








Modelling Financial Markets

Chapter 10: Market Microstructure



Daniel Traian PELE

Bucharest University of Economic Studies
IDA Institute Digital Assets
Blockchain Research Center
AI4EFin Artificial Intelligence for Energy Finance
Romanian Academy, Institute for Economic Forecasting
MSCA Digital Finance

What Does Trading Cost?

> *"This paper presents a method for inferring the effective bid-ask spread directly from a time series of market prices. The method requires no data other than the prices themselves; so it is very cheap."*
>
> — *Roll (1984)*

- One tick (USD 0.01) on SPY (SPDR S&P 500 ETF Trust, an exchange-traded fund on the S&P 500) is 0.21 bp: the floor of the quoted spread
- From daily bars, the SPY spread estimates lie two orders of magnitude above one tick
  - Roll 64.3 bp, EDGE (Efficient Discrete Generalized Estimator of the spread) 40.1 bp: they are dominated by volatility
- From 5-minute bars, EDGE gives 0.46 bp, the closest to one tick
- Corwin–Schultz puts the BVB (median 54 bp) next to the US (median 53 bp)
  - median daily traded value: USD 0.34 million on the BVB, USD 10 505 million in the US

▸ Central chart: the SPY spread at two frequencies

What We Will Explain: from Orders to Prices and Costs



Nasdaq MarketSite, Times Square, New York: a studio, not a trading floor

Photo: Luca Marfè (2015); CC BY 2.0; Wikimedia Commons

What you should be able to do

- compute the Roll, Corwin–Schultz, Abdi–Ranaldo and EDGE spread estimates
- compute the Amihud ratio and interpret its units
- derive the Glosten–Milgrom spread, the Kyle equilibrium and the Almgren–Chriss trajectory
- interpret information shares and the latency arbitrage tax
- Further Topics and Appendix: reading for familiarity

- **Questions**: when computers match orders in microseconds
  - what does it cost to trade?
  - how do prices absorb information?
- The rules of the matching engine (priority, tick size, auctions) shape prices and costs Gould et al. (2013)
- Models: order book and spread (Roll, Corwin–Schultz, EDGE)
  - Amihud ratio, Glosten–Milgrom, Kyle, Almgren–Chriss
- Applications: stress episodes, crypto, the Bucharest Stock Exchange (BVB), the speed race, intraday momentum
- Worked exercises: Seminar 10
- Detailed examples and advanced topics are in Further Topics, before the Appendix
  - the footer button leads there, and the Back button returns to the slide you came from

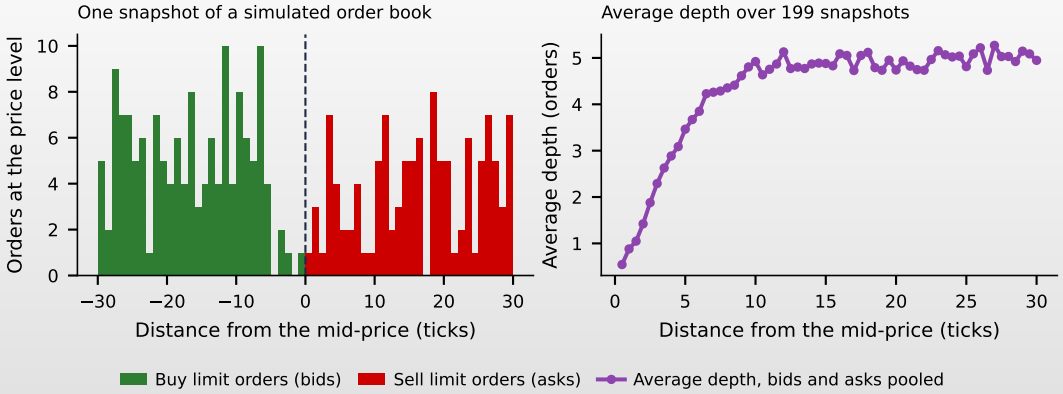
Liquidity and the Limit Order Book: Definitions

- **Question**: what does a market order buy, and why does its cost depend on the shape of the book?
- **Liquidity**: trading a quantity quickly at a price close to fair value
  - no single number captures it
  - **Tightness** (the spread), **depth** (size needed to move the price), **resiliency** (speed of recovery) Kyle (1985); **immediacy** Harris (2003)
- **Limit order book** (LOB): all standing limit orders, by price and time of arrival Gould et al. (2013)
- At time t: best bid $b_t <$ best ask a_t
  - mid-price $m_t = (a_t + b_t)/2$
  - relative spread $(a_t - b_t)/m_t$ in basis points (1 bp $= 0.01\%$)
  - **Depth**: quantity waiting at a price level; **tick size**: smallest price step (USD 0.01 for most US stocks above USD 1)

Orders, Priority and Matching

- **Market order**: immediate execution against the available liquidity, uncertain price
  - **limit order**: price bounded by the limit, uncertain execution
  - other order types: Appendix
- Resting limit orders supply liquidity
  - market orders and marketable limit orders consume it
- **Price-time priority**: the better price first, then the earlier order at the same price
  - orders at the same price form a queue
  - queue position is valuable (case study: the high-frequency trading arms race)
- **Call auctions** at the open and close: orders are pooled
  - they are matched at the single price that maximises the executed quantity
- A market buy of size Q walks up the ask levels
  - it pays the half-spread plus the cost of the deeper levels

A Simulated Order Book (1/2)



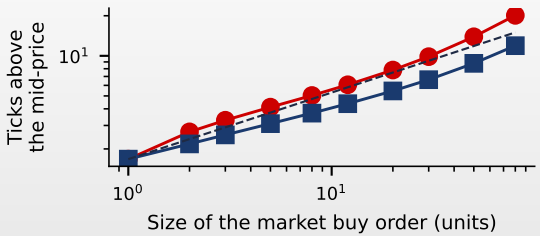
MFM_ch10_order_book ▶ Interactive

A Simulated Order Book (2/2)

- **Simulated**: zero-intelligence order flow Farmer, Patelli & Zovko (2005)
  - limit orders, market orders and cancellations arriving at random
  - rates: Appendix
  - Left: one snapshot (bids green, asks red); right: mean depth by distance from the mid
- Mean spread 3.23 ticks, median 3, above 6 ticks 10% of the time
  - depth 0.8 orders at 1 tick, 4.8 at 10, 5.0 at 25 (199 snapshots)
- Market orders eat the best levels first, so liquidity near the price is scarce Tóth et al. (2011)
  - real books show the same rise away from the mid Bouchaud, Bonart, Donier & Gould (2018)

MFM_ch10_order_book

The Price Impact of a Market Order



- Market buys of 1 to 80 units sent against each simulated snapshot; log-log axes
- Marginal price: a 1-unit order reaches 1.7 ticks above the mid, an 80-unit order 20.1 ticks
- **Concave** impact: slope 0.53, close to 1/2
- Average execution price: slope 0.44; 80 units pay 11.9 ticks per unit on average
- Doubling the order: cost per unit $\times 2^{\mathbf{0.44}} \approx 1.36$, not 2
  - total cost (size $\times$ cost per unit) $\times 2^{\mathbf{1+0.44}} \approx 2.7$
- The impact is concave because book depth grows with distance from the mid
- Large orders split over time (metaorders) show a similar law
  - section *Price Impact and Optimal Execution*

MFM_ch10_order_book ▶ Interactive

Why Is There a Spread?

- **Question**: why is there a spread, and how do we measure it with and without quotes?
- A liquidity provider buys at the bid and sells at the ask
- The spread pays three costs Huang & Stoll (1997)
  - **Order processing**: fees, technology, clearing
  - **Inventory**: holding an unwanted position exposes the provider to price risk
  - **Adverse selection**: some counterparties know more Bagehot (1971)
- The provider incurs losses when trading with informed traders
  - it must recover them from noise traders (section *Information and Prices*)

Quoted, Effective and Realised Spread (1/2)

- Three measures of what a round trip costs, all relative to the mid-price m_t just before the trade
  - p_t: trade price
  - $q_t = +1$ if the buyer initiated the trade (it hit the ask), -1 if the seller did
- **Quoted** spread, from quotes only: $S_t^Q = (a_t - b_t)/m_t \geq 0$
- **Effective** spread, what the trade paid: $S_t^E = 2q_t(p_t - m_t)/m_t$
  - q_t makes the cost positive on both sides
  - the one-way cost is $S_t^E/2$
- **Realised** spread: $S_t^R = 2q_t(p_t - m_{t+\Delta})/m_t$
  - what the liquidity provider keeps before fees and inventory costs
  - $m_{t+\Delta}$: the mid a time Δ after the trade (e.g. 5 minutes)
  - $S_t^R < 0$ when $q_t(m_{t+\Delta} - p_t) > 0$: the later mid-price passes the trade price, in the trade's direction
  - a smaller continuation lowers S_t^R but can leave it positive

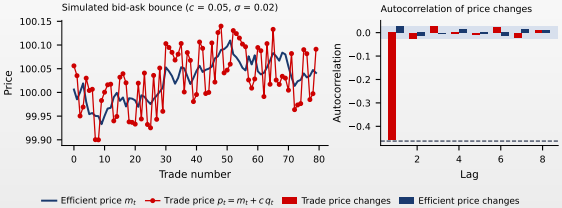
Quoted, Effective and Realised Spread (2/2)

- **Price impact**: $S_t^E - S_t^R = 2q_t(m_{t+\Delta} - m_t)/m_t$, the lasting move of the mid
  - usually attributed to adverse selection
  - inventory and order-flow pressure also contribute
- Effective and realised spreads need trades, the prevailing quotes and the trade direction
  - direction signed for example by the rule of Lee & Ready (1991)
  - above the mid a buy, below it a sell
  - at the mid the tick test (against the last different trade price)
- Without such data the spread is **estimated** from prices only: the next slides

The Roll Model

- Idea Roll (1984): the spread can be read from the way trade prices bounce between bid and ask
  - $t = 1, \dots, n$: successive trades, prices in USD (or log prices)
- **Efficient price**, a random walk that moves only on news: $m_t = m_{t-1} + \varepsilon_t$
  - ε_t i.i.d. (independent, identically distributed), mean 0, variance σ^2
- **Trade price**: $p_t = m_t + c\, q_t$
  - $c > 0$: the half-spread ($s = 2c$)
  - $q_t = +1$ (buy at the ask) or -1 (sell at the bid), each with probability 1/2
  - signs i.i.d. and independent of ε
- Price change: $\Delta p_t = \varepsilon_t + c(q_t - q_{t-1})$, so $\mathrm{Var}(\Delta p_t) = \sigma^2 + 2c^2$ and $\mathrm{Cov}(\Delta p_t, \Delta p_{t-1}) = -c^2$
  - **Bid-ask bounce**: the common term q_{t-1} enters Δp_{t-1} with $+c$ and Δp_t with $-c$, so a jump up to the ask tends to be followed by a fall to the bid, even when the efficient price is a random walk
- **Roll estimator**: $\hat{s} = 2\hat{c} = 2\sqrt{-\widehat{\mathrm{Cov}}(\Delta p_t, \Delta p_{t-1})}$
  - $\widehat{\mathrm{Cov}}$: the sample first-order autocovariance of the $n-1$ price changes
  - defined only if it is negative

The Bid-Ask Bounce (Simulated)



- Left: the first 80 trades, efficient price m_t and trade price p_t
  - right: autocorrelations of the price changes up to lag 8, for both prices
- $n = 5\,000$ simulated trades with $c = 0.05$, $\sigma = 0.02$
  - light-blue band $\pm 1.96/\sqrt{n}$: the 95% band for the autocorrelations of white noise
  - dashed line: theoretical first-order autocorrelation
- True spread $2c = 0.10$, Roll estimate $\hat{s} = 0.0998$
  - lag-1 autocorrelation -0.458 against the theoretical Cov/Var $= -c^2/(\sigma^2 + 2c^2) = -0.463$
- Trade prices, lags 2–8: inside the pointwise 95% band; lag 3 (0.028) sits right at its edge (± 0.028)
- No significant autocorrelation beyond lag 1: compatible with weak-form efficiency at these lags, not a proof of it
  - trade prices are still not a random walk

MFM_ch10_roll ▶ Interactive

The Roll Estimator in Finite Samples

- On real data the sample covariance is often **positive**, so $\hat{s}$ does not exist
  - on SPY 5-minute bars, on 41% of days
  - **Sampling error**: the model is true, but the sample is short and c^2 is small relative to σ^2
  - **Misspecification**: trends after news, orders split into many trades, stale prices, bars instead of trades
- Notation Harris (1990): $\gamma_k = \mathrm{Cov}(\Delta p_t, \Delta p_{t-k})$
  - $\gamma_0 = \sigma^2 + 2c^2$ and $\gamma_1 = -c^2$ in the Roll model
  - T: the number of price changes (77 per SPY day)
  - $\hat{\gamma}_1$: the sample value, from demeaned price changes
- **Bias**: $\mathbb{E}[\hat{\gamma}_1] \approx \gamma_1 - (\gamma_0 + 2\gamma_1)/T$, negative even with no spread ($c = 0$)
  - Roll then finds a spread that is not there
- **Noise**: $T\,\mathrm{Var}(\hat{\gamma}_1) \to \gamma_0^2 + \gamma_1^2$ as $T \to \infty$ (independent signs $q_t = \pm 1$), so the variance shrinks only as $1/T$
  - Bartlett's Gaussian approximation, $\gamma_0^2 + 3\gamma_1^2$, omits the fourth cumulant of binary signs: $-2c^4 = -2\gamma_1^2$
  - the limit requires a finite fourth moment, which Chapter 1 questions
- **Wrong sign**: $\hat{\gamma}_1$ is approximately Normal, so $P(\hat{\gamma}_1 > 0) \approx \Phi(\sqrt{T}\,\gamma_1/\gamma_0)$
  - Φ: distribution function of the standard Normal distribution
  - close to 1/2 whenever $c^2 \ll \sigma^2/\sqrt{T}$

Spreads from Daily Ranges: Corwin–Schultz

- Idea Corwin & Schultz (2012): the daily high is usually a buy at the ask, the low a sell at the bid
  - so the high-low range contains volatility plus the spread
- The volatility part grows with the length of the interval, the spread part does not
  - compare one day with two days
- H_t, L_t: high and low of day t
  - $H_{t,t+1}$, $L_{t,t+1}$: high and low over days t and $t+1$
  - $\ln(H/L)$: the log range
- $\beta = \sum_{j=0}^{1} \left[\ln \frac{H_{t+j}}{L_{t+j}}\right]^2$, the sum of the two squared one-day ranges
  - $\gamma = \left[\ln \frac{H_{t,t+1}}{L_{t,t+1}}\right]^2$, the squared two-day range
  - Without a spread, $\gamma \approx \beta$ on average (variance adds up over two days); the spread inflates β twice, γ only once
- Solving the two moment equations: $\alpha = \frac{\sqrt{2\beta} - \sqrt{\beta}}{3 - 2\sqrt{2}} - \sqrt{\frac{\gamma}{3 - 2\sqrt{2}}}, \qquad \hat{S} = \frac{2(e^{\alpha} - 1)}{1 + e^{\alpha}}$
  - α: the spread part of the log range; $\hat{S}$: relative spread, a fraction of the price ($\times 10^4$ for bp), with $\hat{S} \approx \alpha$ for small α

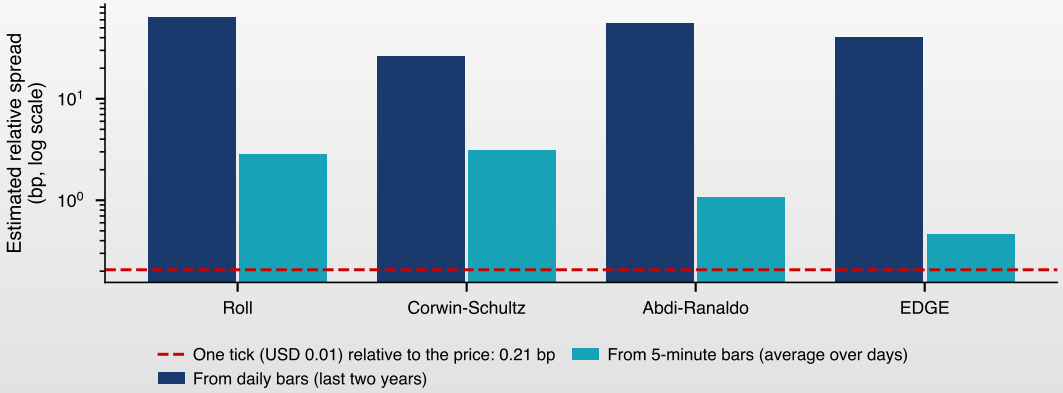
Abdi–Ranaldo and EDGE (1/2)

