected 10–1 spreads have markedly tighter lower tails, with minima clustered around −10.6% under both proxies, while still reaching large positive months of around 26%.

Less regular patterns emerge for the unexpected sorts, and the distributional evidence points to a heavier role of infrequent extreme months in shaping long–short performance. The 10–1 spreads exhibit much higher kurtosis than any decile portfolio (10.28 for the range-based proxy and 15.19 for the daily-based proxy), implying that a small number of large realizations weigh heavily on the factor return. Tail outcomes convey the same message: the daily-based unexpected spread reaches a deep minimum (−26.56%) while its maximum is comparatively lower (about 21.60%). This behavior aligns with the nature of the signal, since the unexpected component is a transitory deviation from the model-implied conditional level; rankings can reshuffle sharply across months, the long and short legs are less stable, and the spread becomes more exposed to rare but sizable moves.

	Decile portfolios										Spread
	1	2	3	4	5	6	7	8	9	10	10–1
Panel A:	*Expected illiquidity (range-based)*										
Mean (%)	0.62	0.73	0.74	0.86	0.86	0.94	1.07	1.09	1.22	2.49	1.88
Std. Dev. (%)	4.81	5.08	5.48	5.88	5.92	6.25	6.33	6.73	6.95	8.07	5.46
Sharpe	0.13	0.14	0.14	0.15	0.14	0.15	0.17	0.16	0.18	0.31	0.34
Skewness	-0.30	-0.18	-0.15	-0.05	0.08	0.22	0.11	0.42	0.47	0.58	0.92
Kurtosis	4.82	6.23	6.01	6.78	5.83	6.22	5.67	6.93	5.93	5.42	4.51
Min (%)	-17.26	-21.53	-24.97	-27.09	-26.64	-25.55	-26.77	-27.59	-24.55	-27.38	-10.68
Max (%)	17.54	23.3	22.42	25.82	25.65	29.92	27.5	32.54	30.24	34.12	25.50
Panel B:	*Expected illiquidity (daily-based)*										
Mean (%)	0.65	0.72	0.75	0.86	0.82	0.95	1.12	1.02	1.27	2.45	1.80
Std. Dev. (%)	4.79	5.07	5.43	5.92	5.96	6.25	6.46	6.65	6.94	8.03	5.46
Sharpe	0.14	0.14	0.14	0.14	0.14	0.15	0.17	0.15	0.18	0.31	0.33
Skewness	-0.27	-0.15	-0.16	-0.05	0.11	0.13	0.24	0.26	0.50	0.56	0.88
Kurtosis	4.83	6.14	6.31	6.40	6.15	6.13	6.05	6.22	6.11	5.34	4.32
Min (%)	-17.19	-21.68	-25.24	-26.59	-26.11	-27.93	-25.06	-28.46	-24.13	-27.34	-10.64
Max (%)	17.85	22.97	24.04	25.02	27.81	26.81	31.94	29.54	30.68	33.61	25.80
Panel C:	*Unexpected illiquidity (range-based)*										
Mean (%)	0.90	0.81	0.84	0.91	1.00	0.94	1.13	1.11	1.33	1.64	0.74
Std. Dev. (%)	5.66	5.57	5.61	5.89	5.73	5.97	6.13	6.27	6.55	7.20	3.59
Sharpe	0.16	0.15	0.15	0.15	0.17	0.16	0.18	0.18	0.20	0.23	0.21
Skewness	-0.10	0.13	0.02	0.26	0.08	0.15	0.14	-0.02	0.08	0.80	0.63
Kurtosis	5.11	6.11	6.34	5.57	5.58	6.01	5.88	6.14	6.28	7.86	10.28
Min (%)	-24.51	-24.14	-24.22	-22.00	-24.64	-23.39	-24.25	-26.14	-28.98	-27.73	-18.43
Max (%)	22.33	25.76	27.62	24.5	24.61	26.90	25.99	26.27	28.03	37.44	22.39
Panel D:	*Unexpected illiquidity (daily-based)*										
Mean (%)	1.11	0.85	0.92	0.85	0.91	0.99	1.06	1.08	1.25	1.59	0.48
Std. Dev. (%)	5.72	5.49	5.67	5.83	5.93	6.00	6.11	6.24	6.52	7.19	3.81
Sharpe	0.19	0.16	0.16	0.15	0.15	0.17	0.17	0.17	0.19	0.22	0.12
Skewness	0.05	-0.02	0.06	0.01	0.21	0.09	0.09	0.14	0.21	0.83	-0.03
Kurtosis	5.27	5.66	6.13	6.61	5.61	6.01	5.88	5.77	6.63	7.94	15.19
Min (%)	-22.86	-22.94	-23.32	-26.27	-23.98	-24.57	-24.93	-26.45	-28.14	-25.76	-26.56
Max (%)	24.26	25.64	27.67	27.74	24.81	25.12	25.07	25.89	29.33	37.00	21.60

Table 2: This table reports monthly summary statistics for illiquidity-sorted decile portfolios over April 1998 to December 2025. Within each panel, decile portfolios are formed by sorting stocks each month on the corresponding illiquidity signal ($\overline{\mu}_\tau$ or $\overline{\epsilon}_\tau$); decile 1 (10) denotes the lowest (highest) signal. The spread portfolio is constructed as the high-minus-low return (10–1). Panel A (B) corresponds to *expected illiquidity* constructed from the range (daily) measure, whereas Panel C (D) corresponds to *unexpected illiquidity* constructed from the range (daily) measure. All statistics are computed from monthly returns. Mean, standard deviation, minimum, and maximum are reported in percent (%) per month.

These descriptive patterns motivate moving from unconditional moments to benchmark-adjusted comovement. The next step therefore examines how the illiquidity-sorted decile portfolios and the corresponding 10–1 spreads load on common risk factors by estimating time-series regressions on the Carhart four-factor model (Carhart, 1997). This analysis clarifies whether the return gradients observed in Table 2 are associated with systematic exposure to market, size, value, and momentum risk, and it provides a disciplined way to compare expected versus unexpected sorts and range- versus daily-based constructions in a unified factor-pricing framework.

Table 3 reports the Carhart estimates for each decile portfolio and for the high-minus-low spread, separately for expected and unexpected sorts and for the two underlying proxies. Across all panels, the decile portfolios exhibit strong market comovement: the market beta is typically close to one in every decile, indicating that each leg behaves like a broad equity portfolio with substantial exposure to aggregate market risk. In contrast, the 10–1 spreads are constructed to be largely market-neutral, and their market loadings are correspondingly close to zero, consistent with the long and short legs carrying similar market exposure.

