

Deep Learning for Limit Order Book Forecasting

Silvia Bartolucci

Department of Computer Science, University College London

The logo for the 11th Annual Machine Learning in Finance Conference features a dark background with a colorful, abstract pattern of green, yellow, and blue dots and lines, resembling a neural network or data visualization. The text is overlaid on the left side of this pattern.

**11th Annual Machine Learning
in Finance Conference.**

September 25, 2025

Complexity in Finance

Financial markets are *complex adaptive systems*, shaped by:

- **A large number of interacting agents** — traders acting on diverse information, strategies, and timescales.
- **Emergent collective behavior** — where individual actions aggregate into patterns that are hard to predict.


To manage this complexity, modern exchanges rely on **Limit Order Books (LOBs)**, which serve as dynamic structures organizing buy and sell orders in real time. These mechanisms are central to ensuring **liquidity, price discovery, and market fairness**.



Finance and AI

Market Microstructure

- **Price formation mechanisms**
- **Asymmetric information**
- **Market anomalies** like price jumps and flash crashes
- **Transaction costs and optimal order execution**
- Understand **trading behaviour** and liquidity dynamics




Automated Learning Modelling

- **Linear Models & Regression:** Early work capturing trends and correlations with interpretable models.
- **Non-linear Deep Learning Models:** Includes CNNs, LSTMs, Transformers, and hybrid architectures.
- **Reinforcement Learning for Trading:** Focuses on decision-making and execution strategies.

Finance and AI


Market Microstructure



Automated Learning
Modelling

Microstructurally-informed deep learning models
for short-term price forecasting

Finance and AI



Kercheval & Zhang (2015)

Early ML (SVM, RF) for LOB, showed predictive information exists.

Tsantekidis et al. (2017)

Forecasting Stock Prices from the Limit Order Book using CNNs

Zhang, Zohren & Roberts (2019)


DeepLOB, with CNN+LSTM hybrid architecture

Lucchese, Pakkanen & Veraart (2024)

Order flow representation and dependence of prediction

Briola, Bartolucci & Aste (2025)

HLOB model: homological CNNs + network filtering



Sirignano (2016)

First deep network on large NASDAQ LOB data, feasibility of DL

Dixon (2018)

RNNs for LOB sequences, capturing temporal dependencies.

Sirignano & Cont (2019)

Universal Approximation of Limit Order Book Dynamics by Deep Neural Networks

Wu et al. (2021)

Robust mid-price-centered representation, reducing instability

Kolm, Turiel & Westray (2023-2024)