- **Abdi–Ranaldo** Abdi & Ranaldo (2017): the close sits at the bid or the ask
  - the middle of the day's range proxies the mid-price
  - c_t: log closing price of day t; $\eta_t = (\ln H_t + \ln L_t)/2$: log mid-range of day t
- The product of the close's distances from today's and tomorrow's mid-range
  - is, on average, a quarter of the squared spread: $\hat{S}^2 = 4\,\mathbb{E}\left[(c_t - \eta_t)(c_t - \eta_{t+1})\right]$
  - News moves η_{t+1} independently of the bounce at c_t, so the volatility terms average out
  - $\mathbb{E}$ is replaced by the sample mean over days; a negative mean product gives $\hat{S} = 0$

Abdi–Ranaldo and EDGE (2/2)

- **EDGE** (Efficient Discrete Generalized Estimator) Ardia, Guidotti & Kroencke (2024): uses open, high, low and close (OHLC) prices
  - Two moment conditions, each a product of log-returns between the open, the close and η_t (exact form: Appendix), give two unbiased estimates $\hat{e}_1$, $\hat{e}_2$ of S^2, with sampling variances v_1, v_2
- Minimum-variance combination (diagonal weighting): $\hat{S}^2 = \dfrac{v_2 \hat{e}_1 + v_1 \hat{e}_2}{v_1 + v_2}$
  - Each estimate is weighted by the **other**'s variance, so the noisier one gets the smaller weight
- Common to both: only OHLC bars are needed
  - a negative mean moment gives $\hat{S} = 0$
  - $\sqrt{\hat{S}^2}$ is not exactly unbiased for S
- For crypto-assets, high-low estimators track the time variation of high-frequency liquidity
  - the Amihud ratio is better for levels Brauneis, Mestel, Riordan & Theissen (2021)

Spread Estimators at Two Frequencies: SPY



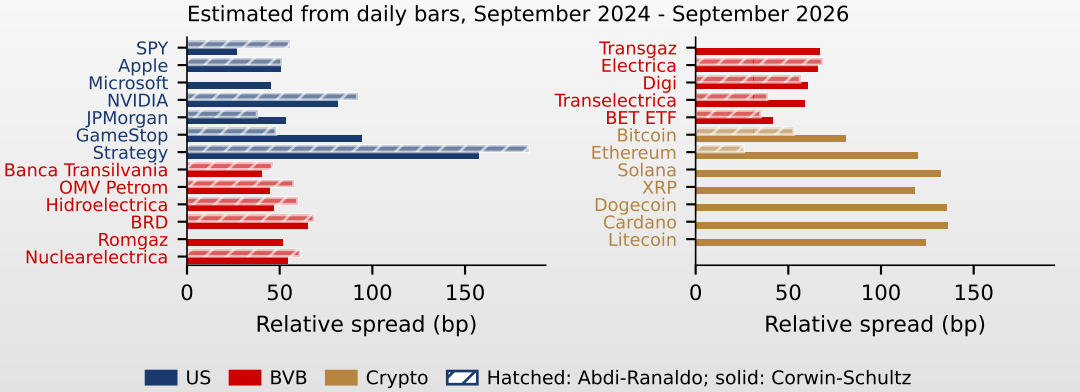
MFM_ch10_spread_estimators ▶ Interactive

SPY Spread Estimates: Daily versus Five-Minute Bars

- Chart: each estimator in bp, log scale
  - dark blue: from daily bars (Sep 2024 – Sep 2026)
  - teal: from 5-minute bars (1 477 days, 2020–2026)
  - red dashed line: one tick
- One tick (USD 0.01) is 0.21 bp: a floor for the **quoted** spread, not for the effective one
- From daily bars the estimates are two orders of magnitude above a tick, dominated by volatility Jahan-Parvar & Zikes (2023)
  - Roll 64.3, Corwin–Schultz 26.5, Abdi–Ranaldo 55.1, EDGE 40.1 bp
- From 5-minute bars EDGE is the closest to a tick
  - Roll 2.82, Corwin–Schultz 3.11, Abdi–Ranaldo 1.06, EDGE 0.46 bp (interval [0.41, 0.51])
- EDGE per day from its 5-minute bars, negative days set to 0, then averaged
  - interval for the mean of the estimator: percentile bootstrap over overlapping 20-day blocks
- A tick is a plausibility check, not a validation
- Type of evidence: comparison of procedures (it shows which method works better on the same data, not why)
  - Roll 2.82 bp: about 54% of the mean autocovariance is demeaning bias (Further Topics)
  - after the correction the Roll spread falls to 1.92 bp, about 32% lower

MFM_ch10_spread_estimators

Spread Estimates Across 25 Assets



Why Daily Spread Proxies Rank BVB Next to US Large Caps

- Chart: Corwin–Schultz (solid bar) and Abdi–Ranaldo (hatched bar), in bp, for each asset
  - colours: US, BVB, crypto
- Daily bars, Sep 2024 – Sep 2026; BVB stocks in lei
- Median Corwin–Schultz: US 53, BVB 54, crypto 124 bp
- BVB blue chips look as liquid as US large caps
  - implausible, given how little value they trade (next section)
  - Median daily traded value, same window and groups: USD 10 505 million in the US, USD 0.34 million on the BVB, a ratio of about 31 000
- The ranking follows volatility
  - Abdi–Ranaldo is 0 for 8 of the 25 assets (negative sample mean)
- Levels need an independent benchmark: the effective spread measured from trades Jahan-Parvar & Zikes (2023)
- Type of evidence: description (it shows what happened in the data, without testing a cause)

MFM_ch10_spread_estimators

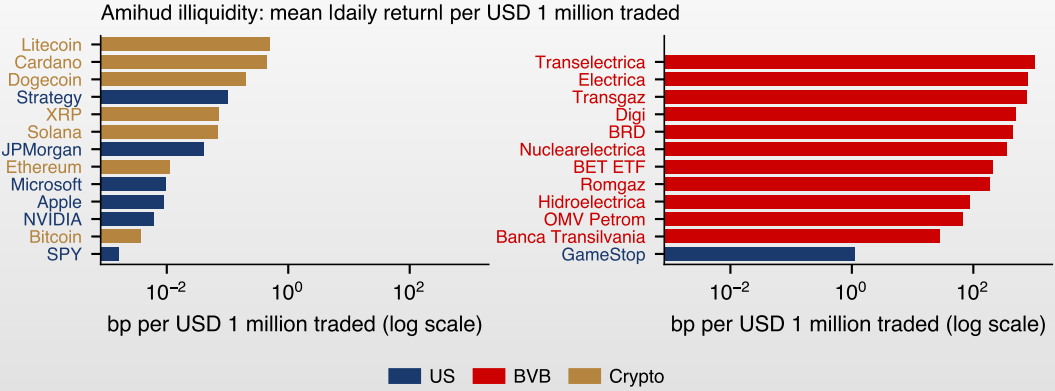
The Amihud Illiquidity Ratio

- **Question**: how do we measure market depth with only daily returns and volumes?
- Daily price impact per unit of money traded Amihud (2002): $\mathrm{ILLIQ}_i = \frac{K}{D_i} \sum_{d=1}^{D_i} \frac{|r_{i,d}|}{\mathrm{DVOL}_{i,d}}$
  - i: the asset; $d = 1, \ldots, D_i$: its trading days with nonzero volume; $r_{i,d}$: daily return; $\mathrm{DVOL}_{i,d}$: traded value (price $\times$ volume)
  - $K = 10^4 \times 10^6 = 10^{10}$, with r as a fraction and DVOL (Dollar VOLume, traded value, price $\times$ volume) in USD: basis points of price move per USD 1 million traded
- Larger means less liquid: the same traded value moves the price more
- Needs only daily returns and volumes, available for every market and period
  - BVB stocks: lei at the BNR (Banca Națională a României, the National Bank of Romania) reference rate of the day
  - crypto: traded value in USD summed across exchanges
- A low-frequency proxy of Kyle's lambda (section *Information and Prices*) Hasbrouck (2009)

Reading the Amihud Ratio: Units and Kyle's Lambda

- A toy day: $|r| = 50$ bp on USD 25 million traded gives $50/25 = 2$ bp per USD 1 million
  - i.e. 200 bp per USD 100 million
  - ILLIQ_i averages such daily ratios
- The intraday SPY measure of this section: bp per USD 100 million
  - the mean of the day's 78 bar ratios
  - Its mean 1.83 is 0.0183 bp per USD 1 million, not comparable with the daily SPY ratio (0.0016): a bar carries about 1/78 of the day's value, but its $|r|$ is not 78 times smaller
- Link to Kyle's λ, the price move per share of net order flow y
  - if all volume were net flow, $\Delta p = \lambda y$
  - at price p: $|r| = \lambda |y|/p$ and $\mathrm{DVOL} = p|y|$, so $|r|/\mathrm{DVOL} = \lambda/p^2$
  - Offsetting buys and sells add volume without moving the price, so the ratio is at most λ/p^2 (seminar A6)

Amihud Illiquidity Across 25 Assets



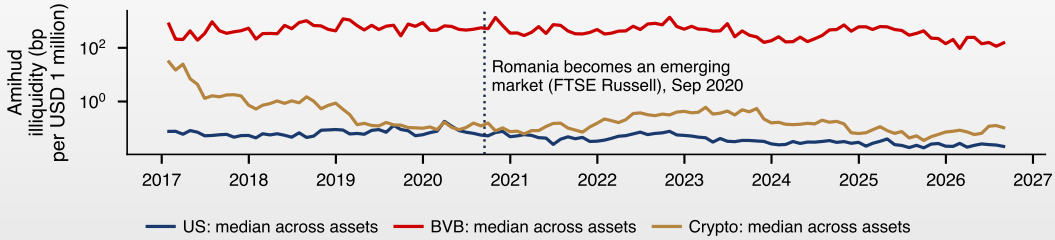
MFM_ch10_amihud ▶ Interactive

Amihud Ratios: Three Clusters Orders of Magnitude Apart

- Sep 2024 – Sep 2026, log scale; mean over about 500 trading days (730 for crypto)
- SPY 0.0016 bp per USD 1 million; Bitcoin 0.0037 bp
  - SPY median traded value USD 40 124 million a day
- Banca Transilvania, the most liquid BVB stock: 28.0 bp, about 17 600 times SPY
- Transelectrica 1 024.0 bp; the ETF on the BET (Bucharest Exchange Trading index, the BVB benchmark): 212.9 bp
- Three clusters, orders of magnitude apart
  - US large caps and top crypto
  - smaller crypto and GameStop
  - the BVB
- Daily spread proxies do not separate markets into these three clusters
- A proxy, not a cost you could actually trade at
- Type of evidence: description

MFM_ch10_amihud

Illiquidity over Time: Three Markets



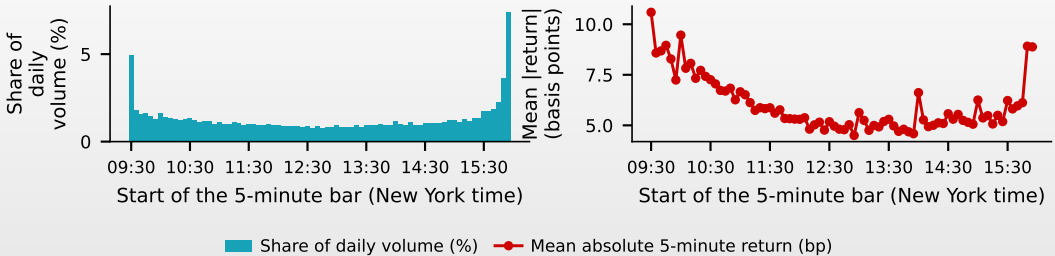
- Monthly Amihud ratio, median across each group (without SPY and the BET ETF), log scale
- Crypto: from 7.6 bp per USD 1 million (2017) to 0.089 (2026); US: from 0.062 to 0.023
- BVB: from 435 to 169, still about 7 000 times the US median
  - no visible break at the move to emerging-market status by the index provider FTSE Russell (FTSE: Financial Times Stock Exchange), September 2020 BVB (2026)

MFM_ch10_amihud ▶ Interactive

Is Liquidity Priced?

- **Level**: investors demand a higher return for assets that are costly to trade Amihud & Mendelson (1986)
  - Expected return increases with the spread, concavely: long-horizon investors hold the illiquid assets
- **Time series**: expected market illiquidity raises future returns Amihud (2002)
  - an unexpected rise lowers current prices
- **Liquidity risk**: stocks that fall when market liquidity dries up earn a premium Pástor & Stambaugh (2003), Acharya & Pedersen (2005)
- Evidence here: monthly market returns fall when illiquidity rises unexpectedly on the BVB, in the US and in crypto (seminar C2)
  - **Unexpected** illiquidity: the residual of an AR(1) (first-order autoregressive) model fitted to the market's monthly illiquidity, as in Amihud (2002); the monthly return is regressed on last month's illiquidity and this residual (Nov 2016 – Aug 2026)
  - t-statistics of the illiquidity-shock coefficient: -4.0 (BVB), -3.8 (US), -4.1 (crypto)
- Type of evidence: association (the variables move together; this does not prove causation)
  - the return and the illiquidity shock are measured in the same month

SPY Within the Trading Day



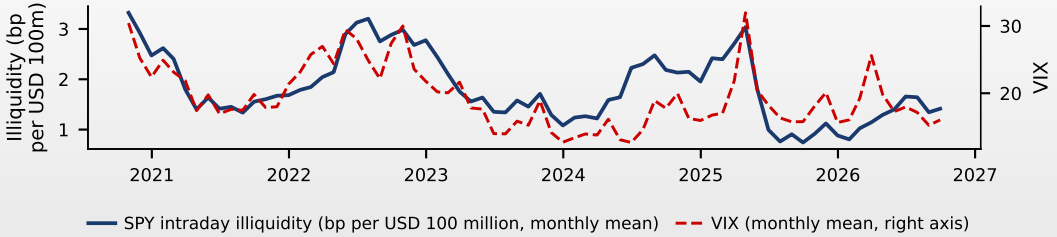
- 78 five-minute bars a day, 1 477 full days (Oct 2020 – Sep 2026)
  - the first bar starts at the opening price
- Volume: first bar 5.0% of the day, a midday bar 0.9%, against $1/78 = 1.28\%$ per bar if even
  - last 30 minutes 18.6% [18.3, 19.0] against $6/78 = 7.69\%$ if even
- Mean absolute 5-minute return $|r|$: 10.6 bp in the first bar, 5.1 bp at midday
  - ratio 2.06 [1.98, 2.15] (95% confidence interval, CI, resampling whole days, seminar B1)
  - the spikes at 10:00 and 14:00 coincide with scheduled macroeconomic data releases

MFM_ch10_intraday ▶ Interactive

Why a U-Shape?