For the *expected illiquidity* sorts, the factor structure of the 10–1 spreads is remarkably stable across proxies. The dominant non-market exposure is to the size factor: the expected spreads load strongly and positively on SMB, with coefficients close to one, indicating that the long–short trade behaves like a pronounced "small-minus-big" tilt in addition to its illiquidity sorting. The expected spreads also display a positive loading on the value factor HML of around one-third, suggesting that the high-illiquidity leg is systematically more value-oriented than the low-illiquidity leg. Finally, the expected spreads load negatively on momentum MOM, indicating exposure to a reversal factor (i.e., returns move opposite to standard momentum). Importantly, these patterns appear in both the range-based and daily-based constructions, which supports the view that the expected-signal sorts are capturing a persistent cross-sectional component that aligns with well-known firm characteristics and their associated factor exposures.

For the *unexpected illiquidity* sorts, the decile portfolios again retain market betas close to one, but the factor profile of the 10–1 spreads changes materially relative to the expected case. The unexpected spreads display a small but statistically meaningful market loading (on the order of a few tenths), while their SMB and HML loadings are econom-

	Decile portfolios										Spread
	1	2	3	4	5	6	7	8	9	10	10–1
Panel A:	*Expected illiquidity (range-based)*										
Alpha	-0.02	0.10	0.10	0.19	0.19	0.26^{c}	0.39^{a}	0.41^{b}	0.59^{a}	1.88^{a}	1.90^{a}
	(-0.21)	(0.82)	(0.78)	(1.39)	(1.40)	(1.85)	(2.76)	(2.52)	(3.60)	(8.43)	(8.97)
MKT–RF	0.97^{a}	0.97^{a}	0.98^{a}	1.00^{a}	0.98^{a}	1.01^{a}	0.97^{a}	0.96^{a}	0.94^{a}	0.97^{a}	0.00
	(41.98)	(34.81)	(32.29)	(30.17)	(30.35)	(30.22)	(28.60)	(24.65)	(23.97)	(18.22)	(-0.09)
SMB	0.03	0.05	0.19^{a}	0.35^{a}	0.48^{a}	0.55^{a}	0.68^{a}	0.75^{a}	0.83^{a}	1.01^{a}	0.99^{a}
	(0.84)	(1.24)	(4.72)	(7.89)	(11.03)	(12.18)	(14.89)	(14.27)	(15.75)	(14.20)	(14.56)
HML	0.01	0.19^{a}	0.28^{a}	0.31^{a}	0.33^{a}	0.36^{a}	0.40^{a}	0.49^{a}	0.40^{a}	0.34^{a}	0.32^{a}
	(0.44)	(5.49)	(7.35)	(7.49)	(8.24)	(8.69)	(9.28)	(10.05)	(8.04)	(5.07)	(5.13)
MOM	0.00	-0.07^{a}	-0.13^{a}	-0.14^{a}	-0.13^{a}	-0.17^{a}	-0.15^{a}	-0.20^{a}	-0.29^{a}	-0.42^{a}	-0.41^{a}
	(-0.18)	(-2.95)	(-4.69)	(-4.74)	(-4.49)	(-5.72)	(-5.06)	(-5.79)	(-8.31)	(-8.75)	(-9.13)
Panel B:	*Expected illiquidity (daily-based)*										
Alpha	0.02	0.09	0.11	0.18	0.15	0.27^{c}	0.44^{a}	0.36^{b}	0.62^{a}	1.84^{a}	1.82^{a}
	(0.22)	(0.79)	(0.89)	(1.32)	(1.06)	(1.92)	(2.97)	(2.31)	(3.70)	(8.28)	(8.56)
MKT–RF	0.96^{a}	0.97^{a}	0.97^{a}	1.02^{a}	0.99^{a}	1.01^{a}	1.00^{a}	0.95^{a}	0.95^{a}	0.96^{a}	0.00
	(41.45)	(34.92)	(31.77)	(31.19)	(29.98)	(29.94)	(28.30)	(25.89)	(23.65)	(18.07)	(-0.05)
SMB	0.02	0.05	0.19^{a}	0.34^{a}	0.48^{a}	0.54^{a}	0.66^{a}	0.80^{a}	0.84^{a}	0.99^{a}	0.97^{a}
	(0.71)	(1.24)	(4.67)	(7.85)	(10.89)	(12.03)	(13.92)	(16.15)	(15.59)	(13.90)	(14.22)
HML	0.00	0.19^{a}	0.28^{a}	0.31^{a}	0.33^{a}	0.36^{a}	0.38^{a}	0.49^{a}	0.43^{a}	0.35^{a}	0.35^{a}
	(0.07)	(5.52)	(7.25)	(7.54)	(8.02)	(8.59)	(8.61)	(10.55)	(8.59)	(5.27)	(5.48)
MOM	-0.01	-0.07^{a}	-0.13^{a}	-0.14^{a}	-0.13^{a}	-0.17^{a}	-0.17^{a}	-0.20^{a}	-0.26^{a}	-0.42^{a}	-0.41^{a}
	(-0.68)	(-2.95)	(-4.74)	(-4.88)	(-4.25)	(-5.62)	(-5.50)	(-6.16)	(-7.23)	(-8.87)	(-8.97)
Panel C:	*Unexpected illiquidity (range-based)*										
Alpha	0.26^{b}	0.19	0.22^{c}	0.26^{c}	0.40^{a}	0.30^{b}	0.46^{a}	0.44^{a}	0.63^{a}	0.96^{a}	0.70^{a}
	(1.99)	(1.37)	(1.70)	(1.88)	(3.07)	(2.19)	(3.38)	(3.13)	(3.95)	(5.06)	(3.88)
MKT–RF	0.89^{a}	0.90^{a}	0.93^{a}	0.98^{a}	0.93^{a}	0.99^{a}	1.00^{a}	1.02^{a}	1.04^{a}	1.05^{a}	0.16^{a}
	(29.00)	(27.88)	(30.37)	(30.11)	(30.32)	(30.52)	(30.96)	(30.57)	(27.38)	(23.33)	(3.68)
SMB	0.64^{a}	0.48^{a}	0.42^{a}	0.41^{a}	0.44^{a}	0.41^{a}	0.50^{a}	0.47^{a}	0.55^{a}	0.62^{a}	-0.02
	(15.47)	(10.95)	(10.15)	(9.39)	(10.67)	(9.30)	(11.35)	(10.35)	(10.69)	(10.27)	(-0.33)
HML	0.30^{a}	0.29^{a}	0.30^{a}	0.30^{a}	0.24^{a}	0.29^{a}	0.31^{a}	0.32^{a}	0.38^{a}	0.38^{a}	0.08
	(7.80)	(7.24)	(7.80)	(7.20)	(6.32)	(7.09)	(7.56)	(7.71)	(7.99)	(6.77)	(1.51)
MOM	-0.06^{b}	-0.11^{a}	-0.15^{a}	-0.16^{a}	-0.20^{a}	-0.19^{a}	-0.18^{a}	-0.21^{a}	-0.18^{a}	-0.29^{a}	-0.23^{a}
	(-2.22)	(-3.69)	(-5.30)	(-5.42)	(-7.13)	(-6.57)	(-6.08)	(-7.01)	(-5.14)	(-7.22)	(-5.99)
Panel D:	*Unexpected illiquidity (daily-based)*										
Alpha	0.45^{a}	0.24^{c}	0.28^{b}	0.20	0.26^{c}	0.35^{a}	0.38^{a}	0.42^{a}	0.56^{a}	0.94^{a}	0.48^{a}
	(3.25)	(1.81)	(2.21)	(1.45)	(1.95)	(2.73)	(2.68)	(2.96)	(3.64)	(4.94)	(2.65)
MKT–RF	0.90^{a}	0.90^{a}	0.94^{a}	0.95^{a}	0.97^{a}	0.98^{a}	1.01^{a}	1.00^{a}	1.05^{a}	1.04^{a}	0.14^{a}
	(26.93)	(28.80)	(31.66)	(28.44)	(30.48)	(32.03)	(29.91)	(29.58)	(28.40)	(22.91)	(3.22)
SMB	0.66^{a}	0.48^{a}	0.45^{a}	0.47^{a}	0.46^{a}	0.45^{a}	0.43^{a}	0.49^{a}	0.46^{a}	0.57^{a}	-0.10
	(14.84)	(11.39)	(11.30)	(10.48)	(10.68)	(10.80)	(9.48)	(10.73)	(9.24)	(9.32)	(-1.67)
HML	0.29^{a}	0.29^{a}	0.32^{a}	0.31^{a}	0.30^{a}	0.29^{a}	0.31^{a}	0.35^{a}	0.34^{a}	0.33^{a}	0.04
	(6.95)	(7.30)	(8.46)	(7.40)	(7.50)	(7.46)	(7.18)	(8.18)	(7.36)	(5.89)	(0.80)
MOM	-0.02	-0.10^{a}	-0.13^{a}	-0.13^{a}	-0.17^{a}	-0.21^{a}	-0.16^{a}	-0.21^{a}	-0.23^{a}	-0.35^{a}	-0.33^{a}
	(-0.79)	(-3.73)	(-4.96)	(-4.31)	(-5.81)	(-7.55)	(-5.38)	(-6.78)	(-6.88)	(-8.61)	(-8.35)