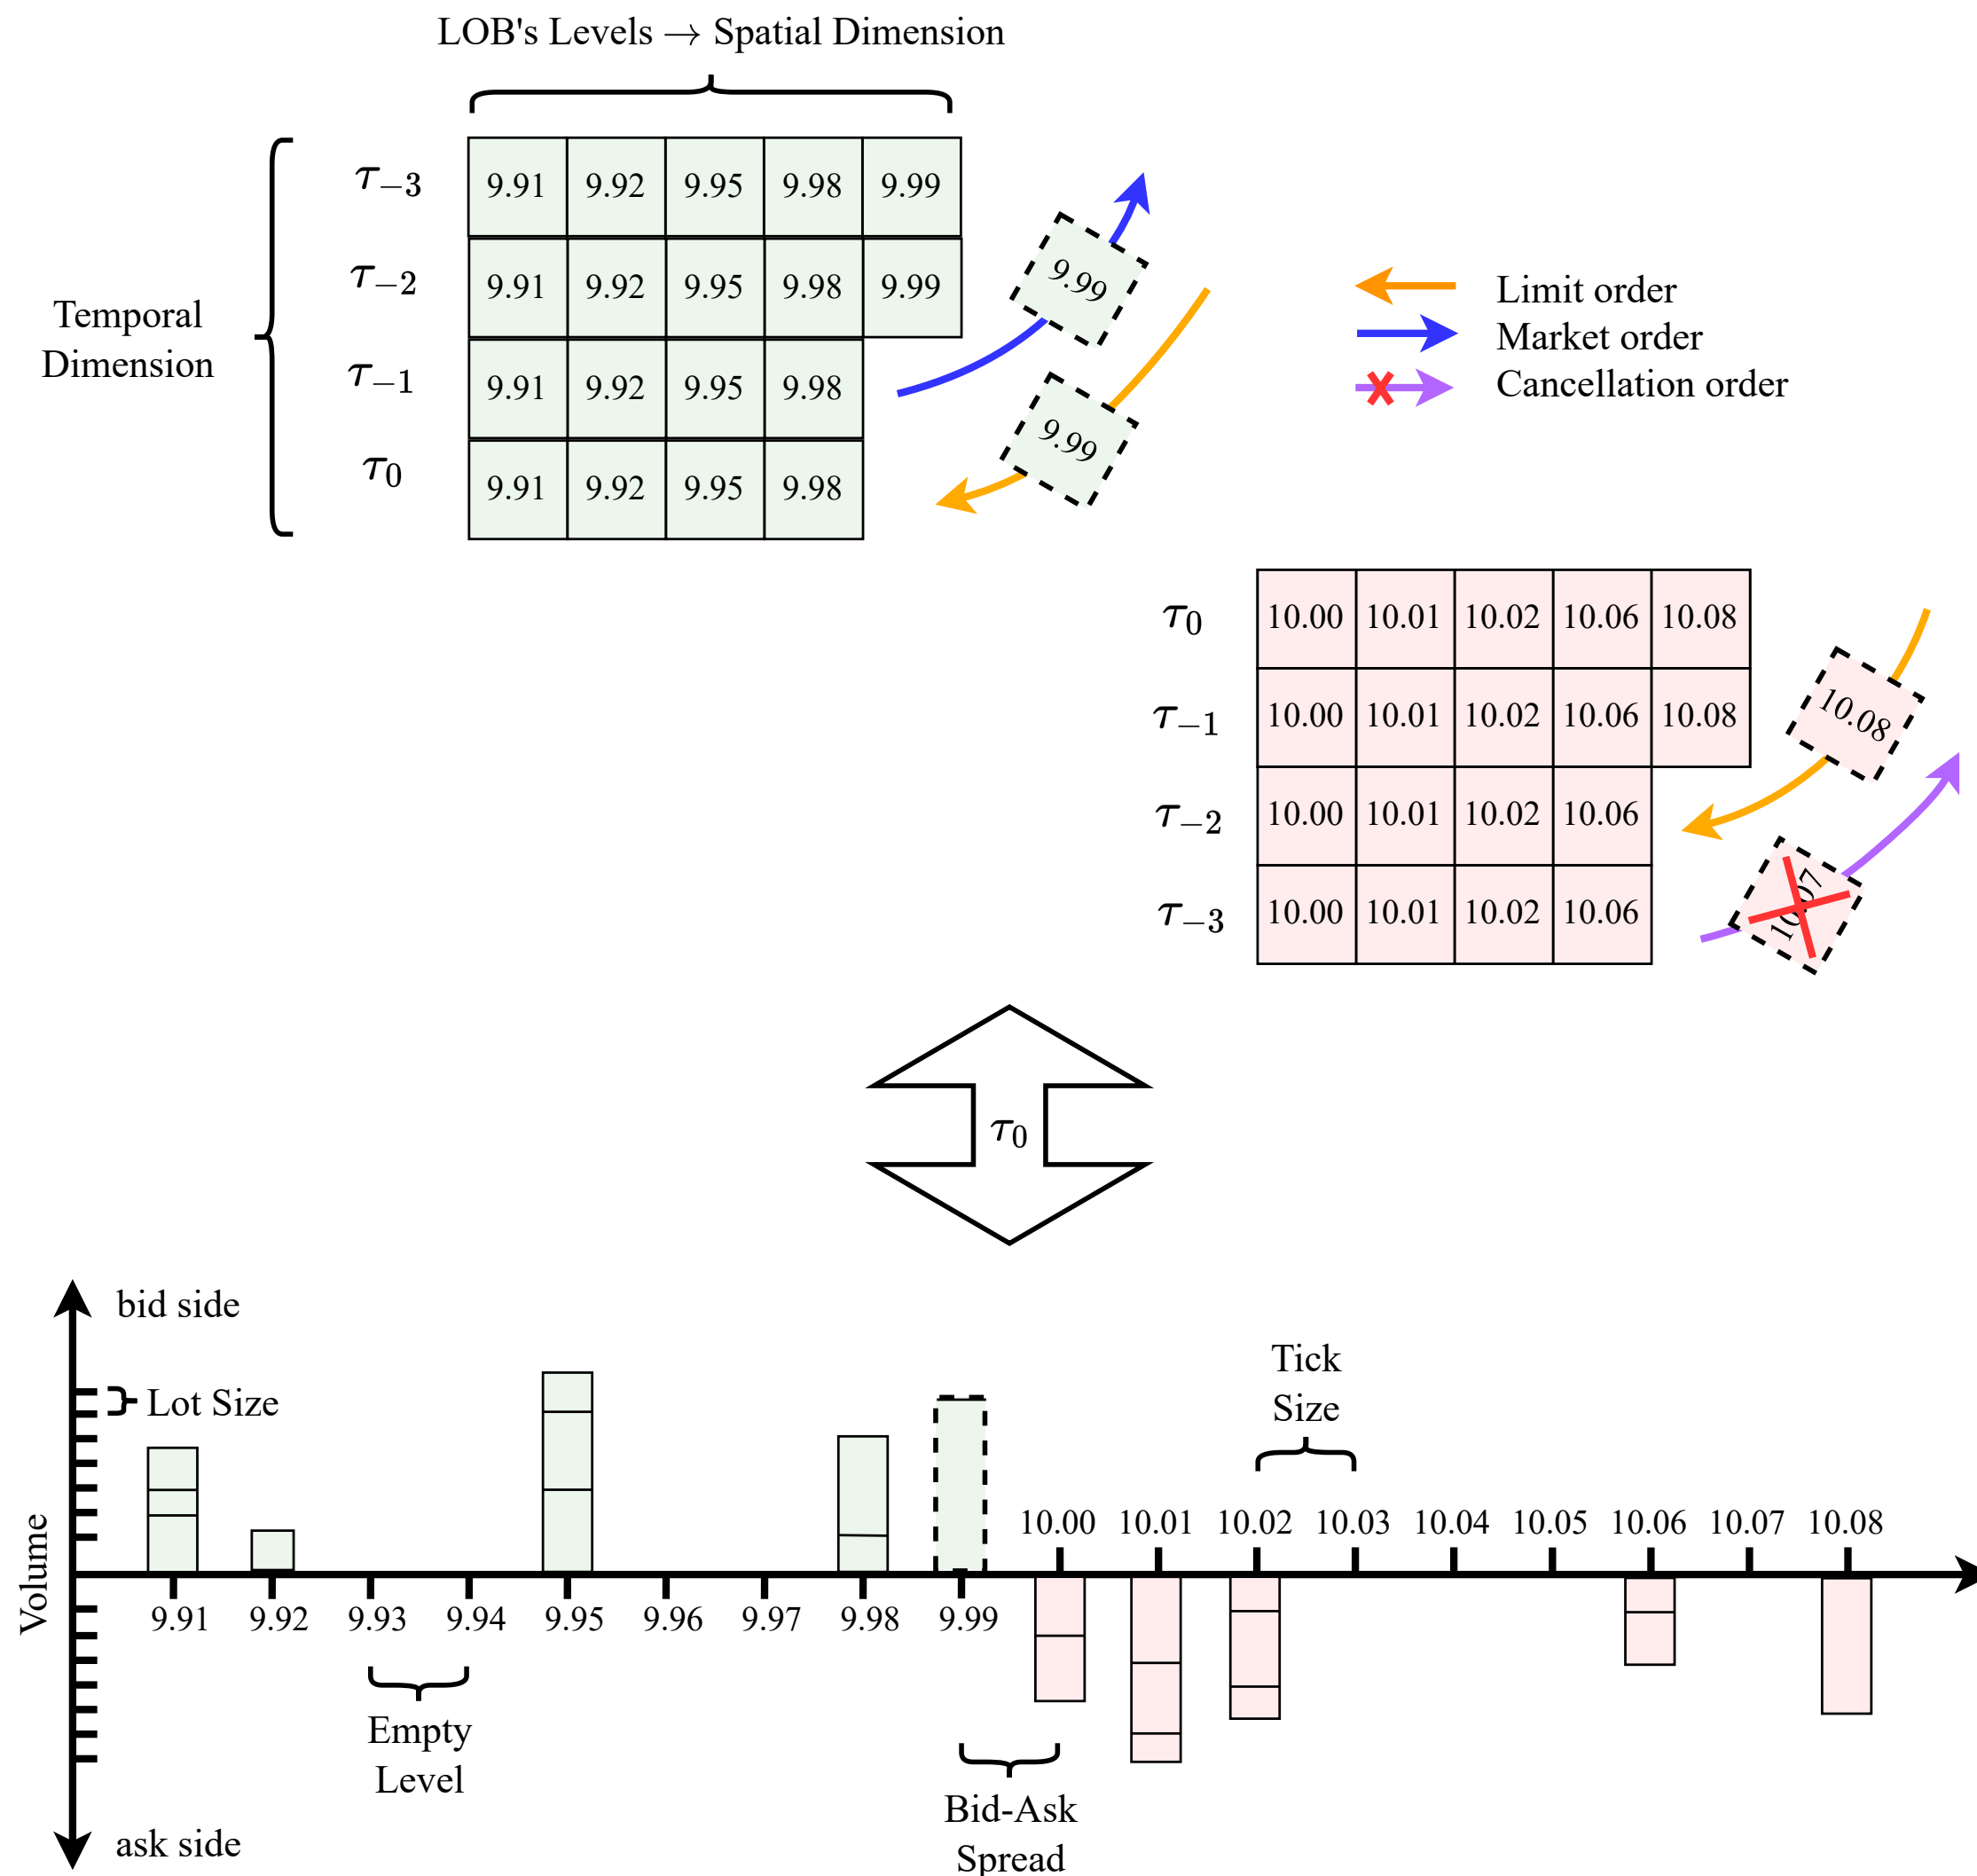
Deep Order Flow Imbalance, horizon-dependent alpha