- Returns and volume are highest at the open and the close Wood, McInish & Ord (1985)
  - documented for stocks of the New York Stock Exchange (NYSE)
- **Open**: overnight news is absorbed; uncertainty is highest
  - Informed traders act early, before the information spreads
- **Clustering**: noise traders who can choose their timing trade together, and liquidity attracts liquidity Admati & Pfleiderer (1988)
- **Close**: index funds and ETFs trade at the closing price, which is fixed by the closing auction
- Practical rule: an execution algorithm follows the volume curve to reduce impact
  - section *Price Impact and Optimal Execution*
- Type of evidence: description
  - the explanations on this slide are not tested here

Liquidity and Volatility: SPY and the VIX (1/2)



- Daily intraday illiquidity: mean over the 78 bars of $|r|$ / traded value
  - bp per USD 100 million (divide by 100 for USD 1 million)
  - monthly means shown
- Correlation with the VIX (CBOE Volatility IndeX): 0.58 daily, 0.69 for monthly means
  - elasticity of log illiquidity on log VIX 0.95
  - HAC standard error, robust to heteroskedasticity and autocorrelation: 0.12
  - fitted relation in seminar B4

MFM_ch10_intraday ▶ Interactive

Liquidity and Volatility: SPY and the VIX (2/2)

- April 2025: 3.04 against a sample mean of 1.83
  - liquidity dries up when investors most need it
  - so liquidity risk is correlated with market risk Acharya & Pedersen (2005)
- Type of evidence: association
  - the illiquidity ratio contains $|r|$, which rises with volatility

The Glosten–Milgrom Model

- **Question**: how does private information enter prices, and why does it alone create a spread?
- Value $V \in \{V_L, V_H\}$, $P(V = V_H) = \theta$; each period one trader arrives Glosten & Milgrom (1985)
  - With probability μ informed (buys if $V = V_H$, sells if $V = V_L$), otherwise a noise trader, uninformed and trading for liquidity needs (buys or sells with probability $1/2$) Bagehot (1971)
- Competitive, risk-neutral market maker with zero expected profit on each trade
  - $a = \mathbb{E}[V \mid \text{buy}]$ (ask), $b = \mathbb{E}[V \mid \text{sell}]$ (bid)
  - A buy signals a high value and a sell a low one, so $a > b$
- Informed traders always buy when $V = V_H$, noise traders half of the time
  - $P(\text{buy} \mid V_H) = \mu + \frac{1-\mu}{2}$, $P(\text{buy} \mid V_L) = \frac{1-\mu}{2}$
- After each trade, Bayes' rule updates θ: a buy raises it, a sell lowers it

Glosten–Milgrom: the Spread

- Bayes' rule after a buy, then the ask as the expected value
  - $P(V_H \mid \text{buy}) = \dfrac{\left(\mu + \frac{1-\mu}{2}\right)\theta}{\left(\mu + \frac{1-\mu}{2}\right)\theta + \frac{1-\mu}{2}(1-\theta)}, \qquad a = V_L + (V_H - V_L)\,P(V_H \mid \text{buy})$
  - Numerator: probability of "high value and a buy"; denominator: total probability of a buy; the bid follows symmetrically from a sell
- At $\theta = 1/2$: $a - b = \mu(V_H - V_L)$, linear in the share of informed traders
  - $V_L = 90$, $V_H = 110$: spread 2 for $\mu = 0.1$, 6 for $\mu = 0.3$
- No inventory, no fees: the spread exists purely because of information
- Trade prices form a martingale: the market is efficient in the semi-strong sense (Chapter 2)

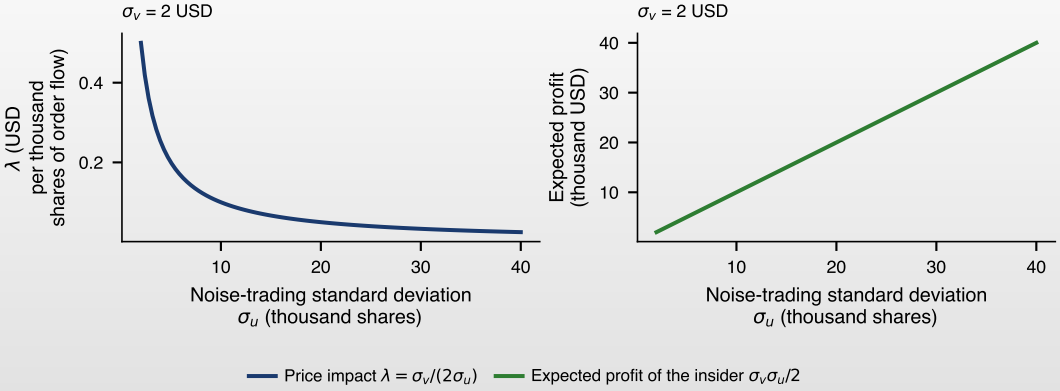
The Kyle Model

- One insider knows $v \sim N(p_0, \sigma_v^2)$, with p_0 the price before trading, and submits a market order x Kyle (1985)
  - Noise traders submit $u \sim N(0, \sigma_u^2)$, independent of v
  - The market maker sees only the total order flow $y = x + u$ and sets $p = \mathbb{E}[v \mid y]$
- Linear equilibrium: $x = \beta(v - p_0)$, $p = p_0 + \lambda y$
  - $\beta > 0$: the insider's trading intensity (not the Corwin–Schultz β)
  - $\lambda > 0$: **Kyle's lambda**, the price move per share of net order flow
- The **insider** takes the pricing rule as given: $\mathbb{E}[(v - p)\,x \mid v] = (v - p_0)x - \lambda x^2$
  - the gain from information minus the cost of moving the price against himself
  - First-order condition: $x = (v - p_0)/(2\lambda)$, so $\beta = 1/(2\lambda)$: he trades less because his order moves the price
- The **market maker**: (v, y) is jointly Normal, so $\lambda = \mathrm{Cov}(v, y)/\mathrm{Var}(y) = \beta\sigma_v^2/(\beta^2\sigma_v^2 + \sigma_u^2)$
  - the population OLS (ordinary least squares) slope of v on y

Kyle: the Equilibrium and What the Price Reveals

- Substituting $\beta = 1/(2\lambda)$ into the slope gives $\beta^2\sigma_v^2 = \sigma_u^2$, hence
  - $\beta = \sigma_u/\sigma_v, \quad \lambda = \sigma_v/(2\sigma_u), \quad \mathbb{E}[\text{insider profit}] = \sigma_v\sigma_u/2$
- More noise trading (σ_u up): a smaller λ (deeper market), a larger β and a larger profit
- More valuable information (σ_v up): a larger λ, the market maker protects himself
- $\mathrm{Var}(x) = \beta^2\sigma_v^2 = \sigma_u^2$: the insider's orders are as variable as noise trading, behind which he hides
- $\mathrm{Var}(v \mid y) = \sigma_v^2 - \lambda^2\mathrm{Var}(y) = \sigma_v^2/2$: the price reveals half of the private information

Kyle's Lambda and the Insider's Profit (1/2)



MFM_ch10_information ▶ Interactive

Kyle's Lambda and the Insider's Profit (2/2)

- $\sigma_v = 2$ USD; noise trading σ_u from 2 000 to 40 000 shares
- Chart units: λ in USD per thousand shares; $1/\lambda$ is **market depth**
- Reading the chart: λ falls like $1/\sigma_u$, the insider's profit rises linearly in σ_u

MFM_ch10_information

How Should We Interpret PIN?

- PIN (probability of informed trading), from the Easley, Kiefer, O'Hara & Paperman (1996) model of daily buy and sell counts: $\mathrm{PIN} = \alpha\mu/(\alpha\mu + \varepsilon_b + \varepsilon_s)$, $\mathrm{PIN} \in [0, 1]$
  - α probability of an information event on a day
  - μ arrival rate of informed orders
  - ε_b, ε_s arrival rates of uninformed buys and sells
- Example: $\alpha = 0.5$, $\mu = 2$, $\varepsilon_b = \varepsilon_s = 2$ give $\mathrm{PIN} = 1/(1 + 2 + 2) = 0.20$: one fifth of expected trades are informed
- PIN describes a model-implied share; it does not identify which particular trade was informed
- VPIN (Volume-synchronised Probability of INformed trading), a volume-time version: Further Topics

From Models to Measurement

Measure	What it identifies	Data needed
Kyle λ	price change per unit of signed order flow; the Amihud ratio is its daily proxy Hasbrouck (2009)	signed flow, or daily bars
MRR	Madhavan–Richardson–Roomans model: permanent (adverse-selection) and transitory parts of the spread Madhavan, Richardson & Roomans (1997)	signed trades and quotes
Hasbrouck VAR	vector autoregression (VAR): long-run quote response to a trade innovation Hasbrouck (1991)	signed trades and quotes
PIN	$\alpha\mu/(\alpha\mu + \varepsilon_b + \varepsilon_s)$, the expected share of informed trades Easley, Kiefer, O'Hara & Paperman (1996)	daily buy and sell counts
Order-flow imbalance	explains most short-horizon price changes Cont, Kukanov & Stoikov (2014)	order-book events

- The public bars used here do not sign trades, so MRR (Madhavan–Richardson–Roomans, structural model of trading costs), the Hasbrouck VAR (Vector AutoRegression) and PIN are not estimated here
- Their equations are in the Appendix: ▸ MRR ▸ PIN ▸ Hasbrouck VAR

Signed Order Flow: Empirical Lambda and OFI

- **Question**: how much does order flow move the price over short intervals, measured directly from data?
- Empirical Kyle lambda: regress the mid-price change on the net signed flow Hasbrouck (2007)
  - over intervals k (e.g. 5 minutes)
  - $\Delta m_k = \lambda \sum_{t \in k} q_t v_t + u_k$; v_t: size of trade t; q_t: its sign (Lee–Ready rule)
- OFI (order-flow imbalance) Cont, Kukanov & Stoikov (2014): built from order-book events n
  - each event changes the best prices P^b_n, P^a_n and the depths d^b_n, d^a_n at the best quotes
  - $e_n = d^b_n \, 1\{P^b_n \geq P^b_{n-\mathbf{1}}\} - d^b_{n-\mathbf{1}} \, 1\{P^b_n \leq P^b_{n-\mathbf{1}}\} - d^a_n \, 1\{P^a_n \leq P^a_{n-\mathbf{1}}\} + d^a_{n-\mathbf{1}} \, 1\{P^a_n \geq P^a_{n-\mathbf{1}}\}$, $\mathrm{OFI}_k = \sum_{n \in k} e_n$; $1\{\cdot\}$: 1 if the condition holds, else 0
  - a rising bid or a rising ask contributes positively; a falling bid or a falling ask contributes negatively
  - at an unchanged ask the quantity matters: a fall in d^a contributes positively
  - $\Delta m_k = \beta \, \mathrm{OFI}_k + \varepsilon_k$, with the slope β (not Kyle's) inversely proportional to the average depth at the best quotes: a linear relation, unlike the concave impact of trade volume
- OFI counts limit orders and cancellations, not only trades
  - so it explains short-horizon price changes better than the trade imbalance
  - it needs order-book events, which the bars used here do not contain

Metaorders and the Square-Root Law

- **Question**: how far does a large order move the price?
  - how should it be split over time to pay as little as possible?
- Signs of market orders are autocorrelated for hundreds of trades Bouchaud, Gefen, Potters & Wyart (2004)
  - large **metaorders** are split into many small trades over hours or days
  - propagator and Hawkes models: Appendix
- Empirical law for the peak impact (from the first to the last trade of the metaorder) Tóth et al. (2011), Bouchaud, Bonart, Donier & Gould (2018)
  - $\Delta P \approx Y \sigma_{\text{day}} \sqrt{Q/V_{\text{day}}}$, with ΔP the relative price move caused by the metaorder, Q its size, V_{day} the daily volume, σ_{day} the daily volatility, Y a constant of order 1
- Concave: trading 1% of daily volume moves the price by about $0.1\,\sigma$, not $0.01\,\sigma$
  - liquidity near the price is scarce and refills slowly (section *The Limit Order Book*)
- Cost: if impact builds as $I(q) = Y\sigma\sqrt{q/V}$ during execution, the average price paid is $\frac{2}{3} I(Q)$
  - q: quantity executed so far
  - before spread and fees
- Testing it needs the traders' own order records
  - public bars allow only an indirect look (Further Topics)

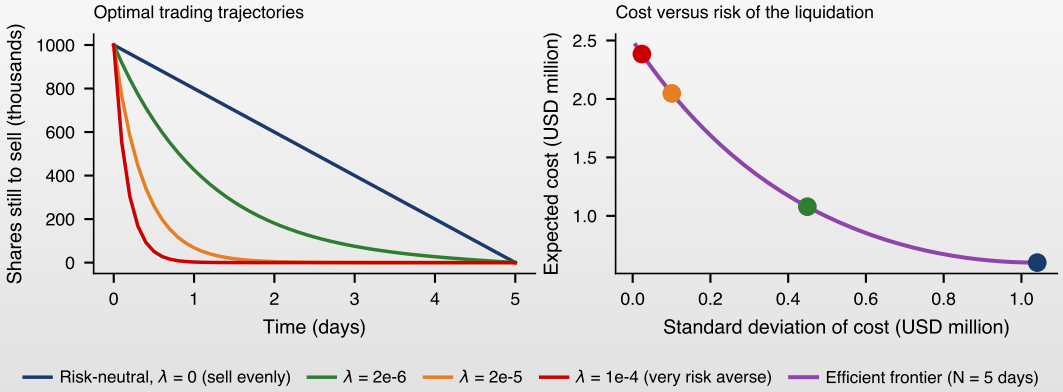
Optimal Execution: the Trade-Off

- Sell X shares within T days, in N steps of length $\tau = T/N$ Almgren & Chriss (2001)
  - holdings x_k from $x_0 = X$ to $x_N = 0$, trades $n_k = x_{k-1} - x_k$
- The market price moves with news and is pushed down by our own sales
  - $S_k = S_{k-1} + \sigma\sqrt{\tau}\,\xi_k - \gamma\, n_k$
  - the price we receive: $\tilde{S}_k = S_{k-1} - \eta\, n_k/\tau$
  - σ: price volatility (USD per day$^{1/2}$); ξ_k: news shock, i.i.d., mean 0, variance 1
  - $\gamma \geq 0$: **permanent** impact, lowers all later prices; $\eta > 0$: **temporary** impact, paid only by the trade itself
- **Trading fast**: high temporary impact, low price risk; **trading slowly**: the opposite
- Objective, over deterministic schedules: min $\mathbb{E}[C] + \lambda \mathrm{Var}[C]$
  - $C = XS_0 - \sum_k n_k \tilde{S}_k$: the implementation cost
  - value at the initial price S_0 minus what the sales bring
  - $\lambda \geq 0$: risk aversion (not Kyle's λ); $\lambda = 0$ minimises the expected cost only

Almgren–Chriss: the Optimal Trajectory

- Continuous time, holdings $x(t)$, $t \in [0, T]$, selling speed $-\dot{x}$
  - $\min \int_0^T (\eta\, \dot{x}^2 + \lambda\sigma^2 x^2)\, dt, \quad x(0) = X,\ x(T) = 0$
  - $\eta\, \dot{x}^2$: selling fast is costly; $\lambda\sigma^2 x^2$: holding a large position is risky
- Solution (derivation: Appendix, seminar A7): $x(t) = X\, \dfrac{\sinh(\kappa(T-t))}{\sinh(\kappa T)}, \qquad \kappa = \sqrt{\lambda\sigma^2/\eta}$, with $\sinh z = (e^z - e^{-z})/2$
  - κ: the urgency of the liquidation, it rises with risk aversion and volatility, falls with impact; $1/\kappa$: the time scale of the liquidation
  - $\lambda = 0$: $\kappa = 0$ and holdings fall linearly (trade evenly)
  - Large λ: most of the order is traded at the start, cutting the exposure to price risk

Optimal Trajectories and the Efficient Frontier (1/2)



MFM_ch10_execution ▶ Interactive

Optimal Trajectories and the Efficient Frontier (2/2)

- Left: shares still to sell over the 5 days, for four values of λ
  - right: expected cost against the standard deviation of cost, USD million (the efficient frontier)
- Numerical example of Almgren & Chriss (2001): $X = 10^6$ shares, $S_0 = 50$ USD, $T = 5$ days
  - $\sigma = 0.95$ USD per day$^{1/2}$, $\eta = 2.5 \cdot 10^{-6}$, $\gamma = 2.5 \cdot 10^{-7}$
- With $\lambda = 2 \cdot 10^{-6}$ the first day sells 57% of the position
  - the standard deviation of cost falls from 1.04 (even selling) to 0.45 million USD
  - for 0.48 million USD more expected cost
- The frontier is convex: the first units of risk reduction are cheap, the last ones expensive
  - full table: Appendix