Table 3: This table reports time-series regressions of monthly decile-portfolio excess returns on the Carhart four-factor model (MKT–RF, SMB, HML, and MOM) over April 1998 to December 2025 Carhart (1997). Within each panel, decile portfolios are formed by sorting stocks each month on the corresponding illiquidity signal ($\overline{\mu}_\tau$ or $\overline{\epsilon}_\tau$); decile 1 (10) denotes the lowest (highest) signal. The spread portfolio is constructed as the high-minus-low return (10–1). Panel A (B) corresponds to *expected illiquidity* constructed from the range (daily) measure, whereas Panel C (D) corresponds to *unexpected illiquidity* constructed from the range (daily) measure. Coefficients are reported with t-statistics in parentheses. Intercepts are expressed in percent per month. Statistical significance is denoted by superscripts a, b, and c, corresponding to $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

ically small and typically not distinguishable from zero. The most robust non-market exposure remains the negative loading on MOM, again consistent with a reversal factor component. Relative to the expected spreads, these results indicate that sorting on innovations produces long–short portfolios whose systematic comovement is less tightly linked to the size and value dimensions and is instead more concentrated in the reversal direction, with the range-based and daily-based versions delivering qualitatively similar conclusions. An additional message is already visible within the Carhart specification itself: the estimated intercepts (the so-called Jensen's alpha; see Jensen (1968)) display a clear gradient for the *expected illiquidity* sorts. Alphas are generally small and statistically weak in the most liquid deciles and become larger and more often significant as one moves toward higher *expected illiquidity*, with the 10–1 spread among the strongest abnormal returns. For the *unexpected illiquidity* sorts, the cross-decile pattern is less sharp, but the spread alpha remains positive and statistically meaningful, consistent with the earlier evidence that these long–short portfolios are less tied to the size and value dimensions and load more cleanly on the reversal factor. Taken together, the long–short *factors* deliver positive alphas that are statistically significant at the 1% level under the Carhart specification across both proxies and both signal definitions, indicating that the spread not only removes market exposure but also concentrates the residual return component relative to the four-factor benchmark. Because alpha is benchmark-dependent (it captures the portion of average returns not explained by the included factors) I next compare decile and spread alphas under three standard specifications: the CAPM (Sharpe, 1964), the Fama–French three-factor model (Fama and French, 1993), and the Carhart four-factor model (Carhart, 1997).

Table 4 reports the estimated intercepts (alphas) from time-series regressions of the illiquidity-sorted decile portfolios and the corresponding 10–1 spreads on three benchmark models: the CAPM (Sharpe, 1964), the Fama–French three-factor model (Fama and French, 1993), and the Carhart four-factor model (Carhart, 1997). The Carhart row is reproduced here for ease of visual comparison with the previous table, where the four-factor regressions were discussed together with the associated factor loadings. The focus in this table is on how benchmark choice affects the estimated abnormal-return component across deciles and, in particular, for the tradable long–short factors.

For the *expected illiquidity* sorts (Panels A–B), the estimated intercepts line up with the