Briola, Bartolucci & Aste (2024)

Short-term predictability depends on microstructure (tick size)

Limit order books in a nutshell

A LOB is the set of all limit orders for a given asset on a given platform at a given time



$$\mathbb{L}(\tau) = \{p_{\ell}^{\text{ask}}(\tau), v_{\ell}^{\text{ask}}(\tau), p_{\ell}^{\text{bid}}(\tau), v_{\ell}^{\text{bid}}(\tau)\}_{\ell=1}^L$$

Mid Price

$$m_{\tau} = \frac{p_1^{\text{ask}}(\tau) + p_1^{\text{bid}}(\tau)}{2}$$

Spread

$$\sigma_{\tau} = p_1^{\text{ask}}(\tau) - p_1^{\text{bid}}(\tau)$$

Tick size

$$\theta = 0.01\$ \text{ (NASDAQ)}$$

Prediction task

We consider 3 different prediction horizons $H\Delta\tau \in \{10, 50, 100\}$

$$(m_{\tau+\Delta\tau} - m_{\tau}) \leq -\theta \rightarrow -1 : \text{Down}$$

$$-\theta < (m_{\tau+\Delta\tau} - m_{\tau}) < +\theta \rightarrow 0 : \text{Stable}$$

$$(m_{\tau+\Delta\tau} - m_{\tau}) \geq +\theta \rightarrow 1 : \text{Up}$$

How do the microstructural properties of an asset and the choice of prediction horizon affect the accuracy of mid-price predictions?

Data

- 15 stocks, all listed on the NASDAQ exchange
- Different sectors (technology, healthcare, finance..)
- 3-years period from January 2017 to December 2019
- Data provider: LOBSTER (NASDAQ Total View ITCH feed)