MFM_ch10_execution

Execution Benchmarks

- **Implementation shortfall**: paper portfolio at the decision price minus the executed portfolio Perold (1988)
- VWAP: the volume-weighted average price over the period
  - TWAP: the time-weighted one, matched by the $\lambda = 0$ schedule
- A participation rule (a fixed share of market volume) ignores price
  - the rule at the heart of the Flash Crash (next section)

High-Frequency Trading: the New Market Makers



A microwave relay tower (Mojave, California); similar line-of-sight links carry trading signals between Chicago and New Jersey

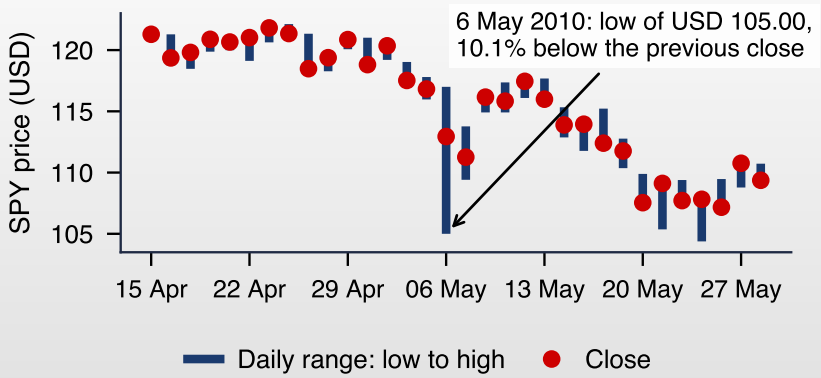
Photo: Tony Webster (2016); CC BY-SA 2.0; Wikimedia Commons

- **Question**: what does speed change in liquidity provision, and who gains from it?
- HFT (high-frequency trading): algorithms that trade in micro- to milliseconds and end the day flat
  - speed from co-location (servers next to the exchange's matching engine) and microwave links
- Low latency speeds up price discovery and narrows spreads Hasbrouck & Saar (2013)
- HFT liquidity is plentiful in calm markets and fragile under stress

High-Frequency Trading: Evidence and Models

- One HFT market maker on Chi-X and Euronext Menkveld (2013)
  - mostly passive orders: it earned the spread and lost on positions
- Order flow and latency become sources of information in their own right O'Hara (2015)
- HFT market makers as intermediaries: Jovanovic & Menkveld (2010), Pele (2025)
- Inventory skews quotes: a dealer who is long lowers both quotes to sell Avellaneda & Stoikov (2008) (Appendix)

The Flash Crash of 6 May 2010



MFM_ch10_stress ▶ Interactive

Flash Crash: Timeline and Regulatory Responses

- Chart: **daily** SPY range and close, 15 Apr – 28 May 2010 (unadjusted)
  - daily bars cannot show the intraday collapse and recovery
- Source: report of the CFTC and the SEC CFTC & SEC (2010)
  - CFTC: Commodity Futures Trading Commission
  - SEC: Securities and Exchange Commission
- 2:32 p.m.: a mutual fund complex starts selling 75 000 E-mini S&P 500 (Standard & Poor's) futures
  - at 9% of volume, regardless of price
- 2:45:28 p.m.: a 5-second pause lets buyers return
- SPY low USD 105.00, 10.1% below the previous close (USD 116.82)
  - the day closed only 3.3% lower, on volume 3.0 times the recent median
- HFT did not trigger the crash but amplified it Kirilenko, Kyle, Samadi & Tuzun (2017)
- Responses: single-stock circuit breakers and rules for erroneous trades
  - circuit breakers replaced in 2013 by the Limit Up–Limit Down mechanism
  - details: Appendix

MFM_ch10_stress

Crypto Venues and Arbitrage

- **Question**: where is the price formed when one asset trades on many venues at once?
- 24/7 trading with no auctions
  - the same asset trades on many order books of centralised exchanges (CEX) at once
  - and in pools of decentralised exchanges (DEX) priced by a formula (automated market maker, AMM, Chapter 16)
  - The same Bitcoin can trade at different prices on different exchanges Makarov & Schoar (2020)
- Price gaps are much larger **across countries** than within a country
  - Capital controls limit the movement of arbitrage capital
- A common component of signed volume across exchanges explains 80% of Bitcoin returns Makarov & Schoar (2020)
  - price discovery is shared across venues
- Crypto liquidity can be forecast from market data Agakishiev, Härdle & Pele (2025)

Price Discovery across Venues: Information Shares (1/2)

- Two log prices of one asset on two venues, $y_t = (p_{1t}, p_{2t})'$
  - each is I(1) (a random walk-like series)
  - their gap $\beta' y_t = p_{1t} - p_{2t}$, with cointegrating vector $\beta = (1, -1)'$, is stationary
- VECM (vector error-correction model) Johansen (1991): each price reacts to yesterday's gap and to past changes
  - $\Delta y_t = \alpha\, \beta' y_{t-1} + \sum_{i=1}^{p} \Gamma_i\, \Delta y_{t-i} + e_t, \qquad \mathrm{Var}(e_t) = \Omega$
  - $\alpha = (\alpha_1, \alpha_2)'$: adjustment coefficients, the speed at which each price closes the gap
  - Γ_i: 2×2 matrices of short-run dynamics, p lags; e_t: innovations of the two prices
- The common efficient price, the stochastic trend shared by both prices, has innovation $\psi' e_t$
  - $\psi' \propto \alpha'_\perp$ ($\propto$: proportional to; $\alpha'_\perp \alpha = 0$)

Price Discovery across Venues: Information Shares (2/2)

- **Hasbrouck information share** Hasbrouck (1995), venue j's share of the variance of the common innovation
  - $IS_j = \dfrac{([\psi' F]_j)^2}{\psi' \Omega \psi} \in [0,1], \quad IS_1 + IS_2 = 1$
  - F: Cholesky factor of Ω (lower triangular, $FF' = \Omega$); $[\cdot]_j$: the j-th element
  - Depends on the ordering when Ω is not diagonal: report the bounds over both orderings
- **Gonzalo–Granger component share** Gonzalo & Granger (1995): $\alpha_\perp$ normalised to sum 1
  - for two venues $CS_1 = \alpha_2/(\alpha_2 - \alpha_1)$
  - in $[0,1]$ only if the two adjustment coefficients have opposite signs
- A price that does **not** adjust ($\alpha_j = 0$, weak exogeneity) has a component share of 1
  - Hasbrouck's bounds can stay below 1 when the innovations are correlated, and they depend on the ordering: with $\alpha_1 = 0$, equal variances and correlation ρ, the bounds of the first market can be $[1 - \rho^2, 1]$
  - Inference: a Wald test of $\alpha_j = 0$; non-rejection does not prove $\alpha_j = 0$

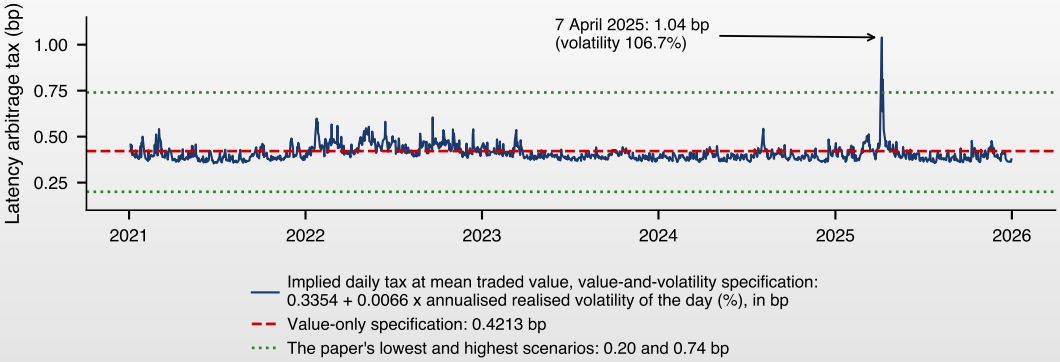
Case Study: the Speed Race and the Question

- **Latency arbitrage** (sniping) Budish, Cramton & Shim (2015)
  - trading against quotes made stale by a public price move before their provider can cancel them
  - In a continuous book the fastest wins; liquidity providers pay through wider spreads and thinner books; proposal: **frequent batch auctions**, e.g. every second, moving competition from speed to price
- **Question** Aquilina, Budish & O'Neill (2022): how much does the race cost a whole stock market?
- Why it was hard: order-book data record only the winners
  - a failed order or a cancel that arrives too late leaves the book unchanged
  - **Message data**: every message to and from the exchange, including "too late" replies, so both winners and losers are seen
- Is sniping on **public** news a cost of liquidity?
  - link to the section *The Bid-Ask Spread*
  - next to **private**-information adverse selection (section *Information and Prices*)

Case Study: Measurement and Three Central Results

- **Measurement chain**: London Stock Exchange messages, 43 days of 2015
  - FTSE 350 (the 350 largest London-listed companies)
  - about 2.2 billion messages (definitions: Appendix)
  - Race: 2+ participants, same stock, price and side, within the information horizon (at most 500 μs, microseconds), at least one success and one failure; race profit marked to the mid 10 s later
  - **Latency arbitrage tax**: daily race profits / daily traded value, in bp
- **1. Frequent and fast**: 537 races a day per FTSE 100 stock, about one a minute
  - decided by 5–10 μs
  - 22% of FTSE 100 traded value is in races
- **2. Concentrated**: the top six firms win 82% of races (a count of races)
  - a race is mostly a transfer from slower liquidity providers to a few fast takers
- **3. Small prizes, large sums**: tax 0.419 bp of all traded value
  - eliminating races would cut the cost of liquidity by about 17%
  - USD 4.8 billion a year in global equities
- Type of evidence: description
  - one exchange, 43 days; the 17% cut is a counterfactual computed from this measurement

Case Study: the Implied Daily Tax on SPY (1/2)



MFM_ch10_race ▶ Interactive

Case Study: the Implied Daily Tax on SPY (2/2)

- Implied tax $0.3354 + 0.0066\,\sigma_t$ bp, at mean traded value
  - from the specification with traded value and volatility
- σ_t: annualised realised volatility (%) from SPY 5-minute bars, 1 244 full days of 2021–2025
- The median tax is 0.40 bp, close to the 0.419 bp measured in London
- On 30% of days the tax exceeds 0.4213 bp, the value of the specification with traded value only
- The peak, 1.04 bp, is on 7 April 2025 ($\sigma_t = 106.7\%$)
  - only 2 days exceed the authors' highest scenario (0.74 bp)
- The volatility term matters only on stress days: most of the prize scales with traded value
- Type of evidence: description
  - a transfer of the London coefficients, not a measurement on SPY (next slide)
  - the authors' tax scenarios differ by a factor of 3.7: the tax level is the main uncertainty (Further Topics)

MFM_ch10_race

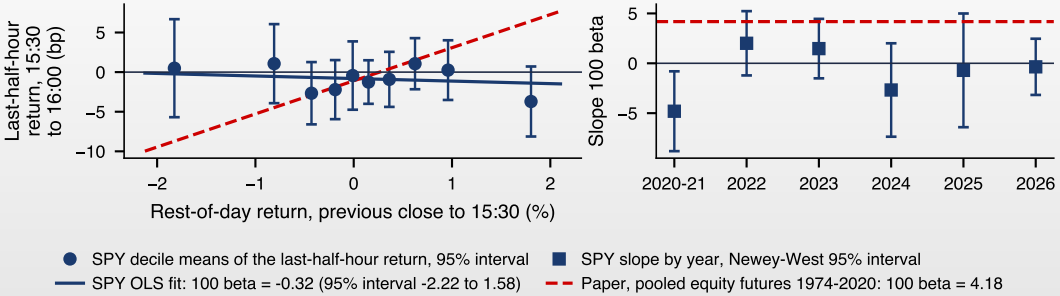
Case Study: What the Transfer Assumes

- **Method**: coefficient-transfer extrapolation
  - the coefficients estimated in London are applied to SPY's traded value and volatility
  - races are not re-estimated on SPY (no US message data)
  - It assumes the same race technology and the same tax per unit of value; volatility here is from 5-minute returns, the authors use 1-minute returns
- **Limits stated by the authors**: one exchange, nine weeks
  - market structure (fragmentation, ETFs) and other asset classes ignored
  - long races and silent losers missed
  - For scale: the futures–SPY arbitrage of Budish, Cramton & Shim (2015) was about USD 75 million a year; the implied SPY prize is four to six times larger
- **Open extension**: does SPY's daily EDGE spread co-move with the implied tax?
  - both depend on volatility, so this is a diagnostic
  - identification needs a message-based race measure and a matched effective-spread benchmark

Case Study: Market Intraday Momentum

- Baltussen, Da, Lammers & Martens (2021), *Journal of Financial Economics*: **market intraday momentum**
  - the return of the last half hour before the close (r_{LH}) follows the sign of the return over the rest of the day (r_{ROD}: from the previous close to 30 minutes before the close)
  - the paper's data: over 60 futures on equity indices, bonds, commodities and currencies, Dec 1974 – May 2020
- The proposed mechanism: **gamma hedging**
  - delta: the sensitivity of an option's value to the price of the underlying; gamma: the change in delta when the price moves
  - short-gamma market makers rebalance their delta hedge: they buy after rises and sell after falls
  - leveraged ETFs rebalance daily towards their target leveraged exposure (e.g. 2×), also in the direction of the day's move
  - much of this trading happens near the close, when volume peaks (the U-shape of SPY volume)
- The paper's results, with slopes and R^2 multiplied by 100
  - equity index futures, pooled regression: $100\beta = 4.18$ ($t = 7.29$), adjusted $R^2 = 2.45\%$ (Table 2)
  - S&P 500, days with negative net gamma exposure of market makers: $100\beta = 6.63$; with positive exposure: 0.82, not significant (Table 7)
  - the price pressure is transitory: the return reverses over the next days
- The question: does the effect hold for SPY after the end of the paper's sample?