	Decile portfolios										Spread
	1	2	3	4	5	6	7	8	9	10	10–1
Panel A:	*Expected illiquidity (range-based)*										
CAPM alpha	-0.02	0.08	0.06	0.14	0.14	0.19	0.33^{c}	0.33	0.45^{c}	1.65^{a}	1.67^{a}
	(-0.23)	(0.63)	(0.39)	(0.85)	(0.81)	(1.01)	(1.66)	(1.44)	(1.88)	(5.44)	(5.72)
FF3 alpha	-0.02	0.05	0.03	0.11	0.11	0.16	0.30^{b}	0.30^{c}	0.42^{b}	1.63^{a}	1.66^{a}
	(-0.23)	(0.53)	(0.30)	(0.83)	(0.80)	(1.10)	(2.06)	(1.61)	(2.38)	(6.70)	(6.99)
Carhart alpha	-0.02	0.10	0.10	0.19	0.19	0.26^{c}	0.39^{a}	0.41^{b}	0.59^{a}	1.88^{a}	1.90^{a}
	(-0.21)	(0.82)	(0.78)	(1.39)	(1.40)	(1.85)	(2.76)	(2.52)	(3.60)	(8.43)	(8.97)
Panel B:	*Expected illiquidity (daily-based)*										
CAPM alpha	0.01	0.07	0.07	0.13	0.10	0.20	0.36^{c}	0.27	0.50^{b}	1.61^{a}	1.60^{a}
	(0.13)	(0.60)	(0.48)	(0.77)	(0.56)	(1.08)	(1.81)	(1.19)	(2.07)	(5.33)	(5.48)
FF3 alpha	0.01	0.05	0.04	0.10	0.07	0.17	0.34^{c}	0.24	0.47^{b}	1.60^{a}	1.59^{a}
	(0.13)	(0.43)	(0.30)	(0.70)	(0.52)	(1.18)	(2.21)	(1.48)	(2.63)	(6.51)	(6.74)
Carhart alpha	0.02	0.09	0.11	0.18	0.15	0.27^{c}	0.44^{a}	0.36^{b}	0.62^{a}	1.84^{a}	1.82^{a}
	(0.22)	(0.79)	(0.89)	(1.32)	(1.06)	(1.92)	(2.97)	(2.31)	(3.70)	(8.28)	(8.56)
Panel C:	*Unexpected illiquidity (range-based)*										
CAPM alpha	0.23	0.14	0.16	0.19	0.30^{c}	0.21	0.38^{b}	0.35^{c}	0.56^{a}	0.82^{a}	0.59^{a}
	(1.33)	(0.86)	(0.99)	(1.14)	(1.85)	(1.27)	(2.19)	(1.93)	(2.78)	(3.43)	(3.09)
FF3 alpha	0.22^{c}	0.15	0.18	0.21	0.32^{b}	0.21	0.41^{a}	0.36^{b}	0.60^{a}	0.86^{a}	0.64^{a}
	(1.72)	(1.10)	(1.33)	(1.53)	(2.37)	(1.55)	(2.89)	(2.40)	(3.53)	(4.12)	(3.57)
Carhart alpha	0.26^{b}	0.19	0.22^{c}	0.26^{c}	0.40^{a}	0.30^{b}	0.46^{a}	0.44^{a}	0.63^{a}	0.96^{a}	0.70^{a}
	(1.99)	(1.37)	(1.70)	(1.88)	(3.07)	(2.19)	(3.38)	(3.13)	(3.95)	(5.06)	(3.88)
Panel D:	*Unexpected illiquidity (daily-based)*										
CAPM alpha	0.45^{b}	0.19	0.22	0.15	0.19	0.26	0.31^{c}	0.33^{c}	0.46^{b}	0.77^{a}	0.32
	(2.43)	(1.20)	(1.40)	(0.87)	(1.12)	(1.53)	(1.80)	(1.81)	(2.39)	(3.21)	(1.59)
FF3 alpha	0.44^{a}	0.18	0.20	0.13	0.16	0.22^{c}	0.25^{b}	0.25^{b}	0.33^{a}	0.81^{a}	0.37^{b}
	(3.18)	(1.34)	(1.55)	(0.94)	(1.20)	(1.66)	(1.90)	(1.85)	(2.50)	(4.09)	(2.32)
Carhart alpha	0.45^{a}	0.24^{c}	0.28^{b}	0.20	0.26^{c}	0.35^{a}	0.38^{a}	0.42^{a}	0.56^{a}	0.94^{a}	0.48^{a}
	(3.25)	(1.81)	(2.21)	(1.45)	(1.95)	(2.73)	(2.68)	(2.96)	(3.64)	(4.94)	(2.65)

Table 4: This table reports the intercepts (alphas) from time-series regressions of monthly decile-portfolio excess returns on the CAPM, the Fama–French three-factor model (FF3), and the Carhart four-factor model (FF4; MKT–RF, SMB, HML, and MOM) over April 1998 to December 2025 (Carhart, 1997). Within each panel, decile portfolios are formed by sorting stocks each month on the corresponding illiquidity signal ($\overline{\mu}_{\tau}$ or $\overline{\epsilon}_{\tau}$); decile 1 (10) denotes the lowest (highest) signal. The spread portfolio is constructed as the high-minus-low return (10–1). Panel A (B) corresponds to *expected illiquidity* constructed from the range (daily) measure, whereas Panel C (D) corresponds to *unexpected illiquidity* constructed from the range (daily) measure. Coefficients are reported with t-statistics in parentheses. Intercepts are expressed in percent per month. Statistical significance is denoted by superscripts a, b, and c, corresponding to $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

sorting dimension: moving from the most liquid to the most illiquid portfolios, alphas tend to shift upward and gain statistical support, and this qualitative ordering is preserved when the benchmark is expanded from CAPM to FF3 and then to Carhart. Economically, the effect is concentrated in the high-illiquidity tail: for example, in Panel A the decile 10 alpha rises from about 1.65% per month under CAPM to about 1.88% under Carhart (both highly significant), while the corresponding 10–1 spread remains close to 1.6–1.9% per month and is statistically significant at the 1% level in all three benchmarks. Panel B delivers the same message for the daily-based proxy, with the 10–1 alpha around 1.6% under CAPM/FF3 and about 1.82% under Carhart, again significant at the 1% level across models. These patterns indicate that the long–short construction concentrates a residual return component that is not explained by standard equity factors.

For the *unexpected illiquidity* sorts (Panels C–D), the cross-decile pattern is less regular, consistent with the more unstable ranking induced by innovation-based signals. The long–short spreads remain positive and are generally statistically meaningful, with particularly strong evidence under the range-based construction: in Panel C the 10–1 alpha is about 0.59% under CAPM, about 0.64% under FF3, and about 0.70% under Carhart, and it is significant at the 1% level in each case. A more benchmark-sensitive behavior appears for the daily-based unexpected factor in Panel D. Here the 10–1 alpha is positive but not statistically significant under CAPM (about 0.32% per month), becomes statistically significant at the 5% level under FF3 (about 0.37% per month), and is statistically significant at the 1% level under Carhart (about 0.48% per month). This sensitivity to benchmark choice aligns with the earlier loading evidence: once momentum (which empirically acts in a reversal direction for these portfolios) is included among the controls, the remaining intercept changes both in magnitude and precision, suggesting that an economically relevant share of comovement for the daily-based unexpected spread is tied to the momentum/reversal dimension.

Taken together, the portfolio evidence indicates that both the *expected-* and *unexpected-illiquidity* long–short strategies generate positive average spreads and, in several cases, non-trivial benchmark-adjusted intercepts. These time-series regressions, however, are portfolio-level diagnostics: they summarize abnormal performance of the constructed factor returns relative to a chosen benchmark, but they do not directly answer the asset-pricing question of whether *exposure* to illiquidity risk is rewarded in the cross-section of individual stocks. To move from performance attribution to risk-premium estimation, I therefore implement cross-sectional pricing tests using the two-step Fama–MacBeth methodology (Fama and MacBeth, 1973). The econometric setup and all implementation details for the two-step procedure are described in Section 3.5, including the first-step time-series regressions used to estimate stock-level factor sensitivities and the second-step cross-sectional regressions used to estimate the price of illiquidity risk. I now turn to the empirical Fama–MacBeth results and test whether the estimated illiquidity beta commands a positive price of risk on average, and whether pricing differs between *expected* versus *unexpected illiquidity* and between range- versus daily-based constructions.