CHTR, GOOG, GS, IBM,
MCD, NVDA

Small tick

$$\langle \sigma \rangle \geq 3\theta$$

AAPL, ABBV, PM

Medium tick

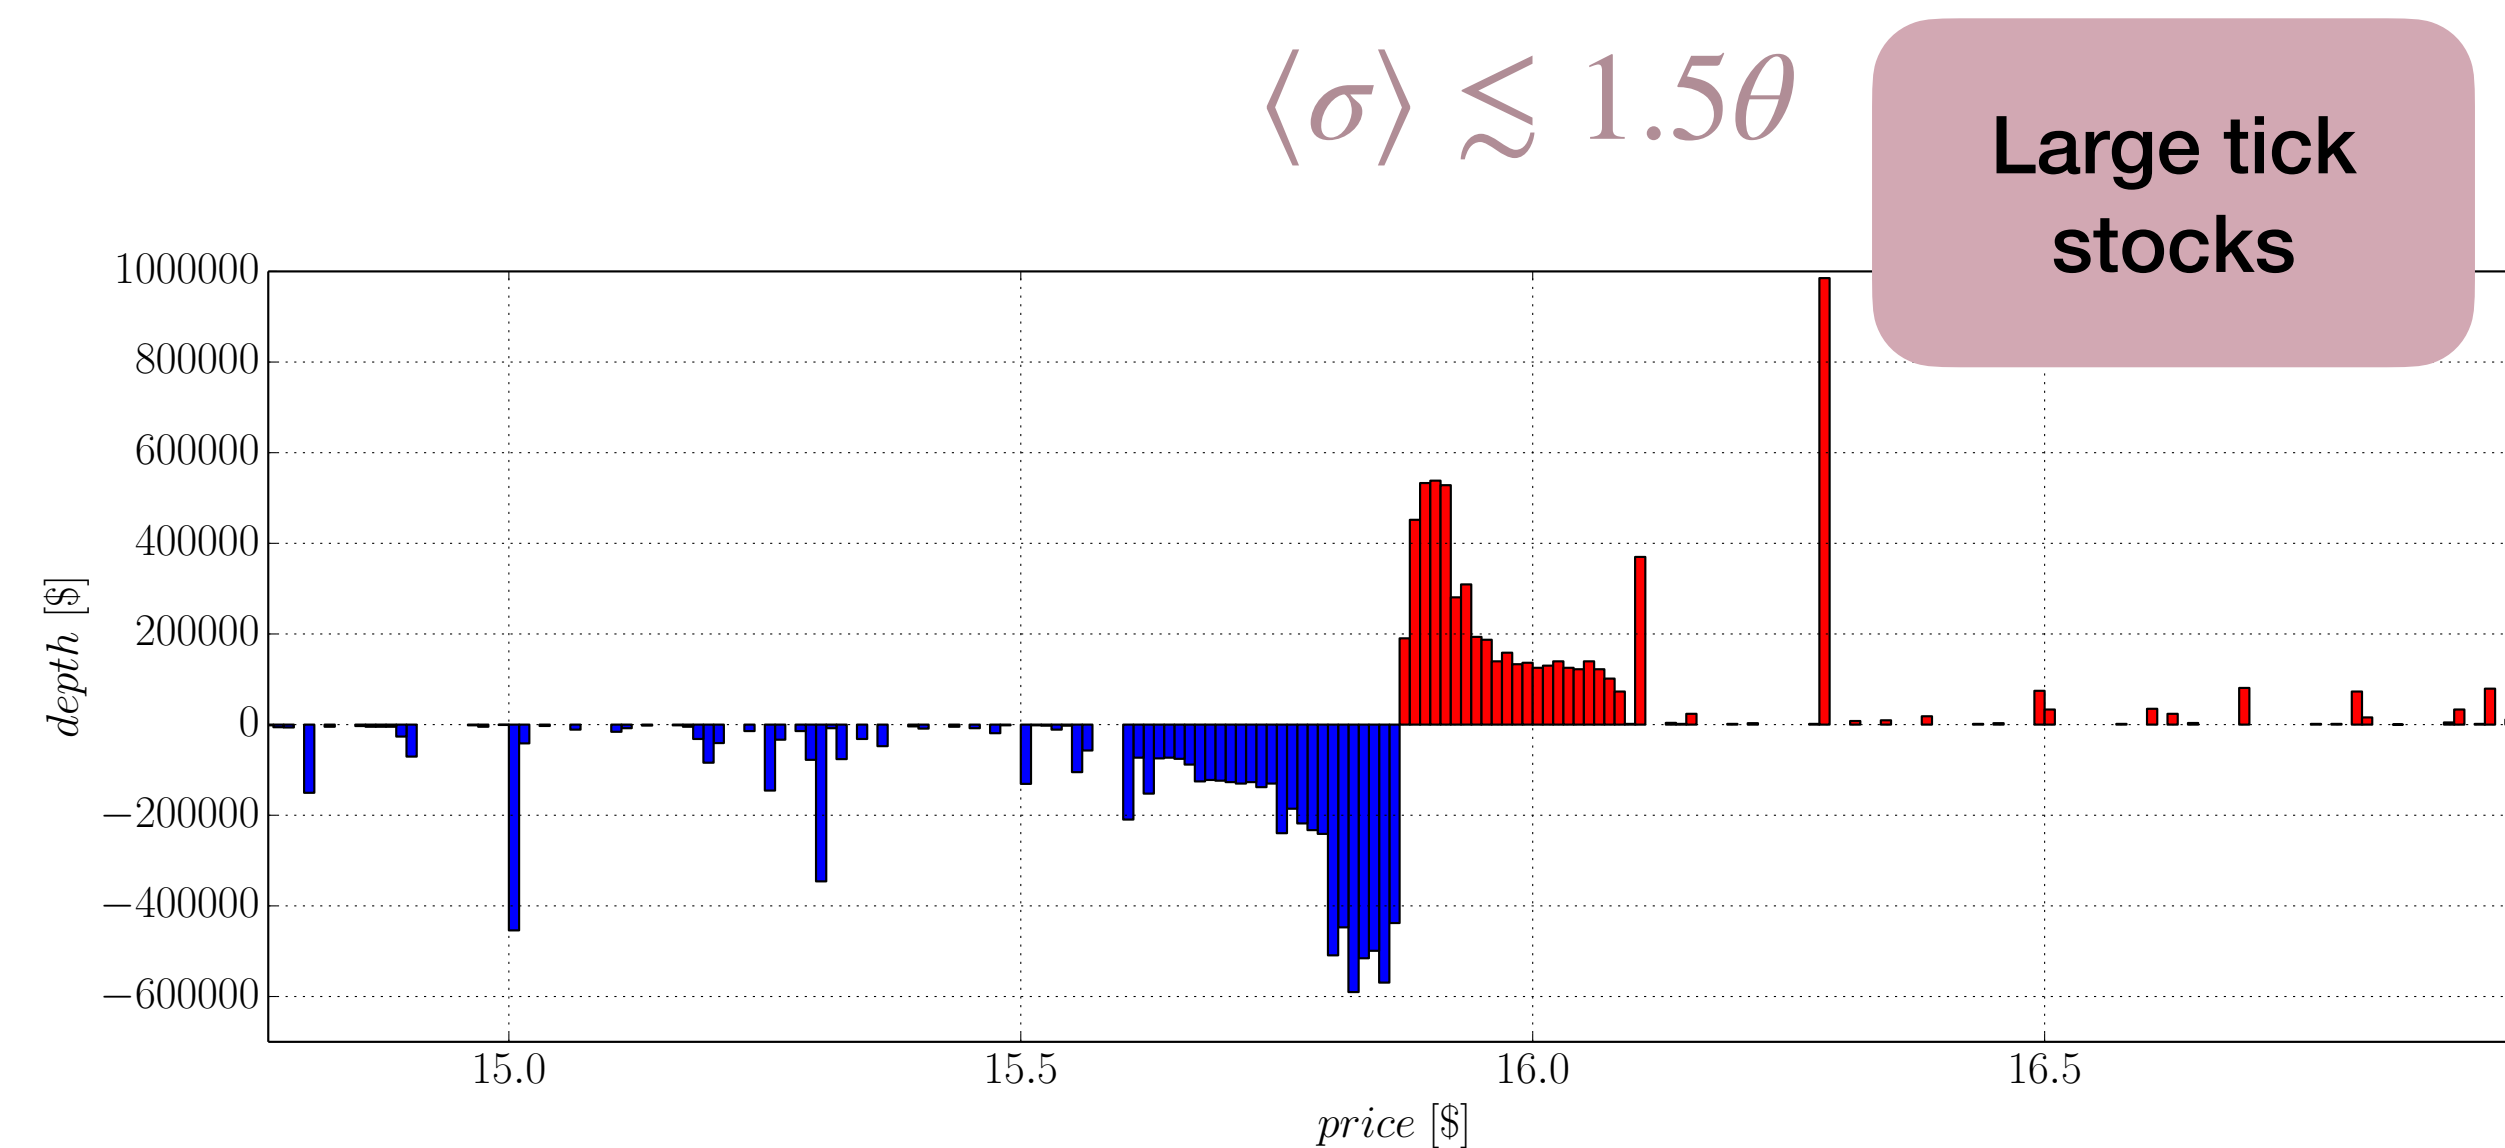
$$1.5\theta \lesssim \langle \sigma \rangle \lesssim 3\theta$$