Case Study: SPY Intraday Momentum, 2020–2026



- Left: mean of r_{LH} (15:30–16:00, in bp) by decile of r_{ROD}, 1 460 days, Oct 2020 – Sep 2026
- Right: the slope 100β by year (2020-21: Oct 2020 – Dec 2021; 2026: Jan – Sep)
  - the only interval that excludes 0 is that of 2020-21: slope -4.81, the opposite sign to the paper's ($p = 0.019$)
  - after the Holm correction for the six subperiods, $p = 0.11$: an exploratory result, not significant

MFM_ch10_intraday_momentum ▶ Interactive

Case Study: the Paper's Regression on SPY Bars (1/2)

- Equation (7) of Baltussen, Da, Lammers & Martens (2021): $r_{LH,t} = \alpha + \beta\, r_{ROD,t} + \varepsilon_t$
  - β: regression slope (not Kyle's β); $\beta > 0$: the last half hour continues the move of the day (momentum); $\beta < 0$: it reverses it
- Course data: SPY 5-minute bars, regular session 09:30–16:00
  - 1 460 days with all 78 bars whose previous trading day is also complete
  - a bar's timestamp marks its start: the 15:25 bar closes at 15:30, the 15:55 bar at 16:00
  - r_{ROD}: from the previous close (close of the 15:55 bar) to the close of the 15:25 bar, including the night; on ex-dividend dates the dividend yield is added
  - r_{LH}: from the close of the 15:25 bar to the close of the 15:55 bar, i.e. six 5-minute bars
  - whether the last bar contains the closing auction depends on the data vendor's convention, which the data do not show
- The main test: the regression on the last 30 minutes, all days
  - the subperiods and the other lengths of the last interval are exploratory: their p-values get the Holm correction in families of six tests

Case Study: the Paper's Regression on SPY Bars (2/2)

- What the replication does differently
  - the SPY ETF instead of the S&P 500 futures contract; a single asset
  - a replication over a later period: the sample starts five months after the end of the paper's
  - coefficients are estimated on the whole sample, not by rolling or expanding forecasts
  - Newey–West standard errors with 5 lags, as in Table 7
- The net gamma exposure, computed by the authors from S&P 500 option positions, is not available
  - the split by the sign of r_{ROD} is a descriptive analysis by the sign of the return, without identifying gamma exposure
  - the mechanism of Baltussen, Da, Lammers & Martens (2021) remains a hypothesis these data do not test
- Two variants from the paper
  - the predictor of Gao, Han, Li & Zhou (2018): r_{ONFH}, from the previous close to 10:00
  - the timing strategy: $R_t = w_t \left(e^{r_{LH,t}} - 1\right)$, with $w_t = +1$ if $r_{ROD,t} > 0$ and $w_t = -1$ otherwise; simple gross return, before costs

Case Study: No Significant Positive Intraday Momentum in SPY

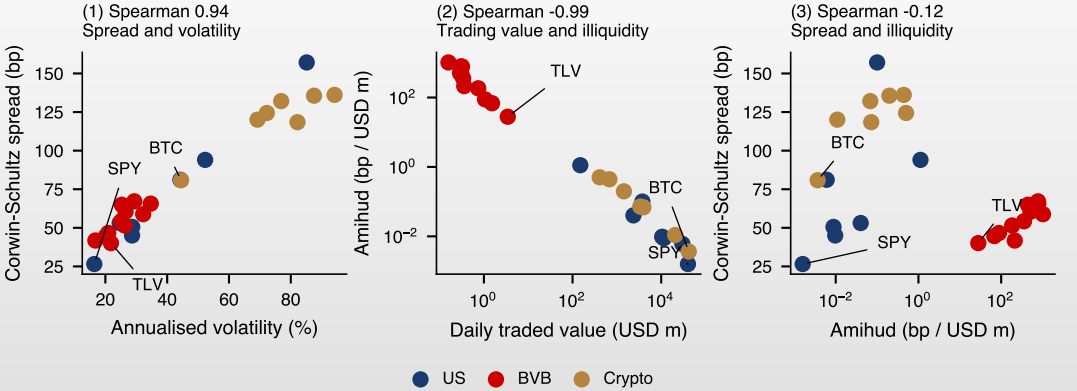
Test	Course data	Interpretation
r_{ROD}, all days (main test)	$100\beta = -0.32$ ($t = -0.33$), 95% interval: $[-2.22, 1.58]$; $R^2 = 0.02\%$, adjusted $R^2 = -0.053\%$	the paper's 4.18 lies outside the interval
r_{ONFH}, Gao, Han, Li & Zhou (2018)	$100\beta = -2.68$ ($t = -1.87$)	this predictor shows no significant positive momentum either
Down days / up days	$100\beta = -0.65$ ($t = -0.26$) / 0.50 ($t = 0.32$); difference, common calendar: -1.14 (HAC SE 2.98, $p = 0.70$)	we find no difference between the slopes
Last 10 – 60 minutes (exploratory)	100β between -1.48 (last 10 minutes, $p = 0.043$) and 1.15 (last 60 minutes)	negative coefficient, not significant after the Holm correction ($p = 0.26$)
Timing strategy, gross	-0.47 bp per day (HAC interval: $[-1.76, 0.81]$); -1.20% per year; Sharpe -0.29; 48.8% of days profitable	negative mean performance in the sample, not significant

- Type of evidence: description
  - a replication over a later period on one asset; the interval excludes the paper's slope, but not a small positive slope (up to 1.58)
  - reversal to the next close: $100\beta = -9.04$ ($t = -0.62$, 1 445 days), not significant
- Limits: one ETF, not futures; the gamma exposure is missing; six years of data
- Open question: did the effect disappear after publication, or did the gamma position of market makers change once S&P 500 options expiring every day appeared (2022)?

An Open Question: Do Price-Based Proxies Measure Trading Costs?

- Where trades and quotes are unavailable, costs are estimated from daily OHLC bars
  - e.g. BVB history, many crypto venues
  - Jahan-Parvar & Zikes (2023): low-frequency estimators track effective spreads only when the spread is large relative to volatility; otherwise volatility dominates
  - Ardia, Guidotti & Kroencke (2024): EDGE is unbiased and efficient under their assumptions
- Formal question, across assets i: $\hat{S}_i = S_i + b(\sigma_i) + \varepsilon_i$
  - i.e. the proxy estimate is the true spread plus a bias that depends on the volatility σ_i
  - plus estimation noise
- H_0: $b \equiv 0$ (the proxy ranks assets by cost; $\equiv$: for every σ_i) against H_1: b increasing in σ_i
- Identification needs a cost benchmark measured independently of the proxy
  - effective spreads from matched trades and quotes
  - Amihud measures price impact, not the spread, and itself contains $|r|$
- Open: is the spread-to-volatility ratio high enough for reliable estimation?
  - where direct cost data are scarce
  - thin trading alone does not determine that ratio

Mini-Case: Do Spread Estimators Measure Costs or Volatility?



MFM_ch10_seminar ▶ Interactive

Testing the AI Claim with Spearman Rank Correlations

- Chart: three panels, log scales for Amihud and traded value
  - (1) Corwin–Schultz against annualised volatility
  - (2) the Amihud ratio against daily traded value
  - (3) Corwin–Schultz against Amihud
- A typical claim of an artificial intelligence (AI) assistant
  - "the estimates are similar, so BVB blue chips are about as liquid as US large caps"
  - test on the 25 assets, Sep 2024 – Sep 2026
  - with Spearman rank correlations ρ and their p-values p
- Corwin–Schultz with volatility $\rho = 0.94$ ($p < 0.001$), with Amihud $\rho = -0.12$ ($p = 0.58$)
  - Difference between the two Spearman correlations 1.06, 95% CI $[0.58, 1.44]$ over assets, $[0.71, 1.43]$ within markets
  - Partial Corwin–Schultz–Amihud ρ controlling for volatility 0.27
- Daily Corwin–Schultz ranks track volatility, as Jahan-Parvar & Zikes (2023) predict
  - a diagnostic, not a test of H_0: Amihud is not an independent spread benchmark
- Type of evidence: association

MFM_ch10_seminar

The Discovery Loop and a Project Seed

- **Loop**, with an AI assistant (prompts: Appendix)
  - 1. Literature: collect validation studies and check every DOI (digital object identifier) on Crossref; 2. rival hypotheses, each with a cross-sectional prediction
  - 3. Replicate the chapter's numbers first; 4. a pre-registered robustness plan; 5. ask the assistant to attack the benchmark; 6. report everything tried
- **Seed**: which estimator from daily bars, if any, measures the level of trading costs?
  - on the BVB and for crypto-assets
- Deliverables
  - A dated pre-registration; EDGE implemented after Ardia, Guidotti & Kroencke (2024) and validated on simulated prices with a known spread; code and a two-page report with every specification tried
  - Declare AI use in `AI_USE.md` and the errors caught in `AI_ERRORS.md`: github.com/danpele/MFM/tree/main/project

Which Liquidity Measure for Which Data?

Available data	Suggested measure
Trades and quotes with timestamps	Effective and realised spread; MRR by the generalised method of moments (GMM); Hasbrouck VAR; PIN
Prices of one asset on several venues	VECM; information-share bounds and component shares
Intraday bars (1–5 minutes)	Corwin–Schultz, Abdi–Ranaldo or EDGE per day (EDGE is the closest to one tick for SPY); intraday Amihud
Daily prices and volumes	Amihud ratio (price impact) across assets; EDGE for spread levels only if validated against effective spreads; high-low estimators mainly for changes over time
Own executions	Implementation shortfall; calibrate the square-root law
Theory, what-if analysis	Simulated order books; Glosten–Milgrom, Kyle, Almgren–Chriss

Check Your Understanding, and the Road to Continuous Time

Can you answer these?

- ⊡ Why can a positive Roll covariance occur even when the Roll model is true?
- ⊡ Why do daily high-low estimators rank BVB stocks next to US large caps?
- ⊡ In Kyle's model, what happens to λ and to the insider's profit when noise trading doubles?
- ⊡ Why does a risk-averse seller front-load?

Chapter 11: continuous time

- ⊡ From trades and order books back to continuous-time price models: Brownian motion and Itô calculus
- ⊡ At high frequency the bid-ask bounce is microstructure noise
  - ▶ it contaminates realised volatility (Chapter 9)

The Practical Decision

- **What I choose**
  - for spread levels from bars: EDGE on 5-minute bars (0.46 bp for SPY, the closest to the 0.21 bp tick); from daily bars, only changes over time, not levels
  - for comparing liquidity across markets: the Amihud ratio (Banca Transilvania: 28.0 bp per USD 1 million, about 17 600 times SPY)
  - for a large order: the Almgren–Chriss schedule; with $\lambda = 2 \cdot 10^{-6}$ it sells 57% on the first day, and the standard deviation of cost falls from 1.04 to 0.45 million USD for 0.48 million USD more expected cost
- **Under which conditions**
  - daily proxies follow volatility (Corwin–Schultz and volatility: $\rho = 0.94$), so levels need the effective spread from trades and quotes
  - Amihud measures price impact, not the spread, and itself contains $|r|$
  - the impact parameters are calibrated on my own executions
- **What result would make me drop it**
  - EDGE for levels: the difference EDGE minus effective spread, on the same days and with comparable weights, far from 0 under inference robust to dependence across days
  - the EDGE interval for SPY, $[0.41, 0.51]$ bp, concerns the mean of the estimator, not the true spread
  - Almgren–Chriss: realised costs on my own executions that deviate systematically from those predicted with the calibrated linear impact; a nonlinear impact requires another cost functional and a new optimisation

The Evidence: SPY, BVB, Crypto and the VIX (1/2)

Series	Content	Period
SPY, 5-minute bars	ETF (exchange-traded fund) on the S&P 500; regular session 09:30–16:00 New York time; 1 477 full days	2020–2026
Bitcoin, 5-minute bars	24 hours a day, 7 days a week; 208 596 bars	2024–2026
US stocks, daily	SPY, Apple, Microsoft, Nvidia, JPMorgan, GameStop, Strategy	2016–2026
BVB stocks, daily	Ten blue chips and the ETF on the BET (Bucharest Exchange Trading, the BVB benchmark index); values in lei converted to USD at the reference rate of the National Bank of Romania (BNR)	2016–2026
Crypto-assets, daily	Bitcoin, Ethereum, Solana, XRP, Dogecoin, Cardano, Litecoin	2016–2026
VIX, daily	Index of the 30-day implied volatility of the S&P 500, computed from option prices	2020–2026

- Daily data: open, high, low, close and volume
  - prices adjusted for dividends and splits; days without trades removed
  - traded value: unadjusted price $\times$ volume, on the end-of-sample share count (documented splits and bonus shares)
- For intraday charts, r is the change in the log close from one bar to the next

The Evidence: SPY, BVB, Crypto and the VIX (2/2)

- Some order books and the bid-ask bounce are **simulated**, and labelled as such
  - order books: section *The Limit Order Book*
  - bid-ask bounce: section *The Bid-Ask Spread*

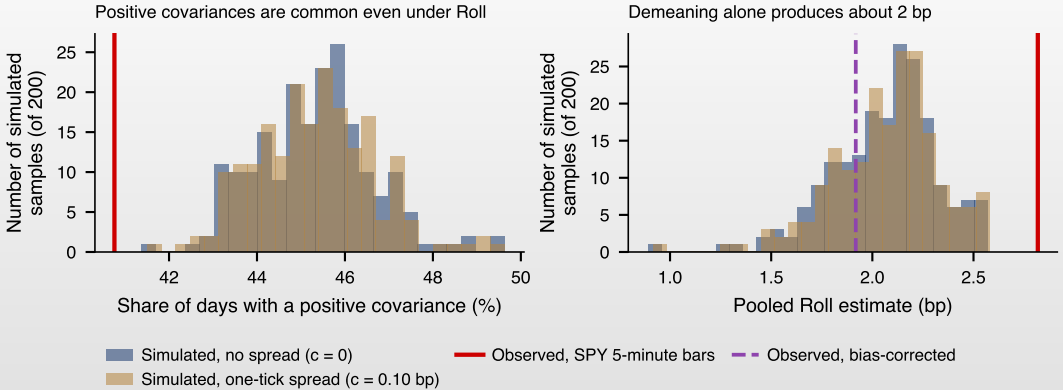
Notation across Microstructure Models

Symbol	Meaning, by model
T	Roll in finite samples: number of price changes (77 per SPY day); Almgren–Chriss, Avellaneda–Stoikov: the horizon
α	Corwin–Schultz: spread part of the log range; PIN: probability of an information event; VECM: adjustment coefficients; simulated book: limit-order rate
β	Corwin–Schultz: sum of the squared one-day ranges; Kyle: the insider's trading intensity; VECM: cointegrating vector; OFI: impact slope; intraday momentum: slope of regression (7)
γ	Roll: autocovariances γ_k; Corwin–Schultz: squared two-day range; Almgren–Chriss: permanent impact; Avellaneda–Stoikov: risk aversion
λ	Kyle: price move per share of net order flow; Almgren–Chriss: risk aversion; PIN: Poisson rate in $\pi(n;\lambda)$; Hawkes: intensity $\lambda(t)$
μ	Glosten–Milgrom: share of informed traders; PIN: arrival rate of informed orders; simulated book: market-order rate; Hawkes: baseline rate
η	Abdi–Ranaldo and EDGE: log mid-range η_t; Almgren–Chriss: temporary impact
δ, ε	δ: PIN bad-news probability, cancellation rate, quote distance (Avellaneda–Stoikov); ε: Roll news shock, PIN uninformed rates $\varepsilon_b, \varepsilon_s$, order sign (propagator)

Demeaning Pulls the Wrong-Sign Share Below One Half

- The form $\Phi(\sqrt{T}\,\gamma_1/\gamma_0)$ centres $\hat{\gamma}_1$ at γ_1 and neglects γ_1^2
  - demeaning moves the centre to $\gamma_1 - (\gamma_0 + 2\gamma_1)/T$
  - With the binary-sign standard deviation $s_\gamma = \sqrt{(\gamma_0^2 + \gamma_1^2)/T}$: $P(\hat{\gamma}_1 > 0) \approx \Phi\Big(\dfrac{\gamma_1 - (\gamma_0 + 2\gamma_1)/T}{s_\gamma}\Big)$
- No spread ($\gamma_1 = 0$): the argument is $-1/\sqrt{T}$, so $P \approx \Phi(-1/\sqrt{T}) < 1/2$
  - positive covariances on fewer than half of the days, from demeaning alone
- SPY, $T = 77$, one-tick spread
  - $\gamma_0 \approx 83.9$ bp^2 (the 9.2 bp bar standard deviation, squared)
  - $\gamma_1 = -c^2 \approx -0.011$ bp^2
  - shift $(\gamma_0 + 2\gamma_1)/T \approx 1.09$ bp^2
  - $\Phi(-0.115) \approx 45.4\%$, inside the simulated band of the next slides; the unshifted formula gives 49.96%, outside it
- Delta-method standard error and the daily-data case: Appendix
  - seminar A1–A2 derive them step by step

SPY: Sampling Error or Misspecification? (1/2)



MFM_ch10_seminar

SPY: Sampling Error or Misspecification? (2/2)

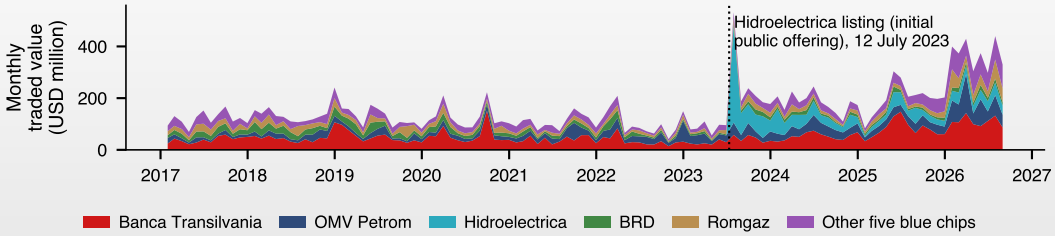
- Left: share of days with $\hat{\gamma}_1 > 0$ in 200 samples simulated under the Roll model
  - 77 five-minute returns a day, with no spread and with a one-tick spread
  - red line: the value in the data
- Right: the Roll estimate from the autocovariance averaged over days, same simulations
  - red line: observed (2.82 bp)
  - dashed line: bias-corrected (1.92 bp)
- With a one-tick spread, $\hat{\gamma}_1 > 0$ on 45.3% of days (band: [43.0, 48.1]), close to the approximation with spread, $\Phi(-0.115) \approx 45.4\%$ (no spread: $\Phi(-1/\sqrt{77}) \approx 45.5\%$)
- In the data the share is 40.8%, below the band
  - the negative autocorrelation is stronger than a one-tick bid-ask bounce
- No spread: the Roll estimator from the averaged autocovariance still gives 2.06 bp
  - band: [1.50, 2.53]
  - The mean over days $\bar{g}_1 = -1.99$ bp^2 gives the observed 2.82 bp; about 54% of it (-1.07) is demeaning bias
  - Correction, with $\hat{\gamma}_0 \approx 83.9$: $(\bar{g}_1 + \hat{\gamma}_0/T)/(1 - 2/T) \approx -0.92$, i.e. a spread of $2\sqrt{0.92} \approx 1.92$ bp: after the approximate demeaning correction, negative autocovariance remains that the one-tick Roll benchmark does not explain; squares add up: $2.06^2 + 1.92^2 \approx 2.82^2$