Table 5 reports the time-series averages of the monthly estimated risk premia (the $\overline{\gamma}$ coefficients) for the illiquidity factor under alternative control specifications, obtained from Fama–MacBeth cross-sectional regressions of monthly stock excess returns.

A first observation is that the intercept terms are generally positive and statistically significant at conventional levels across specifications. For instance, under the Carhart four-factor specification, the intercept is around 0.73% per month for *expected illiquidity* (both range- and daily-based, $t = 4.4$) and remains strongly significant also for the unexpected constructions (about 0.72–0.73%, t between 4.4 and 4.5). This indicates that, on average, a non-negligible component of cross-sectional expected returns is not absorbed by the included benchmark factors[21].

The central parameter of interest, however, is the price of illiquidity risk (row "Illiq."). Here the evidence is markedly weaker, in the baseline specification (which includes only the illiquidity beta), the expected range-based factor delivers a $\overline{\gamma}^{\text{Illiq}}$ of about 0.61% with a t-statistic of 1.68, corresponding to statistical significance at the 10% level. The daily-based expected factor yields a similar magnitude (0.59%, $t = 1.63$), again only marginally significant. Once standard risk controls are introduced, however, statistical significance disappears: under CAPM, FF3, and FF4, the estimated $\overline{\gamma}^{\text{Illiq}}$ for *expected illiquidity* becomes statistically indistinguishable from zero. For the range-based expected factor, $\overline{\gamma}^{\text{Illiq}}$ falls to roughly 0.35–0.43% and is no longer significant (e.g., $t = 1.11$ under FF3 and FF4), and the daily-based expected factor delivers the same qualitative conclusion of an insignificant

[21]The large and growing "factor zoo" suggests that standard benchmark models may leave systematic return variation unexplained. See Harvey et al. (2016) for a discussion of the proliferation of proposed factors and the resulting data-mining concerns, and Hou et al. (2017) for evidence that many published anomalies weaken under stricter replication standards. Related evidence on which characteristics remain robust once many candidates are included is provided by Green et al. (2017).

	Expected illiquidity (range-based)				*Expected illiquidity (daily-based)*				*Unexpected illiquidity (range-based)*				*Unexpected illiquidity (daily-based)*			
	Baseline	CAPM	FF3	FF4	Baseline	CAPM	FF3	FF4	Baseline	CAPM	FF3	FF4	Baseline	CAPM	FF3	FF4
Intercept	0.68^{a}	0.46^{b}	0.73^{a}	0.73^{a}	0.69^{a}	0.45^{b}	0.74^{a}	0.73^{a}	0.95^{a}	0.48^{b}	0.69^{a}	0.73^{a}	0.91^{a}	0.48^{b}	0.72^{a}	0.72^{a}
	(3.21)	(2.16)	(4.41)	(4.45)	(3.27)	(2.16)	(4.40)	(4.44)	(4.18)	(2.49)	(3.98)	(4.49)	(3.72)	(2.50)	(4.37)	(4.48)
Illiq.	0.61^{c}	0.35	0.43	0.42	0.59	0.31	0.41	0.41	0.11	-0.18	0.02	0.00	0.18	-0.02	0.25	0.23
	(1.68)	(0.93)	(1.11)	(1.11)	(1.63)	(0.83)	(1.07)	(1.07)	(0.44)	(-0.60)	(0.08)	(0.02)	(0.63)	(-0.05)	(0.81)	(0.78)
MKT–RF		0.48	0.32	0.32		0.49	0.32	0.32		0.62^{c}	0.40	0.33		0.56^{c}	0.33	0.32
		(1.36)	(0.98)	(1.02)		(1.38)	(0.98)	(1.02)		(1.78)	(1.20)	(1.05)		(1.66)	(1.03)	(1.02)
SMB			0.25	0.25			0.25	0.25			0.29	0.29			0.28	0.28
			(1.13)	(1.14)			(1.14)	(1.15)			(1.34)	(1.36)			(1.28)	(1.29)
HML			-0.52^{b}	-0.50^{b}			-0.52^{b}	-0.50^{b}			-0.48^{b}	-0.47^{b}			-0.54^{b}	-0.51^{b}
			(-2.13)	(-2.08)			(-2.14)	(-2.09)			(-2.03)	(-1.98)			(-2.24)	(-2.19)
MOM				0.02				0.02				-0.08				-0.02
				(0.05)				(0.05)				(-0.20)				(-0.06)

Table 5: Fama–MacBeth cross-sectional estimates under alternative control sets. The table reports the time-series averages of the monthly estimated risk premia (i.e., the $\overline{\gamma}$ coefficients) from Fama–MacBeth cross-sectional regressions (Fama and MacBeth, 1973) of monthly stock excess returns for 1,893 U.S. equities over April 1998 to December 2025. Columns report specifications for *expected* and *unexpected illiquidity* factors constructed from the range-based and daily-based measures. Within each block, "Baseline" includes only the illiquidity factor; "CAPM" adds MKT–RF; "FF3" adds SMB and HML; "FF4" further adds MOM. Coefficients are expressed in percent per month (i.e., multiplied by 100), with t-statistics in parentheses. Statistical significance is denoted by superscripts a, b, and c, corresponding to $p < 0.01$, $p < 0.05$, and $p < 0.10$, respectively.

price of risk once controls are included. For the *unexpected illiquidity* factors, the evidence is even weaker: the baseline $\overline{\gamma}^{\text{Illiq}}$ is small and statistically insignificant (e.g., 0.11%, t = 0.44 for range-based; 0.18%, t = 0.63 for daily-based), and remains close to zero once controls are added. The only clear sign reversal occurs for the range-based unexpected factor under CAPM (about −0.18%), but it is again far from statistically significant (t = −0.60).

Overall, the evidence shows that the illiquidity sorted long–short portfolios earn positive and economically large alphas in time-series regressions, the corresponding stock-level illiquidity betas do not command a robust cross-sectional risk premium once standard factors are accounted for. In other words, the abnormal performance documented in the portfolio-based analysis does not translate into a statistically reliable price of risk in the Fama–MacBeth framework.

Additional coefficients help rationalize this outcome. The HML price of risk is consistently negative and statistically significant (around −0.50%, $t \approx -2$), indicating a cross-sectional premium tilted toward growth rather than value stocks in this sample. The market and size premia are generally positive but not statistically significant in most specifications. These patterns suggest that part of the return dispersion across stocks is captured by traditional dimensions, leaving limited residual variation to be priced by illiquidity exposure.

4.4 Interpretation of the Pricing Evidence

Several considerations can rationalize why the portfolio-level patterns do not translate into a stable cross-sectional price of illiquidity risk in the Fama–MacBeth stage. A first issue is data provenance. The sample is built from publicly available sources rather than CRSP/Compustat, which limits the ability to apply the standard screens and corporate-event treatments that are routine in the asset-pricing literature, and it also constrains the set of characteristics that can be used to refine the cross-section (e.g., historical market capitalization and book equity for systematic value-weighting or characteristic-controlled portfolio formation). These limitations matter because the illiquidity effect is documented to be heterogeneous across firms and to vary with market-wide conditions, so estimates of an average price of illiquidity risk can be sensitive to cross-sectional composition. In particular, Amihud (2002) emphasizes that increases in *expected market illiquidity* can simultaneously raise required returns for stocks (a broad discount-rate effect) and induce substitution toward more liquid stocks ("flight to liquidity"), implying that the net return impact can differ sharply between liquid and illiquid (often smaller) firms and across stress episodes[22]. The present setting is consistent with this concern because data availability may reduce coverage precisely in the segment of the cross-section where liquidity effects are often strongest.

A second issue concerns portfolio construction under data constraints. Because historical market capitalization and book equity are not reliably available for the full cross-section in the publicly sourced dataset, the decile portfolios and long–short factors are formed using equal weights. This choice is transparent and comparable over long horizons, but it can alter factor exposures and the magnitude and precision of benchmark-adjusted performance. In closely related contexts, Plyakha et al. (2012) emphasize that equal-weighted and value-weighted implementations are not innocuous alternatives: equal-weighting mechanically embeds frequent rebalancing and shifts portfolio emphasis toward smaller firms, which can change comovement with standard risk factors and affect estimated intercepts even when the underlying sorting signal is unchanged[23]. In practice, the inability to implement value-weighted sorts or characteristic-based reweighting may therefore attenuate the precision and comparability of the Fama–MacBeth pricing tests.

A third issue is selection and survivorship. Reliance on public data makes it difficult to incorporate delistings and corporate events systematically, and the asset-pricing literature shows that missing delisting returns are often large and negative and that delistings are concentrated among the smallest stocks, so omitting them can distort the return distribution precisely where illiquidity is most pronounced[24]. The direction of the bias is not

[22] Amihud (2002) describes two forces when *expected market illiquidity* rises: (i) a general effect that depresses stock prices and raises expected returns, and (ii) substitution toward more liquid stocks ("flight to liquidity"), which mitigates price declines for liquid stocks and can weaken the average relation between illiquidity and expected returns for that segment of the cross-section. This mechanism implies that empirical evidence on an illiquidity premium is sensitive to how well the sample represents the smallest and least liquid firms and to whether it captures episodes in which the substitution channel is particularly strong. The broader asset-pricing perspective that liquidity compensation can be state-dependent is consistent with models where expected returns incorporate liquidity risk and its interaction with market conditions (Acharya and Pedersen, 2005).

[23] The key point for the present application is not that equal-weighting is "worse", but that it changes the economic object being tested: relative to value-weighting, it increases exposure to small-cap dynamics and introduces a systematic rebalancing component that can affect factor attribution and the sampling variability of abnormal returns. This makes cross-sectional pricing estimates less directly comparable to studies that can replicate standard market- or value-weighted constructions.

[24] See Shumway and Warther (1999) on delisting bias, the magnitude of missing delisting returns, and the implications for cross-sectional tests that load heavily on small and distressed firms.

mechanically one-sided for long–short portfolios: excluding delisted firms can remove extreme negative outcomes from the long leg, but it can also remove episodes in which the short leg would have benefited from defaults or delistings. Either way, incomplete coverage of exits can weaken inference by reducing variation in distress states and by disconnecting the measured illiquidity signal from the full economic support of realized returns.

Lastly, measurement choices can reduce power even when liquidity is economically important. Different proxies capture different dimensions of liquidity, and not all dimensions must be priced; if a proxy only imperfectly tracks price impact, the resulting illiquidity betas and their estimated premia become noisy. This concern is consistent with the liquidity-risk perspective of Pástor and Stambaugh (2003) and is reinforced by the replication and measurement-sensitivity evidence emphasized in Pástor and Stambaugh (2019), which highlights that alternative liquidity series can deliver both false positives and false negatives depending on how trading frictions are proxied and aggregated.

CHAPTER 5

CONCLUSIONS

This thesis investigates whether illiquidity is systematically priced in the cross-section of U.S. equities, with an explicit focus on constructing *tradable* illiquidity factors derived from model-based signals. The empirical design combines (i) alternative illiquidity proxies, (ii) a dynamic modeling step based on a Multiplicative Error Model (MEM), and (iii) standard asset-pricing tests (portfolio sorts, time-series factor regressions, and Fama–MacBeth cross-sectional regressions). A central feature of the framework is the decomposition of illiquidity into an *expected* component (the conditional mean) and an *unexpected* component (the innovation), which are then mapped into investable long–short strategies.

The first conclusion is a measurement one. Illiquidity proxies that exploit intraday information are empirically superior. Range-based measures display cleaner time-series behaviour and more credible persistence than close-to-close daily proxies, consistent with lower measurement noise. The comparison with realized measures is crucial in this respect: when a realized benchmark is available, it behaves as the "clean" reference signal, and the range-based proxy is much closer to it than the daily-based proxy. This validation step matters because the subsequent modeling and pricing tests inherit the quality of the underlying illiquidity signal.

The second conclusion is about dynamics and interpretation. Modeling illiquidity with a MEM is not only appropriate for a strictly positive series; it also provides a clean economic split between a slow-moving component and a shock component. Across specifications, the conditional mean behaves like an illiquidity "state" that evolves gradually, while the innovation captures short-run liquidity disturbances. Importantly, when the underlying proxy is improved (moving from close-to-close to range-based, and closer to realized benchmarks when available), the estimated dynamics become less mechanically persistent and more informative about genuine innovations.

The asset-pricing evidence is consistent with this distinction. *Expected illiquidity* spreads load most on size, reversal (negative momentum), and value, while *unexpected illiquidity* spreads load mainly on reversal (negative momentum) with only weak market exposure. In the cross-sectional Fama–MacBeth tests, however, the results are weaker once standard benchmarks are included. In practice, the portfolio spread that appears in the sorts does not turn into a stable and clearly significant "price of illiquidity risk" in the cross-section. After controlling for common factors, it becomes difficult to separate what is truly liquidity-specific from what is already captured by the usual equity dimensions, and statistical significance tends to fade.