MFM_ch10_seminar

Corwin–Schultz: Reading an Estimate, and Two Limit Cases

- $\hat{S} = 0.001$: a round trip costs about 10 bp
- $\alpha < 0$ (volatility larger than the model expects) gives $\hat{S} < 0$
  - negative two-day estimates are set to 0, then averaged over all pairs of days
- For daily data the paper adjusts for overnight moves (Appendix)
- With $3 - 2\sqrt{2} = (\sqrt{2}-1)^2$ and $\sqrt{2\beta} - \sqrt{\beta} = (\sqrt{2}-1)\sqrt{\beta}$, the formula becomes $\alpha = (\sqrt{\beta} - \sqrt{\gamma})/(\sqrt{2}-1)$
- **Test 1, pure volatility**: no spread is found where there is none
  - one-day log ranges 0.01 and 0.01, two-day range $0.01\sqrt{2}$
  - so $\beta = 2 \times 0.01^2 = \gamma$, $\alpha = 0$ and $\hat{S} = 0$
- **Test 2, pure spread**: the whole range is spread
  - the price only bounces between a bid and an ask 1% apart, so every range is 0.01
  - $\beta = 2 \times 0.01^2$, $\gamma = 0.01^2$, $\alpha = 0.01$ and $\hat{S} \approx 0.01$ (100 bp)

Traded Value of BVB Blue Chips



- Chart: monthly traded value of the ten blue chips, USD at the BNR reference rate
  - annual totals (sums of the monthly values): 1 428 million in 2017, 1 248 million in 2022, 2 380 million in 2025
- Concentration: Banca Transilvania 41% of the 2025 value, Hidroelectrica (listed July 2023) 14%
- One median SPY day (USD 40 124 million) is about 17 times the 2025 total of the ten blue chips
  - a position of a few million USD takes days to build or unwind
- The BVB in brief: Appendix

MFM_ch10_bvb

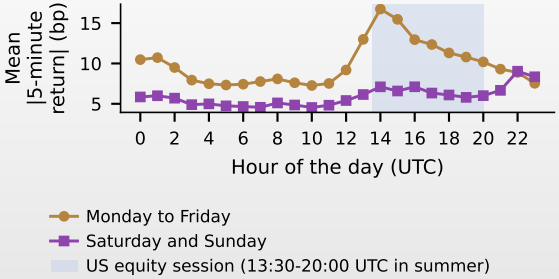
Nonsynchronous Trading: the Lo–MacKinlay Model

- **Question**: what do observed returns look like when an asset does not trade every period? Lo & MacKinlay (1990)
  - the case of many BVB stocks
- Virtual returns of the unobserved price: $r_{it} = \mu_i + \beta_i f_t + \varepsilon_{it}$
  - f_t: common factor, i.i.d., with variance σ_f^2, ε_{it}: idiosyncratic part, independent of f
  - In each period asset i does not trade with probability π_i, independently over time and across assets
- Observed return r_{it}^o: 0 in a period without trades
  - otherwise the sum of the virtual returns accumulated since the last trade
- Moments of one asset, for $n \geq 1$, with $\sigma_i^2 = \mathrm{Var}(r_{it})$: $\mathrm{Var}(r_{it}^o) = \sigma_i^2 + \dfrac{2\pi_i \mu_i^2}{1-\pi_i}$, $\qquad \mathrm{Cov}(r_{it}^o, r_{i,t+n}^o) = -\mu_i^2 \pi_i^n$
- The own autocorrelation is negative but negligible
  - μ_i^2, the squared daily mean, is small relative to σ_i^2

Nonsynchronous Trading: Cross-Autocorrelation and Beta

- Between two assets, for $n \geq 1$: the thinly traded asset j picks up the common factor with a lag
  - $\mathrm{Cov}(r^o_{it}, r^o_{j,t+n}) = \dfrac{(1-\pi_i)(1-\pi_j)}{1-\pi_i\pi_j}\,\beta_i\beta_j\sigma_f^2\,\pi_j^n$
- An equally weighted portfolio of many assets with the same π: $\mathrm{Corr}(r^o_{pt}, r^o_{p,t+n}) = \pi^n$
  - With $\pi = 0.3$ the observed index has first-order autocorrelation 0.30, although the virtual returns are independent over time
- The same-day regression beta of a thinly traded asset is biased towards zero; the correction of Scholes & Williams (1977):
  - $\hat\beta^{SW} = \dfrac{\hat b_{-1} + \hat b_0 + \hat b_{+1}}{1 + 2\hat\rho_m}$; $\hat b_k$: slope of r_{it} on $r_{m,t+k}$; $\hat\rho_m$: first-order autocorrelation of the market return r_m
- The Amihud ratio uses only days with volume
  - after a pause, $|r|$ accumulates several days of news
  - this tends to raise the ratio of thinly traded assets

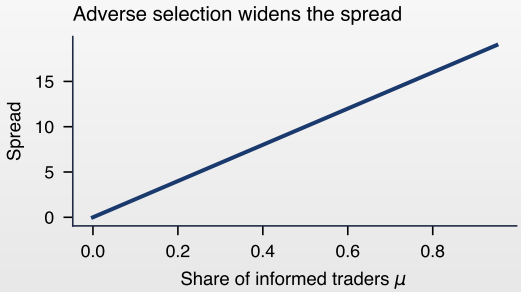
Bitcoin Around the Clock

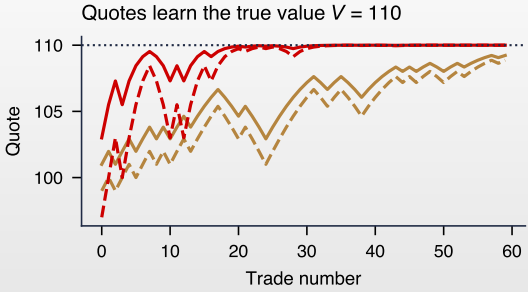


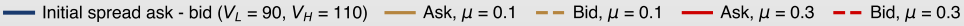
- Mean absolute 5-minute return by hour of Coordinated Universal Time (UTC)
  - Sep 2024 – Sep 2026, 208 596 bars
- Weekdays: peak 16.7 bp at 14:00 UTC
  - US open, macro releases; low 7.3 bp at 10:00 UTC
- Weekend 5.9 bp against 9.9 bp on weekdays
  - 1.69 times calmer (confidence intervals: seminar B7)
- A 24/7 market still follows the clock of banks, ETFs and traditional investors

MFM_ch10_intraday

Glosten–Milgrom in Action (1/2)







MFM_ch10_information

Glosten–Milgrom in Action (2/2)

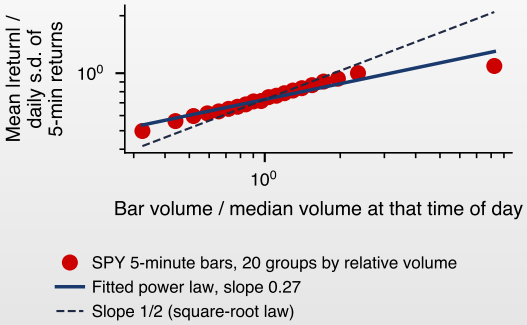
- Left: initial spread as a function of μ
- Right: one simulated sequence of 60 trades with true value $V = 110$ (same random draws for both μ)
- On these individual paths the bid exceeds 108 after 7 trades for $\mu = 0.3$, only after 48 for $\mu = 0.1$
- More informed traders: wider spreads, but faster price discovery
  - a trade-off between liquidity and informational efficiency
  - empirical counterpart: PIN (probability of informed trading) Easley, Kiefer, O'Hara & Paperman (1996)

MFM_ch10_information

VPIN: Order Imbalance in Volume Time

- **Why**: PIN needs maximum-likelihood estimation on daily order counts
  - fragile at today's volumes (Appendix)
- VPIN (Volume-synchronised Probability of INformed trading) Easley, López de Prado & O'Hara (2012) replaces the estimation with a direct average in volume time
  - trades are grouped into volume buckets of equal volume V
- Bulk volume classification: in time bar i of bucket τ, the buy fraction is $\Phi(\Delta P_i/\sigma_{\Delta P})$
  - $V_\tau^B = \sum_{i\in\tau} V_i\,\Phi(\Delta P_i/\sigma_{\Delta P})$, $\quad V_\tau^S = V - V_\tau^B$; ΔP_i: price change over bar i; $\sigma_{\Delta P}$: its standard deviation
- $\mathrm{VPIN} = \dfrac{1}{nV}\displaystyle\sum_{\tau=\mathbf{1}}^{n} |V_\tau^S - V_\tau^B|$, on a rolling window of the last n buckets
  - the authors: $n = 50$, $V = 1/50$ of average daily volume
- Link to PIN: in the PIN model $\mathbb{E}|V^S - V^B| \approx \alpha\mu$ and $\mathbb{E}[V^S + V^B] = \alpha\mu + \varepsilon_b + \varepsilon_s$, so VPIN approximates PIN
- Contested: correctly computed, VPIN did not predict the 2010 Flash Crash Andersen & Bondarenko (2014)
  - it largely tracks volume and volatility

Volume and Price Moves in SPY Bars



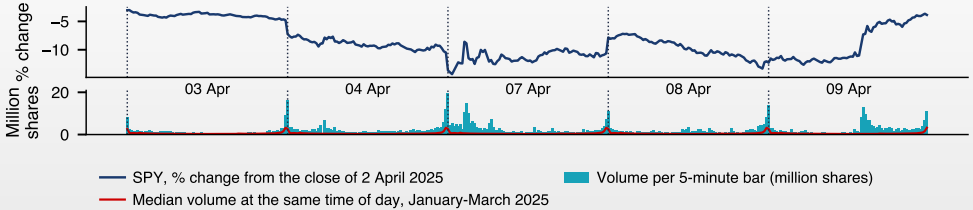
- 115 052 five-minute bars in 20 equal-size groups by relative volume
  - relative volume: bar volume / median volume at that time of day
- Mean $|r|$, scaled by the day's volatility, rises with slope 0.27 on log-log axes
  - concave, but flatter than 1/2
- An association between bars, not a metaorder law
  - news drives both volume and $|r|$ (simultaneity), and opposite trades in a bar net out
- Seminar B6 tries an instrumental variable and questions its exclusion restriction
  - for execution costs calibrate the metaorder law on your own trades

MFM_ch10_impact

Almgren–Chriss: the Optimal Trajectory

- Discrete time, with x_k the holdings after trade k, at $t_k = k\tau$: the same shape, $x_k = X\,\dfrac{\sinh(\kappa(T - t_k))}{\sinh(\kappa T)}$
  - κ solves $\frac{2}{\tau^2}(\cosh\kappa\tau - 1) = \frac{\lambda\sigma^2}{\tilde\eta}$, with $\cosh z = (e^z + e^{-z})/2$ and $\tilde\eta = \eta - \frac{\gamma\tau}{2} > 0$, the temporary impact corrected for the discrete steps
  - The equation tends to $\kappa^2 = \lambda\sigma^2/\eta$ as $\tau \to 0$
- Mean and variance of the cost of any schedule: $\mathbb{E}[C] = \frac{1}{2}\gamma X^2 + \frac{\tilde\eta}{\tau}\sum_{k=1}^{N} n_k^2$, $\quad \mathrm{Var}[C] = \sigma^2\tau\sum_{k=1}^{N} x_k^2$
  - $\frac{1}{2}\gamma X^2$: permanent impact, the same for every schedule
  - $\sum n_k^2$ penalises large trades, $\sum x_k^2$ holding the position long

A Recent Stress Episode: April 2025



- SPY 5-minute bars, 3–9 April 2025
  - bottom: bar volume against the Jan–Mar 2025 median at the same time of day
- 7 April: range 8.6%, realised volatility 6.7% (median 0.61%)
  - one bar +4.28% and one ten minutes later −2.92%
  - 9 April: +10.0% open to close; volume 5.4 times the usual benchmark
  - mean of the five days: volume 4.6 times the usual benchmark
- Bars document sharp moves and heavy trading
  - they cannot establish book depth, price continuity or the absence of halts

MFM_ch10_stress

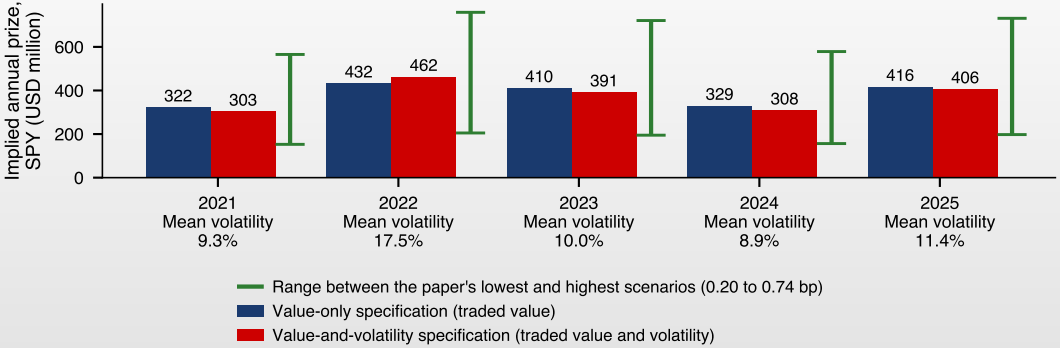
Price Discovery on the BVB: the BET ETF and the BET-TR Index (1/2)

- Daily log closes of the Patria-TVBETETF and the BET-TR index
  - the ETF reinvests dividends
  - BET-TR: BET Total Return, dividends reinvested
  - 495 common trading days, September 2024 – September 2026
- The basis (log ETF minus log index) has standard deviation 0.84%
  - BET-TR is computed from its constituents' prices, not traded on a venue
- **Johansen trace test**: tests the number r of cointegrating vectors, step by step from $r = 0$
  - 49.6 for $r = 0$ (5% critical value 15.49): rejected; 0.10 for $r \leq 1$ (3.84): not rejected
  - Hence one cointegrating vector; VECM with 1 lag, chosen by BIC (Bayesian information criterion)

Price Discovery on the BVB: the BET ETF and the BET-TR Index (2/2)

- Adjustment: ETF $\alpha = -0.121$ ($t = -2.11$); index $\alpha = -0.034$ ($t = -0.63$)
  - The index adjustment is not distinguishable from zero, the ETF's is: suggestive evidence that the ETF adjusts toward the index, with wide bootstrap intervals (seminar B8)
- Information share of the index: between 0.25 and 0.98, with a residual correlation of 0.93
- Component share of the index: $-0.121/(-0.121 + 0.034) = 1.39$, outside $[0, 1]$ because both α are negative
- Lesson: with daily closes the innovations are almost collinear, so the bounds are wide
  - price discovery needs intraday data (seminar B8)

Case Study: the Implied Annual Prize for SPY (1/2)



MFM_ch10_race

Case Study: the Implied Annual Prize for SPY (2/2)

- Annual prize $\sum_t \Pi_t$, USD million; Π_t: implied race profit on day t (tax $\times$ traded value)
- Bars: the specification with traded value only (0.4213 bp) and the one with value and volatility
  - whiskers: the authors' scenarios, 0.20–0.74 bp
- 2025: 416 million (value only) and 406 million (value and volatility)
- The volatility specification gives more only in high-volatility years
  - when the year's mean volatility exceeds $(0.4213 - 0.3354)/0.0066 = 13.0\%$
  - in 2022 (17.5%), not in the other years (8.9–11.4%)
- The scenarios differ by a factor of 3.7, the specifications by at most 7%
  - the tax level is the main uncertainty

MFM_ch10_race

Order Types and the Simulated Book

- Conditions on execution and lifetime
  - IOC (immediate-or-cancel): executes at once as far as possible, the rest is cancelled; FOK (fill-or-kill): executes in full at once or is cancelled; GTC (good-till-cancelled): waits in the book until cancelled
- **Stop** order: becomes a market order once a trigger price trades
  - **iceberg**: displays only part of its size, the rest waits hidden
- A marketable limit order consumes liquidity like a market order, with a price cap
  - example: a buy limit at or above the best ask
- **Zero-intelligence** simulation Farmer, Patelli & Zovko (2005)
  - Rates: limit orders $\alpha = 0.1$ per price level, uniform up to 40 ticks from the opposite best quote; market orders $\mu = 0.5$ per side; cancellations $\delta = 0.02$ per standing order; 300 000 events, unit orders
  - Such a model already predicts much of the cross-section of spreads on the London Stock Exchange

Back

Statistical Properties of the Roll Estimator

- Delta method Harris (1990), a first-order Taylor expansion of $\hat{s} = 2\sqrt{-\hat{\gamma}_1}$ around γ_1:
  $\mathrm{se}(\hat{s}) \approx \sqrt{(\gamma_0^2 + \gamma_1^2)/T}\,/\,c$, for signs $q_t = \pm 1$
  - Bartlett's Gaussian approximation has $3\gamma_1^2$ instead of γ_1^2
  - Large when c is small: the smaller the spread, the less precisely it is estimated
  - $\hat{s}$ is also biased (Jensen's inequality, truncation at 0)
- With daily data the bounce (a few bp) is buried under the daily volatility
  - 103 bp for SPY over the last two years
  - hence estimators based on the daily high and low

Back

A Structural Model of Trading Costs: MRR (1/2)

- Roll assumes i.i.d. trade signs; real signs are persistent (a buy tends to follow a buy)
  - First-order Markov signs: $\text{Corr}(q_t, q_{t-k}) = \rho^k$, $\rho \in [0, 1)$ the persistence; Corr: correlation
  - Signs independent of ε and no information: $\text{Cov}(\Delta p_t, \Delta p_{t-1}) = -c^2(1-\rho)^2$, closer to 0 than $-c^2$, so Roll underestimates c
- MRR model Madhavan, Richardson & Roomans (1997): the trade price is the efficient price plus a transitory cost
  - the efficient price moves with the **surprise** in the order flow
  - $p_t = m_t + \phi\, q_t + \xi_t, \qquad m_t = m_{t-1} + \theta\,(q_t - \mathbb{E}[q_t \mid q_{t-1}]) + \varepsilon_t$
  - $\theta \geq 0$: permanent (adverse-selection) impact of the surprise $q_t - \mathbb{E}[q_t \mid q_{t-1}] = q_t - \rho q_{t-1}$; $\phi \geq 0$: transitory cost (processing, inventory)
  - q_t a first-order Markov chain; public news ε_t and rounding noise ξ_t i.i.d., mean zero, independent of each other and of q

Back

A Structural Model of Trading Costs: MRR (2/2)

- Estimation by GMM (Generalized Method of Moments): choose (θ, ϕ, ρ, a) so that the sample versions of four moments are zero
  - Residual: $u_t = \Delta p_t - (\phi + \theta) q_t + (\phi + \rho\theta) q_{t-1}$, the price change not explained by the trades; a: its drift (constant)
  - $\mathbb{E}[q_t q_{t-1} - \rho q_{t-1}^2] = 0, \quad \mathbb{E}[u_t - a] = 0, \quad \mathbb{E}[(u_t - a) q_t] = 0, \quad \mathbb{E}[(u_t - a) q_{t-1}] = 0$
  - First: defines ρ as the autocorrelation of signs; the others: the residual has mean a and is uncorrelated with current and past signs
- Outputs
  - Implied effective spread $2(\phi + \theta)$; adverse-selection share $\theta/(\phi + \theta) \in [0, 1]$: 0 means pure processing and inventory costs, 1 pure information
- Needs signed trades (trade-and-quote data)
  - the public 5-minute bars do not identify q_t, so MRR is not estimated here

Back

Corwin–Schultz and EDGE in Practice

- Corwin–Schultz on daily data: the derivation assumes no price move overnight Corwin & Schultz (2012)
  - the paper shifts H_{t+1}, L_{t+1} down by $L_{t+1} - C_t$ if $L_{t+1} > C_t$
  - up by $C_t - H_{t+1}$ if $H_{t+1} < C_t$ (C_t the close)
- Within one session of 5-minute bars no overnight adjustment is needed
  - negative two-bar estimates are set to 0
- EDGE over windows: set each negative window estimate to zero, then average Ardia, Guidotti & Kroencke (2024)
  - the authors' recommendation
  - more negative estimates go with larger spreads

Back

The EDGE Moment Conditions

- Log prices of bar t: open o_t, high h_t, low l_t, close c_t
  - mid-range $\eta_t = (h_t + l_t)/2$ as for Abdi–Ranaldo Ardia, Guidotti & Kroencke (2024)
  - Five log-returns around the open and the previous close: $r_{1t} = \eta_t - o_t$, $r_{2t} = o_t - \eta_{t-1}$, $r_{3t} = \eta_t - c_{t-1}$, $r_{4t} = c_{t-1} - \eta_{t-1}$, $r_{5t} = o_t - c_{t-1}$
- Two moment series
  - their sample means are the estimates $\hat{e}_1$, $\hat{e}_2$ of S^2
  - their sample variances v_1, v_2
  - $x_{1t} = -\frac{4}{p_o}\, d_{1t}\, r_{2t} - \frac{4}{p_c}\, d_{3t}\, r_{4t}, \qquad x_{2t} = -\frac{4}{p_o}\, d_{1t}\, r_{5t} - \frac{4}{p_c}\, d_{5t}\, r_{4t}$
  - $d_{kt} = r_{kt} - \tau_t\, \bar{r}_k/\bar{\tau}$: r_{kt} centred on the bars that moved; $\tau_t = 0$ only for a flat bar ($h_t = l_t = c_{t-1}$)
  - p_o: share of bars with $o_t \neq h_t$ plus share with $o_t \neq l_t$; p_c: the same for c_{t-1} against h_{t-1}, l_{t-1}; both equal 2 when the open and close never touch an extreme
- Why it works: the open and the previous close are trades at the bid or the ask
  - their distances to the surrounding mid-ranges have opposite signs on average, as in Roll
  - $-4/p$ turns this negative product into S^2

Back

Estimating PIN, and Why It Is Fragile (1/2)

- EKOP model Easley, Kiefer, O'Hara & Paperman (1996): each day an information event occurs with probability α
  - EKOP: initials of Easley, Kiefer, O'Hara and Paperman
  - bad news with probability δ
  - given the day's information state, daily buy and sell counts B and S are independent Poisson
  - Rates (expected orders per day): uninformed ε_b, ε_s; informed μ added to the buy side on good-news days, to the sell side on bad-news days
- Likelihood of one day: a mixture over the three states (no news, bad news, good news)
  - $L = (1-\alpha)\,\pi(B;\varepsilon_b)\pi(S;\varepsilon_s) + \alpha\delta\,\pi(B;\varepsilon_b)\pi(S;\mu+\varepsilon_s) + \alpha(1-\delta)\,\pi(B;\mu+\varepsilon_b)\pi(S;\varepsilon_s)$
  - $\pi(n;\lambda) = e^{-\lambda}\lambda^n/n!$: Poisson probability of n orders at rate λ (not Kyle's λ); each term: probability of the state times that of the observed (B,S)
  - The sample likelihood is the product of L over days, maximised in $(\alpha,\delta,\mu,\varepsilon_b,\varepsilon_s)$

Back

Estimating PIN, and Why It Is Fragile (2/2)

- From the estimated parameters: $\mathrm{PIN} = \alpha\mu/(\alpha\mu + \varepsilon_b + \varepsilon_s)$
  - expected informed orders over expected total orders
- Numerics: for liquid stocks ($B \approx 10^4$) $e^{-\varepsilon}\varepsilon^B/B!$ under- or overflows
  - factor out common terms and use log-sum-exp Lin & Ke (2011)
  - Without it, estimates pile up at boundary values
  - The likelihood is also multimodal: use many starting values (seminar A9–A10)

Back

The Hasbrouck Trade–Quote VAR

- VAR in the mid-quote return r_t and the signed trade x_t Hasbrouck (1991)
  - $r_t = \sum_{i\geq 1} a_i r_{t-i} + \sum_{i\geq 0} b_i x_{t-i} + \varepsilon_{1t}, \qquad x_t = \sum_{i\geq 1} c_i r_{t-i} + \sum_{i\geq 1} d_i x_{t-i} + \varepsilon_{2t}$
  - t: trade time; $x_t = q_t$ (± 1) or signed volume; a_i, b_i, c_i, d_i: lag coefficients; ε_{1t}: quote shock (public news), ε_{2t}: trade shock (the unexpected part of the order flow)
- Identification: a recursive (Cholesky) ordering
  - the trade comes first within the interval
  - b_0 enters, the reverse contemporaneous term does not
- **Permanent impact**: the long-run cumulative response of the quote to a trade shock
  - $\lim_{h\to\infty} \sum_{k=0}^{h} \partial r_{t+k} / \partial \varepsilon_{2t}$
  - i.e. the information content of trades
- Variance decomposition of the efficient price: the share explained by trades
  - it generalises the adverse-selection share of MRR
- Like MRR, it needs signed trades; in Kyle's model the permanent impact is exactly λ

Back

Order Flow Is Long-Memory and Self-Exciting (1/2)

- Signs of market orders are autocorrelated with a power-law decay: metaorders are split Bouchaud, Gefen, Potters & Wyart (2004)
  - $\mathrm{Corr}(\varepsilon_t, \varepsilon_{t+\ell}) \sim \ell^{-g}, \qquad 0 < g < 1$
  - Here $\varepsilon_t = \pm 1$: sign of the t-th market order; ℓ: lag in number of orders; $\sim$: behaves like, for large ℓ
  - $g < 1$: the correlations decay so slowly that their sum diverges (long memory)
- Yet prices are close to a random walk: the propagator model reconciles the two
  - $p_t = \sum_{s<t} G(t-s)\,\varepsilon_s + \text{noise}$
  - $G(\tau)$: the impact of an order τ steps later; a decaying kernel G offsets the persistent signs

Back

Order Flow Is Long-Memory and Self-Exciting (2/2)

- Arrival times cluster: a Hawkes process has an intensity that jumps after each event Bacry, Mastromatteo & Muzy (2015)
  - $\lambda(t) = \mu + \int_{-\infty}^{t} \phi(t-s)\, dN_s = \mu + \sum_{t_i < t} \phi(t - t_i)$
  - $\lambda(t)$ (not Kyle's λ): expected number of events per unit time at t; $\mu > 0$: baseline rate; N_s: number of events up to s, occurring at times t_i
  - $\phi \geq 0$: excitation kernel, how much a past event raises today's rate
- Estimated by maximum likelihood; branching ratio $\|\phi\|_1 = \int_0^\infty \phi(\tau)\, d\tau$
  - The share of events triggered by earlier events (endogeneity); stationary if < 1; close to 1: the market mostly reacts to itself
- Links:
  - durations between exceedances of the value at risk (VaR, not the VAR above) and their clustering (Chapter 8)
  - intraday volatility (Chapter 9)
  - the concave impact (section *Price Impact and Optimal Execution*)

Back

Deriving the Almgren–Chriss Trajectory

- Continuous time, holdings $x(t)$, trading rate $-\dot{x}$
  - temporary impact cost $\eta \int_0^T \dot{x}^2\, dt$
  - price-risk variance $\sigma^2 \int_0^T x^2\, dt$
  - Minimise $\int_0^T (\eta\, \dot{x}^2 + \lambda\sigma^2 x^2)\, dt$ subject to $x(0) = X$, $x(T) = 0$ (the permanent impact $\frac{1}{2}\gamma X^2$ does not depend on the path)
- Euler–Lagrange equation for the integrand $F(x, \dot{x}) = \eta\dot{x}^2 + \lambda\sigma^2 x^2$
  - the first-order condition of a calculus-of-variations problem
  - $\frac{d}{dt}\, \partial F/\partial \dot{x} = \partial F/\partial x \Rightarrow \ddot{x} = \kappa^2 x$, $\kappa^2 = \lambda\sigma^2/\eta$
- General solution $Ae^{\kappa t} + Be^{-\kappa t}$, with constants A, B
  - the boundary conditions give $x(t) = X \sinh(\kappa(T - t))/\sinh(\kappa T)$
- In discrete time κ^2 becomes $\frac{2}{\tau^2}(\cosh \kappa\tau - 1) = \lambda\sigma^2/\tilde{\eta}$
  - section *Price Impact and Optimal Execution*
  - a resilient book with transient impact changes the shape Obizhaeva & Wang (2013)

Back

Almgren–Chriss, the Numerical Table

Strategy (5 daily steps)	Expected cost	Standard deviation
Even selling, $\lambda = 0$	0.60 million USD	1.04 million USD
$\lambda = 2 \cdot 10^{-6}$	1.08	0.45
$\lambda = 2 \cdot 10^{-5}$	2.05	0.10
$\lambda = 10^{-4}$	2.38	0.02

- With $\lambda = 2 \cdot 10^{-6}$, 428 599 shares remain after the first day
- Same parameters as the main example; each row solves the discrete-time κ equation for its λ

Back

Optimal Market Making: Avellaneda–Stoikov (1/2)

- A dealer quotes a bid and an ask around the mid-price until the horizon T Avellaneda & Stoikov (2008)
  - s_t: mid-price, a Brownian motion with volatility σ; q: the dealer's inventory (shares held, negative if short)
  - He maximises expected exponential (CARA, constant absolute risk aversion) utility of terminal wealth, risk aversion $\gamma > 0$ (not the Almgren–Chriss permanent impact γ)
- Order flow: a quote at distance δ from the mid is hit by market orders with intensity
  - $\Lambda(\delta) = Ae^{-k\delta}$
  - $A > 0$: rate of execution at the mid; $k > 0$: how fast the execution rate falls as the quote moves away from the mid

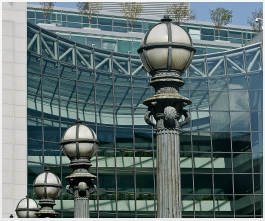
Back

Optimal Market Making: Avellaneda–Stoikov (2/2)

- Reservation (indifference) price: the value of the asset to this dealer, given his inventory
  - $r = s - q\gamma\sigma^2(T-t)$
  - $T-t$: time left; a long dealer ($q > 0$) values the asset below the mid, and more so with more risk aversion, volatility or time left
- Approximate optimal total spread around r
  - from an expansion of the value function, not the exact solution
  - $\delta^a + \delta^b = \gamma\sigma^2(T-t) + \frac{2}{\gamma}\ln\left(1+\frac{\gamma}{k}\right)$
  - δ^a, δ^b: distances of the ask and the bid from r; first term: inventory risk; second term: reflects how fast the probability of execution falls with distance from the mid
- Inventory skews both quotes
  - this formalises the inventory component of the spread (section *The Bid-Ask Spread*)

Back

Flash Crash: Details from the CFTC–SEC Report



Headquarters of the US Securities and Exchange Commission, Washington

Photo: David (dbking) (2008); CC BY 2.0; Wikimedia Commons

- The CFTC–SEC staff report CFTC & SEC (2010)
- The 2:32 p.m. sell order (main slide) was worth about USD 4.1 billion
  - The algorithm took 9% of the previous minute's volume "without regard to price or time"
- By 2:45:28 p.m. buy-side depth in the E-Mini had fallen below 1% of its morning level
- Over 20 000 trades in more than 300 securities were executed more than 60% away from their earlier prices
  - some at a penny
- Single-stock circuit breakers (main slide) pause trading for 5 minutes CFTC & SEC (2010)
  - trigger: a 10% move in 5 minutes