A simple way to read this evidence is that the two MEM components matter in different ways. The *expected* component is very persistent and looks like a slow-moving firm condition. Because of this, it overlaps with exposures that standard factors and related characteristics already capture, so in controlled specifications it adds limited extra pricing information. The *unexpected* component instead captures short-run liquidity *shocks*, meaning a quick worsening in trading conditions, which is closer to the economic idea of liquidity risk; empirically, however, the shock-based signal delivers smaller and less stable spreads, and its cross-sectional price of risk is not robust once standard benchmarks are included.

Overall, the main insight is that measurement drives inference: intraday-informed proxies produce the cleanest signals, and illiquidity sorted portfolios can generate spreads, but the cross-sectional evidence does not support a robust and distinct price of illiquidity risk once standard factor structure is taken into account.

BIBLIOGRAPHY

Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. *Journal of financial Economics*, 77(2):375–410.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625.

Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 64(2):253–280.

Barndorff-Nielsen, O. E. and Shephard, N. (2003). Realized power variation and stochastic volatility models. *Bernoulli*, 9(2):243–265.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327.

Brownlees, C. T., Cipollini, F., and Gallo, G. M. (2011). Multiplicative error models. *Available at SSRN 1852285*.

Brownlees, C. T. and Gallo, G. M. (2010). Comparison of volatility measures: a risk management perspective. *Journal of Financial Econometrics*, 8(1):29–56.

Caporin, M., Rossi, E., and Santucci de Magistris, P. (2017). Chasing volatility: A persistent multiplicative error model with jumps. *Journal of Econometrics*, 198(1):122–145.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of finance*, 52(1):57–82.

Chordia, T., Roll, R., and Subrahmanyam, A. (2000). Commonality in liquidity. *Journal of financial economics*, 56(1):3–28.

Cipollini, F., Engle, R. F., and Gallo, G. M. (2009). A model for multivariate non-negative valued processes in financial econometrics. *Available at SSRN 1333869*.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2):174–196.

Engle, R. (2002). New frontiers for arch models. *Journal of applied econometrics*, 17(5):425–446.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007.

Engle, R. F. and Bollerslev, T. (1986). Modelling the persistence of conditional variances. *Econometric reviews*, 5(1):1–50.

Engle, R. F. and Gallo, G. M. (2006). A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics*, 131(1–2):3–27.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of financial economics*, 33(1):3–56.

Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of political economy*, 81(3):607–636.

Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics*, 14(1):71–100.

Goyenko, R. Y., Holden, C. W., and Trzcinka, C. A. (2009). Do liquidity measures measure liquidity? *Journal of financial Economics*, 92(2):153–181.

Green, J., Hand, J. R., and Zhang, X. F. (2017). The characteristics that provide independent information about average us monthly stock returns. *The Review of Financial Studies*, 30(12):4389–4436.

Hansen, P. R. and Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business & Economic Statistics*, 24(2):127–161.

Harvey, C. R., Liu, Y., and Zhu, H. (2016). … and the cross-section of expected returns. *The Review of financial studies*, 29(1):5–68.

Hou, K., Xue, C., and Zhang, L. (2017). Replicating anomalies. Technical report, National Bureau of Economic Research.

Hwang, S. and Valls Pereira, P. L. (2006). Small sample properties of garch estimates and persistence. *The European Journal of Finance*, 12(6-7):473–494.

Jensen, M. C. (1968). The performance of mutual funds in the period 1945-1964. *The Journal of finance*, 23(2):389–416.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335.

Lacava, D., Ranaldo, A., and Santucci de Magistris, P. (2023). Realized illiquidity. *Swiss Finance Institute Research Paper*, (22-90).

Liang, S. X. and Wei, J. K. (2012). Liquidity risk and stock returns around the world. *Journal of Banking & Finance*, 36(12):3274–3288.

Madhavan, A. (2000). Market microstructure: A survey. *Journal of financial markets*, 3(3):205–258.

Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business*, 53(1):61–65.

Pástor, L. and Stambaugh, R. F. (2003). Liquidity risk and expected stock returns. *Journal of Political economy*, 111(3):642–685.

Pástor, L. and Stambaugh, R. F. (2019). Liquidity risk after 20 years. *Critical Finance Review*, 8(1-2):277–299.

Plyakha, Y., Uppal, R., and Vilkov, G. (2012). Why does an equal-weighted portfolio outperform value-and price-weighted portfolios? *Available at SSRN 2724535*.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance*, 19(3):425–442.

Shumway, T. and Warther, V. A. (1999). The delisting bias in crsp's nasdaq data and its implications for the size effect. *The Journal of Finance*, 54(6):2361–2379.

Tauchen, G. E. and Pitts, M. (1983). The price variability-volume relationship on speculative markets. *Econometrica: Journal of the Econometric Society*, pages 485–505.

FIGURES AND TABLES

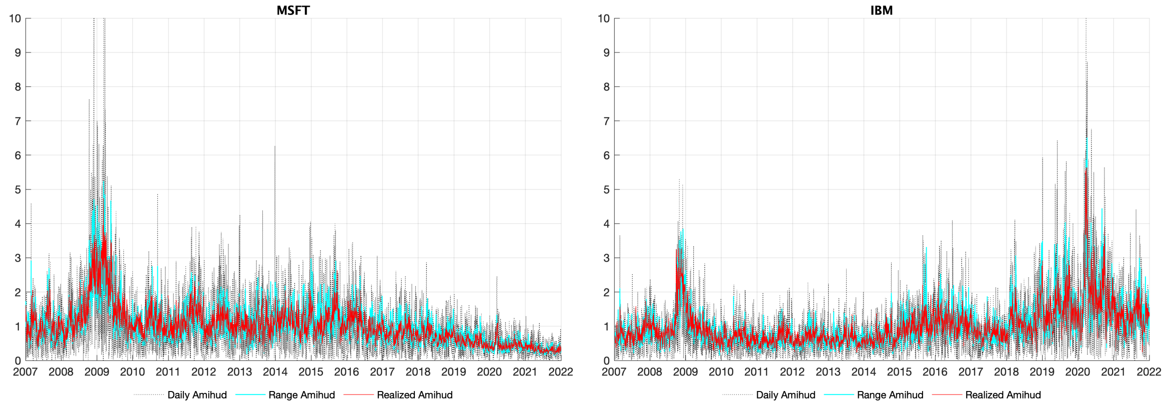

Figure 4: Comparison between the classic daily Amihud illiquidity measure (black dotted line), the range-based Amihud (cyan solid line), and the realized Amihud (red solid line) for MSFT (left panel) and IBM (right panel). The sample period spans from January 3, 2007 to December 31, 2021. For each proxy $\mathcal{A}_t^{(j)}$, with $j \in \{\text{daily}, \text{range}, \text{realized}\}$, the normalized series is defined as $\tilde{\mathcal{A}}_t^{(j)} = \mathcal{A}_t^{(j)} / \left(\frac{1}{T}\sum_{t=1}^{T} \mathcal{A}_t^{(j)}\right)$, where T denotes the available sample size for the given stock. This normalization preserves the relative time-series dynamics while rescaling each proxy to have unit mean, facilitating comparability across measures constructed on different scales.