Back

The Bucharest Stock Exchange



Offices of the Bucharest Stock Exchange

Photo: Za.kurek (2014); CC BY-SA 4.0; Wikimedia Commons

- An electronic limit order book, as in the rest of this chapter BVB (2026)
- Emerging-market status (FTSE Russell) since September 2020
  - Banca Transilvania and Nuclearelectrica were the first stocks in FTSE emerging-market indices
- A thin market, where nonsynchronous trading matters (section *Illiquidity and Intraday Patterns*)
  - Herding and bubbles on the BVB: Pele, Mazurencu-Marinescu & Nijkamp (2013)
  - Bubbles and crashes are the subject of Chapter 17

Back

Case Study: Data and Method

- **Data**: London Stock Exchange message data for all FTSE 350 stocks
  - 17 August – 16 October 2015, 43 trading days
  - about 2.2 billion messages, obtained by the British regulator
  - Time stamps accurate to the microsecond (μs), taken where a message reaches the exchange's network, before any exchange processing
- **Race**: 2+ participants on the same stock, price and side
  - 1+ take and 1+ cancel attempts, or 2+ takes
  - at least one success and one failure
  - "Same time": within the **information horizon**, the first message's observed processing time plus the fastest reaction time seen (29 μs), capped at 500 μs
- **Race profit**: the signed gap between the race price and the mid-price 10 s later
  - times the displayed depth
- **Without races**, the cost of liquidity would fall by about 17%
  - the cut: non-race tax / non-race effective half-spread
  - averaged over dates
- **Extrapolation**: OLS of daily race profits, without a constant
  - on traded value V_t
  - and on V_t and 1-minute realised volatility σ_t times mean traded value $\bar{V}$

Back

Case Study: Further Results

- **Speed**: the winner beats the first loser by 5–10 μs (mode), 46 μs (median)
- **Concentration**: by volume, the top six firms take about 80% of the liquidity taken in races
  - but they provide only 42% of it
  - the 82% of the main slide counts races won
- **Prize per race**: per FTSE 100 race, 0.48 ticks (1.20 bp) per share and £1.95 in total
  - $R^2 = 0.811$ of daily race profits on traded value
  - The USD 4.8 billion a year of the main slide (global equities, 2018) ranges from 2.3 to 8.4 across race definitions
- **Cost of liquidity**: in the full sample, race price impact is 30.58% of all price impact
  - and 32.82% of the effective spread
  - FTSE 100 realised spread -1.93 bp in races, $+0.15$ bp outside
  - one-way spreads, half the measures of the section *The Bid-Ask Spread*
  - This decomposition shows that sniping belongs next to adverse selection in the cost of liquidity

Back

The Discovery Loop, Reusable Prompts (1/2)

- Tools: an AI assistant (a large language model, LLM)
  - such as Claude, ChatGPT, Gemini or GitHub Copilot
  - literature search: Semantic Scholar, Elicit, Google Scholar
- **1. Literature**: collect validation studies, then check every DOI on Crossref (the title must match)
  - `Find studies that validate Roll, Corwin-Schultz, Abdi-Ranaldo and EDGE against effective spreads from trades and quotes; report the benchmark, the correlation in levels and in changes, the markets, and the DOI.`
- **2. Competing hypotheses**, each with a cross-sectional prediction
  - `Give rival explanations for why Corwin-Schultz ranks BVB blue chips next to US large caps, and for each the predicted Spearman correlation of the estimate with volatility, with the Amihud ratio and with traded value.`
- **3. Data and code**: replicate the chapter's numbers first
  - `Write Python that loads the daily OHLC prices and volumes of the 25 assets of this chapter and computes Corwin-Schultz, with the overnight adjustment and negative two-day estimates set to zero.` Target medians: US 53, BVB 54, crypto 124 bp

Back

The Discovery Loop, Reusable Prompts (2/2)

- **4. Robustness**, pre-registered
  - `Write a pre-registered plan: Spearman and partial Spearman correlations, bootstrap over assets, leave-one-market-out (US, BVB, crypto), a second window 2022-2024, and EDGE and Abdi-Ranaldo in place of Corwin-Schultz.`
- **5. Critique**: ask the assistant to attack the benchmark itself
  - `Act as a referee: argue that the Amihud ratio is not a cost measure either (it scales with volume and with volatility); which benchmark would identify the spread, and what would it cost to obtain?`
- **6. Report**: every estimator and window, bootstrap intervals
  - the benchmark's limits, code and the AI log

Back

What the Researcher Must Check

- **Tick-size benchmark**: no **quoted** spread is below one tick (effective spreads can be)
  - compare every estimate with $0.01/P$, one tick over the price P (0.21 bp for SPY)
  - a plausibility check, not a validation
- **Volatility contamination**: test by partial correlation with the benchmark, controlling for volatility
  - and the 5-minute versus daily comparison in the section *The Bid-Ask Spread*
- **Truncation**: negative estimates set to zero bias the mean upwards
  - report the share truncated and a median
- **Units and venues**: lei converted at the BNR rate
  - crypto volume summed over exchanges, some with inflated volume
  - test: rerun on a subset of venues
- **Small, clustered cross-section**: 25 assets in 3 markets are not exchangeable
  - 3 clusters are too few for cluster-robust inference
  - report the asset bootstrap as conditional on independence
  - plus a within-market bootstrap and leave-one-market-out
- **Mistyped formulas**: check the constants of the estimator against the paper
  - run a unit test on simulated prices with a known spread

Back

Microstructure: Readings and Code

Advanced courses on quantinar.com

- Applied Computational Finance

Reading, code and data

- Textbooks: Hasbrouck (2007), Bouchaud, Bonart, Donier & Gould (2018), Harris (2003), Franke, Härdle & Hafner (2019)
- High-frequency Bitcoin data and daily risk forecasts: Pele & Mazurencu-Marinescu-Pele (2019)
- Quantlets: github.com/danpele/MFM/tree/main/Quantlets/Ch_10

Back

Acronyms used in this material (1/2)

- **AI**: Artificial Intelligence
- **AMM**: Automated Market Maker
- **AR**: AutoRegressive (model)
- **BET**: Bucharest Exchange Trading (index)
- **BET-TR**: BET Total Return (index)
- **BIC**: Bayesian Information Criterion
- **BNR**: Banca Națională a României — *the National Bank of Romania*
- **BVB**: Bursa de Valori București — *the Bucharest Stock Exchange*
- **CARA**: Constant Absolute Risk Aversion
- **CEX**: Centralised EXchange
- **CFTC**: Commodity Futures Trading Commission
- **CI**: Confidence Interval
- **DEX**: Decentralised EXchange
- **DOI**: Digital Object Identifier
- **DVOL**: Dollar VOLume (traded value, price $\times$ volume)
- **EDGE**: Efficient Discrete Generalized Estimator (of the bid-ask spread)
- **EKOP**: Easley–Kiefer–O'Hara–Paperman (model of informed trading)
- **ETF**: Exchange-Traded Fund
- **FOK**: Fill-Or-Kill (order)
- **FTSE**: Financial Times Stock Exchange (FTSE Russell, index provider)
- **GMM**: Generalized Method of Moments
- **GTC**: Good-Till-Cancelled (order)
- **HAC**: Heteroskedasticity and Autocorrelation Consistent (standard errors)
- **HFT**: High-Frequency Trading
- **ILLIQ**: ILLIQuidity ratio (Amihud)
- **IOC**: Immediate-Or-Cancel (order)

Acronyms used in this material (2/2)

- **LLM**: Large Language Model
- **LOB**: Limit Order Book
- **MRR**: Madhavan–Richardson–Roomans (structural model of trading costs)
- **NYSE**: New York Stock Exchange
- **OFI**: Order-Flow Imbalance
- **OHLC**: Open, High, Low, Close
- **OLS**: Ordinary Least Squares
- **PIN**: Probability of INformed trading
- **SEC**: Securities and Exchange Commission
- **TR**: Total Return
- **TWAP**: Time-Weighted Average Price
- **US**: United States
- **USD**: US dollar, international currency code
- **UTC**: Coordinated Universal Time (Temps Universel Coordonné)
- **VAR**: Vector AutoRegression
- **VaR**: Value at Risk
- **VECM**: Vector Error-Correction Model
- **VIX**: CBOE Volatility IndeX
- **VPIN**: Volume-synchronised Probability of INformed trading
- **VWAP**: Volume-Weighted Average Price
- **XRP**: XRP (the native token of the XRP Ledger)

References (1/5)

- Abdi, F., Ranaldo, A. (2017). A simple estimation of bid-ask spreads from daily close, high, and low prices. *Review of Financial Studies*, 30(12), 4437–4480.
- Acharya, V.V., Pedersen, L.H. (2005). Asset pricing with liquidity risk. *Journal of Financial Economics*, 77(2), 375–410.
- Admati, A.R., Pfleiderer, P. (1988). A theory of intraday patterns: volume and price variability. *Review of Financial Studies*, 1(1), 3–40.
- Agakishiev, I., Härdle, W.K., Pele, D.T. (2025). Cryptocurrency liquidity forecasting. In C.Y.-H. Chen, W.K. Härdle, H.H.-S. Lu (Eds.), *Handbook of Blockchain Analytics*, 41–60. Springer Handbooks of Computational Statistics. Springer.
- Almgren, R., Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3(2), 5–39.
- Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1), 31–56.
- Amihud, Y., Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2), 223–249.
- Andersen, T.G., Bondarenko, O. (2014). VPIN and the flash crash. *Journal of Financial Markets*, 17, 1–46.
- Aquilina, M., Budish, E., O'Neill, P. (2022). Quantifying the high-frequency trading "arms race". *Quarterly Journal of Economics*, 137(1), 493–564.
- Ardia, D., Guidotti, E., Kroencke, T.A. (2024). Efficient estimation of bid–ask spreads from open, high, low, and close prices. *Journal of Financial Economics*, 161, 103916.
- Avellaneda, M., Stoikov, S. (2008). High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), 217–224.

References (2/5)

- Bacry, E., Mastromatteo, I., Muzy, J.-F. (2015). Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(1), 1550005.
- Bagehot, W. (1971). The only game in town. *Financial Analysts Journal*, 27(2), 12–14.
- Baltussen, G., Da, Z., Lammers, S., Martens, M. (2021). Hedging demand and market intraday momentum. *Journal of Financial Economics*, 142(1), 377–403.
- Bouchaud, J.-P., Bonart, J., Donier, J., Gould, M. (2018). Trades, Quotes and Prices: Financial Markets Under the Microscope. Cambridge University Press.
- Bouchaud, J.-P., Gefen, Y., Potters, M., Wyart, M. (2004). Fluctuations and response in financial markets: the subtle nature of 'random' price changes. *Quantitative Finance*, 4(2), 176–190.
- Brauneis, A., Mestel, R., Riordan, R., Theissen, E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124, 106041.
- Budish, E., Cramton, P., Shim, J. (2015). The high-frequency trading arms race: frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4), 1547–1621.
- BVB (2026). History of the Bucharest Stock Exchange. Bursa de Valori București.
- CFTC, SEC (2010). Findings regarding the market events of May 6, 2010. Report of the staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.
- Cont, R., Kukanov, A., Stoikov, S. (2014). The price impact of order book events. *Journal of Financial Econometrics*, 12(1), 47–88.
- Corwin, S.A., Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67(2), 719–760.
- Easley, D., Kiefer, N.M., O'Hara, M., Paperman, J.B. (1996). Liquidity, information, and infrequently traded stocks. *Journal of Finance*, 51(4), 1405–1436.

References (3/5)

- Easley, D., López de Prado, M.M., O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies*, 25(5), 1457–1493.
- Farmer, J.D., Patelli, P., Zovko, I.I. (2005). The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102(6), 2254–2259.
- Franke, J., Härdle, W.K., Hafner, C.M. (2019). Statistics of Financial Markets: An Introduction, 5th ed. Springer.
- Gao, L., Han, Y., Li, S.Z., Zhou, G. (2018). Market intraday momentum. *Journal of Financial Economics*, 129(2), 394–414.
- Glosten, L.R., Milgrom, P.R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100.
- Gonzalo, J., Granger, C.W.J. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, 13(1), 27–35.
- Gould, M.D., Porter, M.A., Williams, S., McDonald, M., Fenn, D.J., Howison, S.D. (2013). Limit order books. *Quantitative Finance*, 13(11), 1709–1742.
- Harris, L. (1990). Statistical properties of the Roll serial covariance bid/ask spread estimator. *Journal of Finance*, 45(2), 579–590.
- Harris, L. (2003). Trading and Exchanges: Market Microstructure for Practitioners. Oxford University Press.
- Hasbrouck, J. (1991). Measuring the information content of stock trades. *Journal of Finance*, 46(1), 179–207.
- Hasbrouck, J. (1995). One security, many markets: determining the contributions to price discovery. *Journal of Finance*, 50(4), 1175–1199.
- Hasbrouck, J. (2007). Empirical Market Microstructure. Oxford University Press.

References (4/5)

- Hasbrouck, J. (2009). Trading costs and returns for U.S. equities: estimating effective costs from daily data. *Journal of Finance*, 64(3), 1445–1477.
- Hasbrouck, J., Saar, G. (2013). Low-latency trading. *Journal of Financial Markets*, 16(4), 646–679.
- Huang, R.D., Stoll, H.R. (1997). The components of the bid-ask spread: a general approach. *Review of Financial Studies*, 10(4), 995–1034.
- Jahan-Parvar, M.R., Zikes, F. (2023). When do low-frequency measures really measure effective spreads? Evidence from equity and foreign exchange markets. *Review of Financial Studies*, 36(10), 4190–4232.
- Johansen, S. (1991). Estimation and hypothesis testing of cointegration vectors in Gaussian vector autoregressive models. *Econometrica*, 59(6), 1551–1580.
- Jovanovic, B., Menkveld, A.J. (2010). Middlemen in limit order markets. SSRN Working Paper 1624329.
- Kirilenko, A., Kyle, A.S., Samadi, M., Tuzun, T. (2017). The Flash Crash: high-frequency trading in an electronic market. *Journal of Finance*, 72(3), 967–998.
- Kyle, A.S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.
- Lee, C.M.C., Ready, M.J. (1991). Inferring trade direction from intraday data. *Journal of Finance*, 46(2), 733–746.
- Lin, H.-W.W., Ke, W.-C. (2011). A computing bias in estimating the probability of informed trading. *Journal of Financial Markets*, 14(4), 625–640.
- Lo, A.W., MacKinlay, A.C. (1990). An econometric analysis of nonsynchronous trading. *Journal of Econometrics*, 45(1–2), 181–211.

References (5/5)

- Madhavan, A., Richardson, M., Roomans, M. (1997). Why do security prices change? A transaction-level analysis of NYSE stocks. *Review of Financial Studies*, 10(4), 1035–1064.
- Makarov, I., Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2), 293–319.
- Menkveld, A.J. (2013). High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4), 712–740.
- Obizhaeva, A.A., Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1), 1–32.
- O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*, 116(2), 257–270.
- Pástor, Ľ., Stambaugh, R.F. (2003). Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3), 642–685.
- Pele, D.T. (2025). Middlemen in limit-order markets. Quantinar course, Bucharest University of Economic Studies.
- Pele, D.T., Mazurencu-Marinescu, M., Nijkamp, P. (2013). Herding behaviour, bubbles and log periodic power laws in illiquid stock markets: a case study on the Bucharest Stock Exchange. SSRN Working Paper 2305945.
- Pele, D.T., Mazurencu-Marinescu-Pele, M. (2019). Using high-frequency entropy to forecast Bitcoin's daily value at risk. *Entropy*, 21(2), 102.
- Perold, A.F. (1988). The implementation shortfall. *Journal of Portfolio Management*, 14(3), 4–9.
- Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39(4), 1127–1139.
- Scholes, M., Williams, J. (1977). Estimating betas from nonsynchronous data. *Journal of Financial Economics*, 5(3), 309–327.
- Tóth, B., Lempérière, Y., Deremble, C., de Lataillade, J., Kockelkoren, J., Bouchaud, J.-P. (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2), 021006.
- Wood, R.A., McInish, T.H., Ord, J.K. (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance*, 40(3), 723–739.