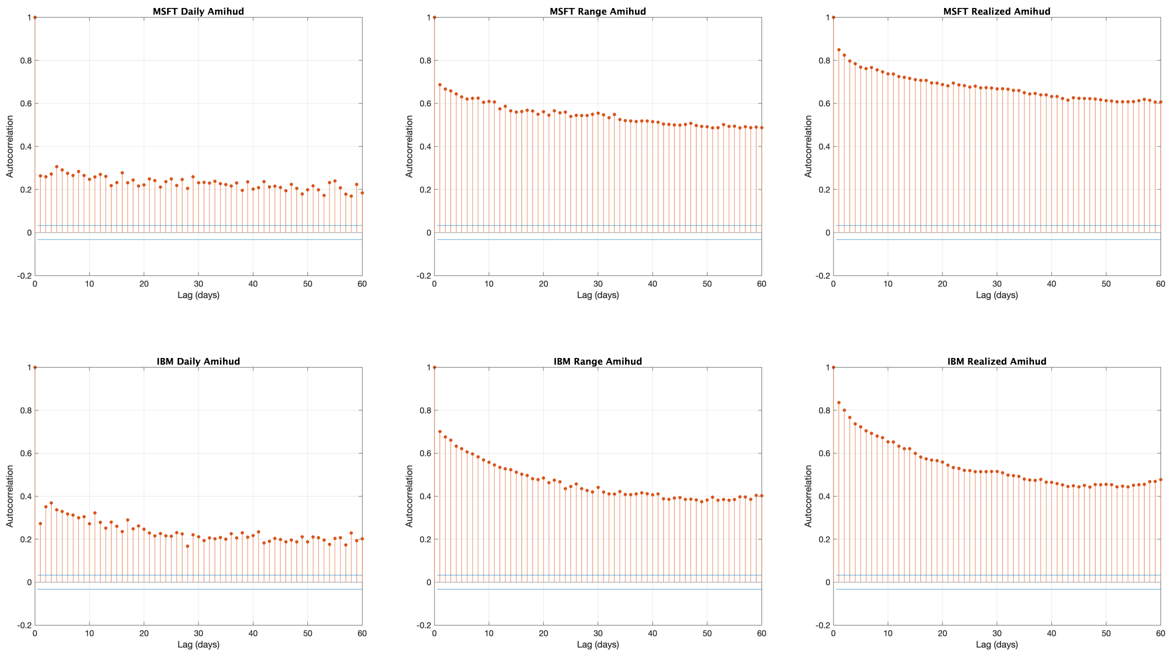

Figure 5: Empirical autocorrelation functions (ACF) of illiquidity measures for MSFT (top row) and IBM (bottom row). In each row, the left panel reports the classic daily Amihud proxy, the middle panel the range-based Amihud, and the right panel the realized Amihud. The sample period spans from January 3, 2007 to December 31, 2021. The horizontal axis reports lags in trading days (up to 60). The blue bands denote approximate 95% confidence bounds under the white-noise null.

Table 6: Descriptive statistics of illiquidity measures across frequencies

	Daily Amihud ($\mathcal{A}_t^D$)			Range Amihud ($\mathcal{A}_t^{HL}$)			Realized Amihud ($\mathcal{A}_t$)		
Frequency	Daily	Weekly	Monthly	Daily	Weekly	Monthly	Daily	Weekly	Monthly
AAPL									
Mean	0.0031	0.0031	0.0031	0.0026	0.0026	0.0026	0.0194	0.0194	0.0194
Median	0.0023	0.0026	0.0025	0.0021	0.0022	0.0022	0.0161	0.0162	0.0164
Std. Dev.	0.0032	0.0022	0.0019	0.0019	0.0017	0.0016	0.0132	0.0127	0.0123
IQR	0.0030	0.0019	0.0016	0.0016	0.0013	0.0012	0.0102	0.0092	0.0092
P95	0.0088	0.0077	0.0077	0.0061	0.0061	0.0064	0.0466	0.0475	0.0494
Max	0.0302	0.0179	0.0118	0.0187	0.0121	0.0098	0.1198	0.0930	0.0800
CV	1.0177	0.6930	0.6202	0.7119	0.6445	0.6133	0.6821	0.6532	0.6338
MSFT									
Mean	0.0075	0.0075	0.0075	0.0067	0.0068	0.0068	0.0514	0.0514	0.0514
Median	0.0057	0.0068	0.0068	0.0060	0.0064	0.0065	0.0479	0.0486	0.0495
Std. Dev.	0.0071	0.0045	0.0039	0.0039	0.0034	0.0031	0.0276	0.0256	0.0244
IQR	0.0077	0.0046	0.0039	0.0045	0.0036	0.0030	0.0307	0.0275	0.0248
P95	0.0200	0.0147	0.0132	0.0139	0.0127	0.0110	0.1010	0.0920	0.0879
Max	0.0868	0.0441	0.0274	0.0354	0.0243	0.0199	0.2437	0.1734	0.1607
CV	0.9481	0.6085	0.5234	0.5799	0.4974	0.4639	0.5364	0.4968	0.4741
IBM									
Mean	0.0187	0.0187	0.0187	0.0168	0.0168	0.0168	0.1265	0.1264	0.1264
Median	0.0145	0.0159	0.0160	0.0140	0.0144	0.0146	0.1085	0.1093	0.1097
Std. Dev.	0.0171	0.0117	0.0099	0.0101	0.0087	0.0080	0.0697	0.0638	0.0593
IQR	0.0183	0.0115	0.0102	0.0108	0.0094	0.0089	0.0723	0.0673	0.0669
P95	0.0502	0.0421	0.0378	0.0351	0.0333	0.0320	0.2596	0.2497	0.2504
Max	0.2348	0.1208	0.0848	0.1095	0.0761	0.0630	0.7135	0.5785	0.4908
CV	0.9150	0.6232	0.5315	0.6012	0.5167	0.4733	0.5508	0.5049	0.4692

Notes: The table reports descriptive statistics of three illiquidity measures Amihud, range-based Amihud, and realized Amihud computed for AAPL, MSFT, and IBM at daily, weekly, and monthly frequencies over the sample period January 3, 2007 to December 31, 2021. Weekly and monthly series are obtained by averaging daily observations over non-overlapping 5-day and 22-day windows, respectively. To improve readability, all reported values are rescaled by a factor 10^9. The coefficient of variation (CV) is defined as the ratio between the standard deviation and the mean.