BAC, CSCO, KO, ORCL,
PFE, VZ

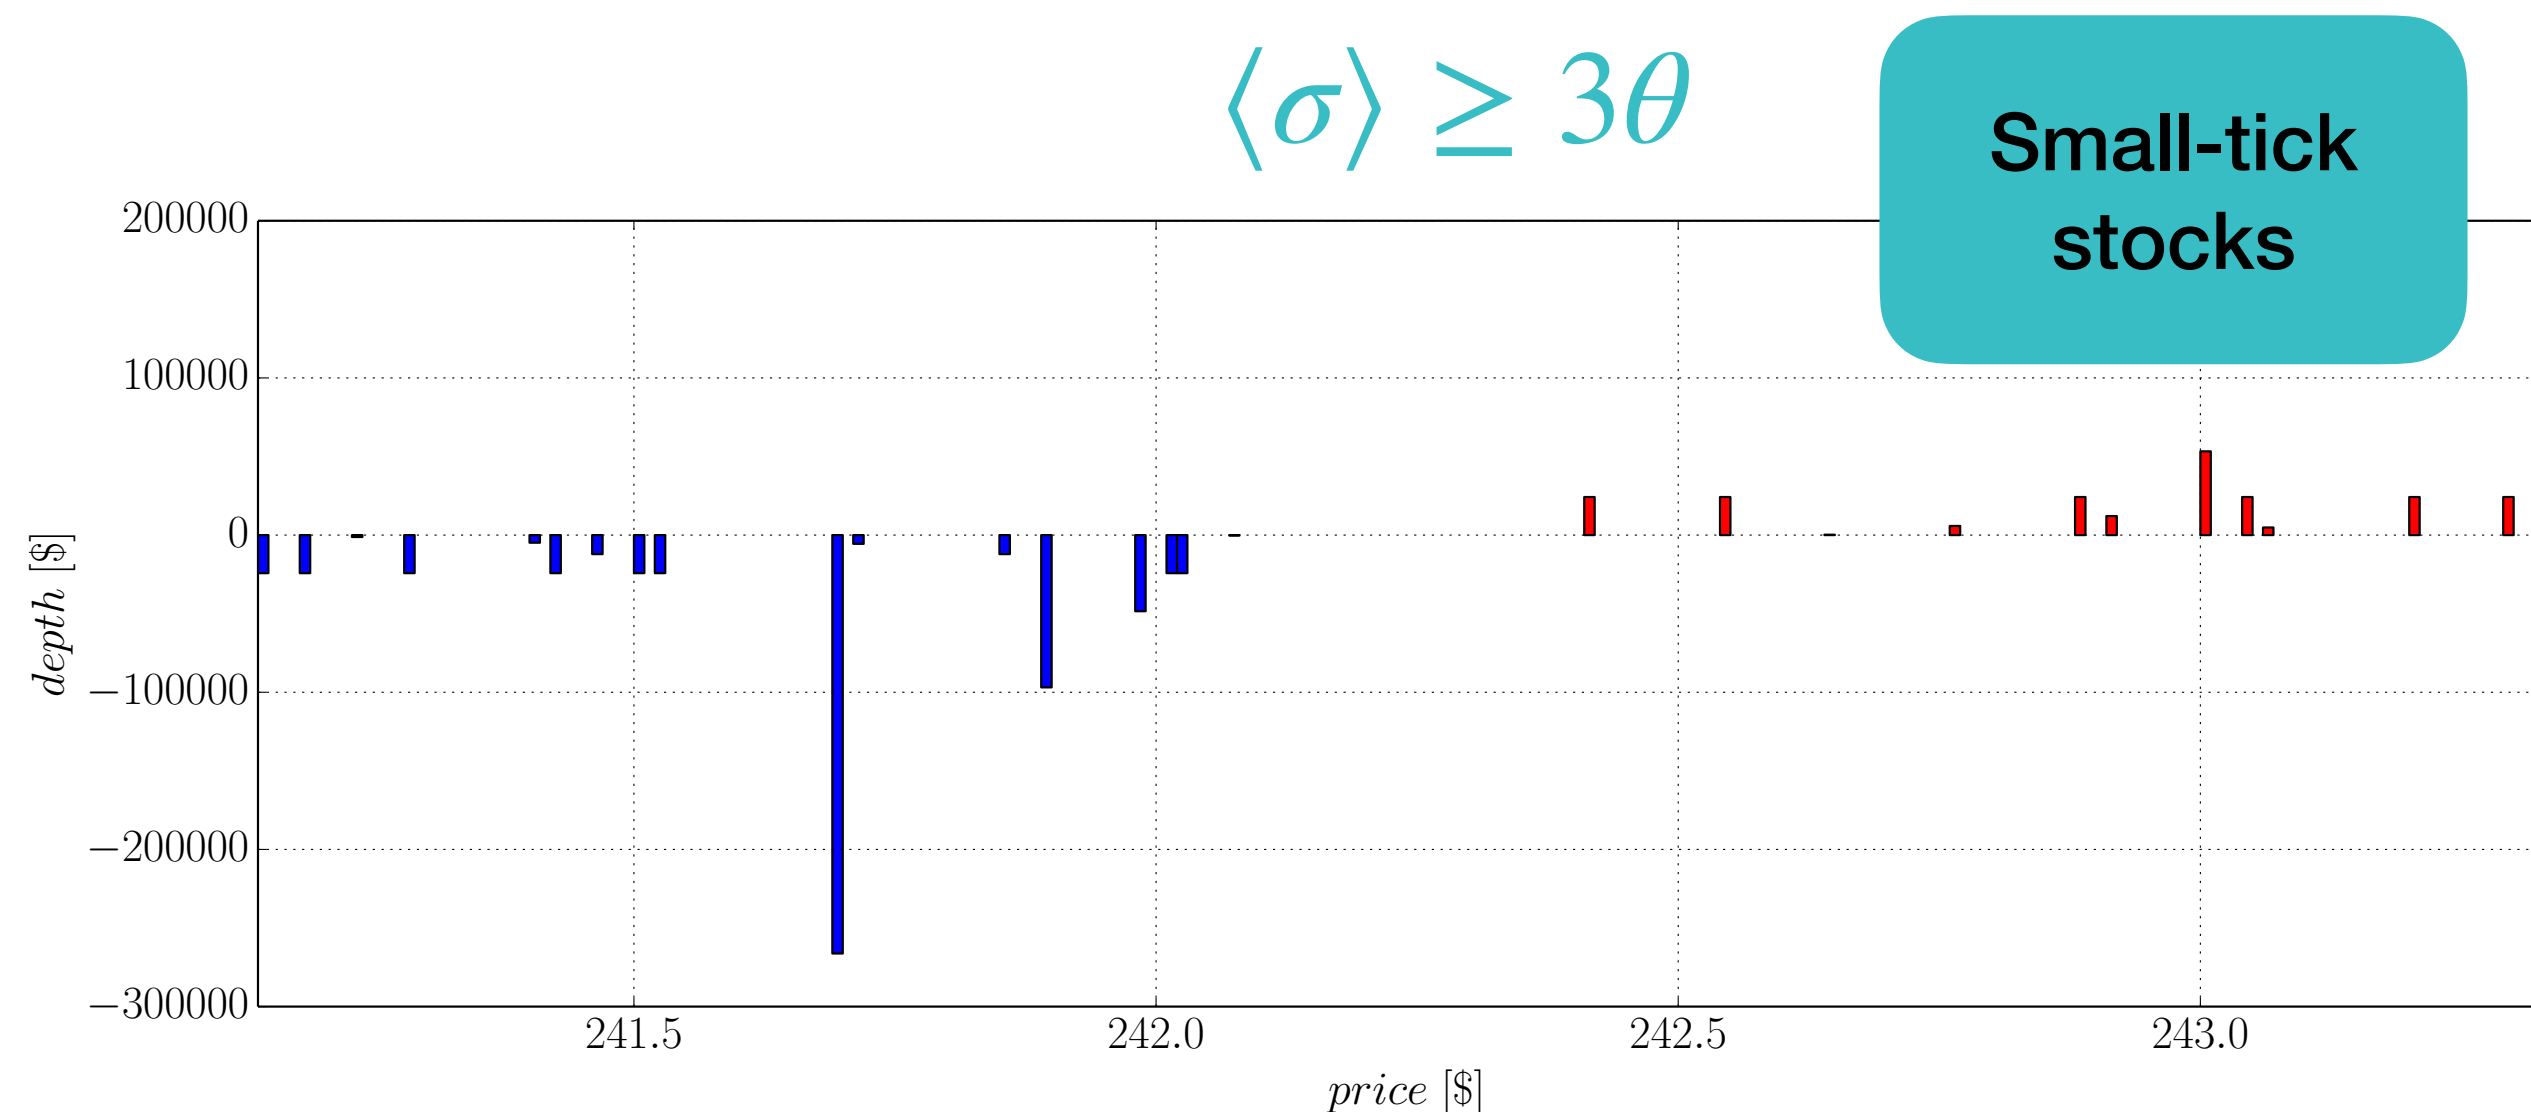
Large tick

$$\langle \sigma \rangle \lesssim 1.5\theta$$

Microstructural properties



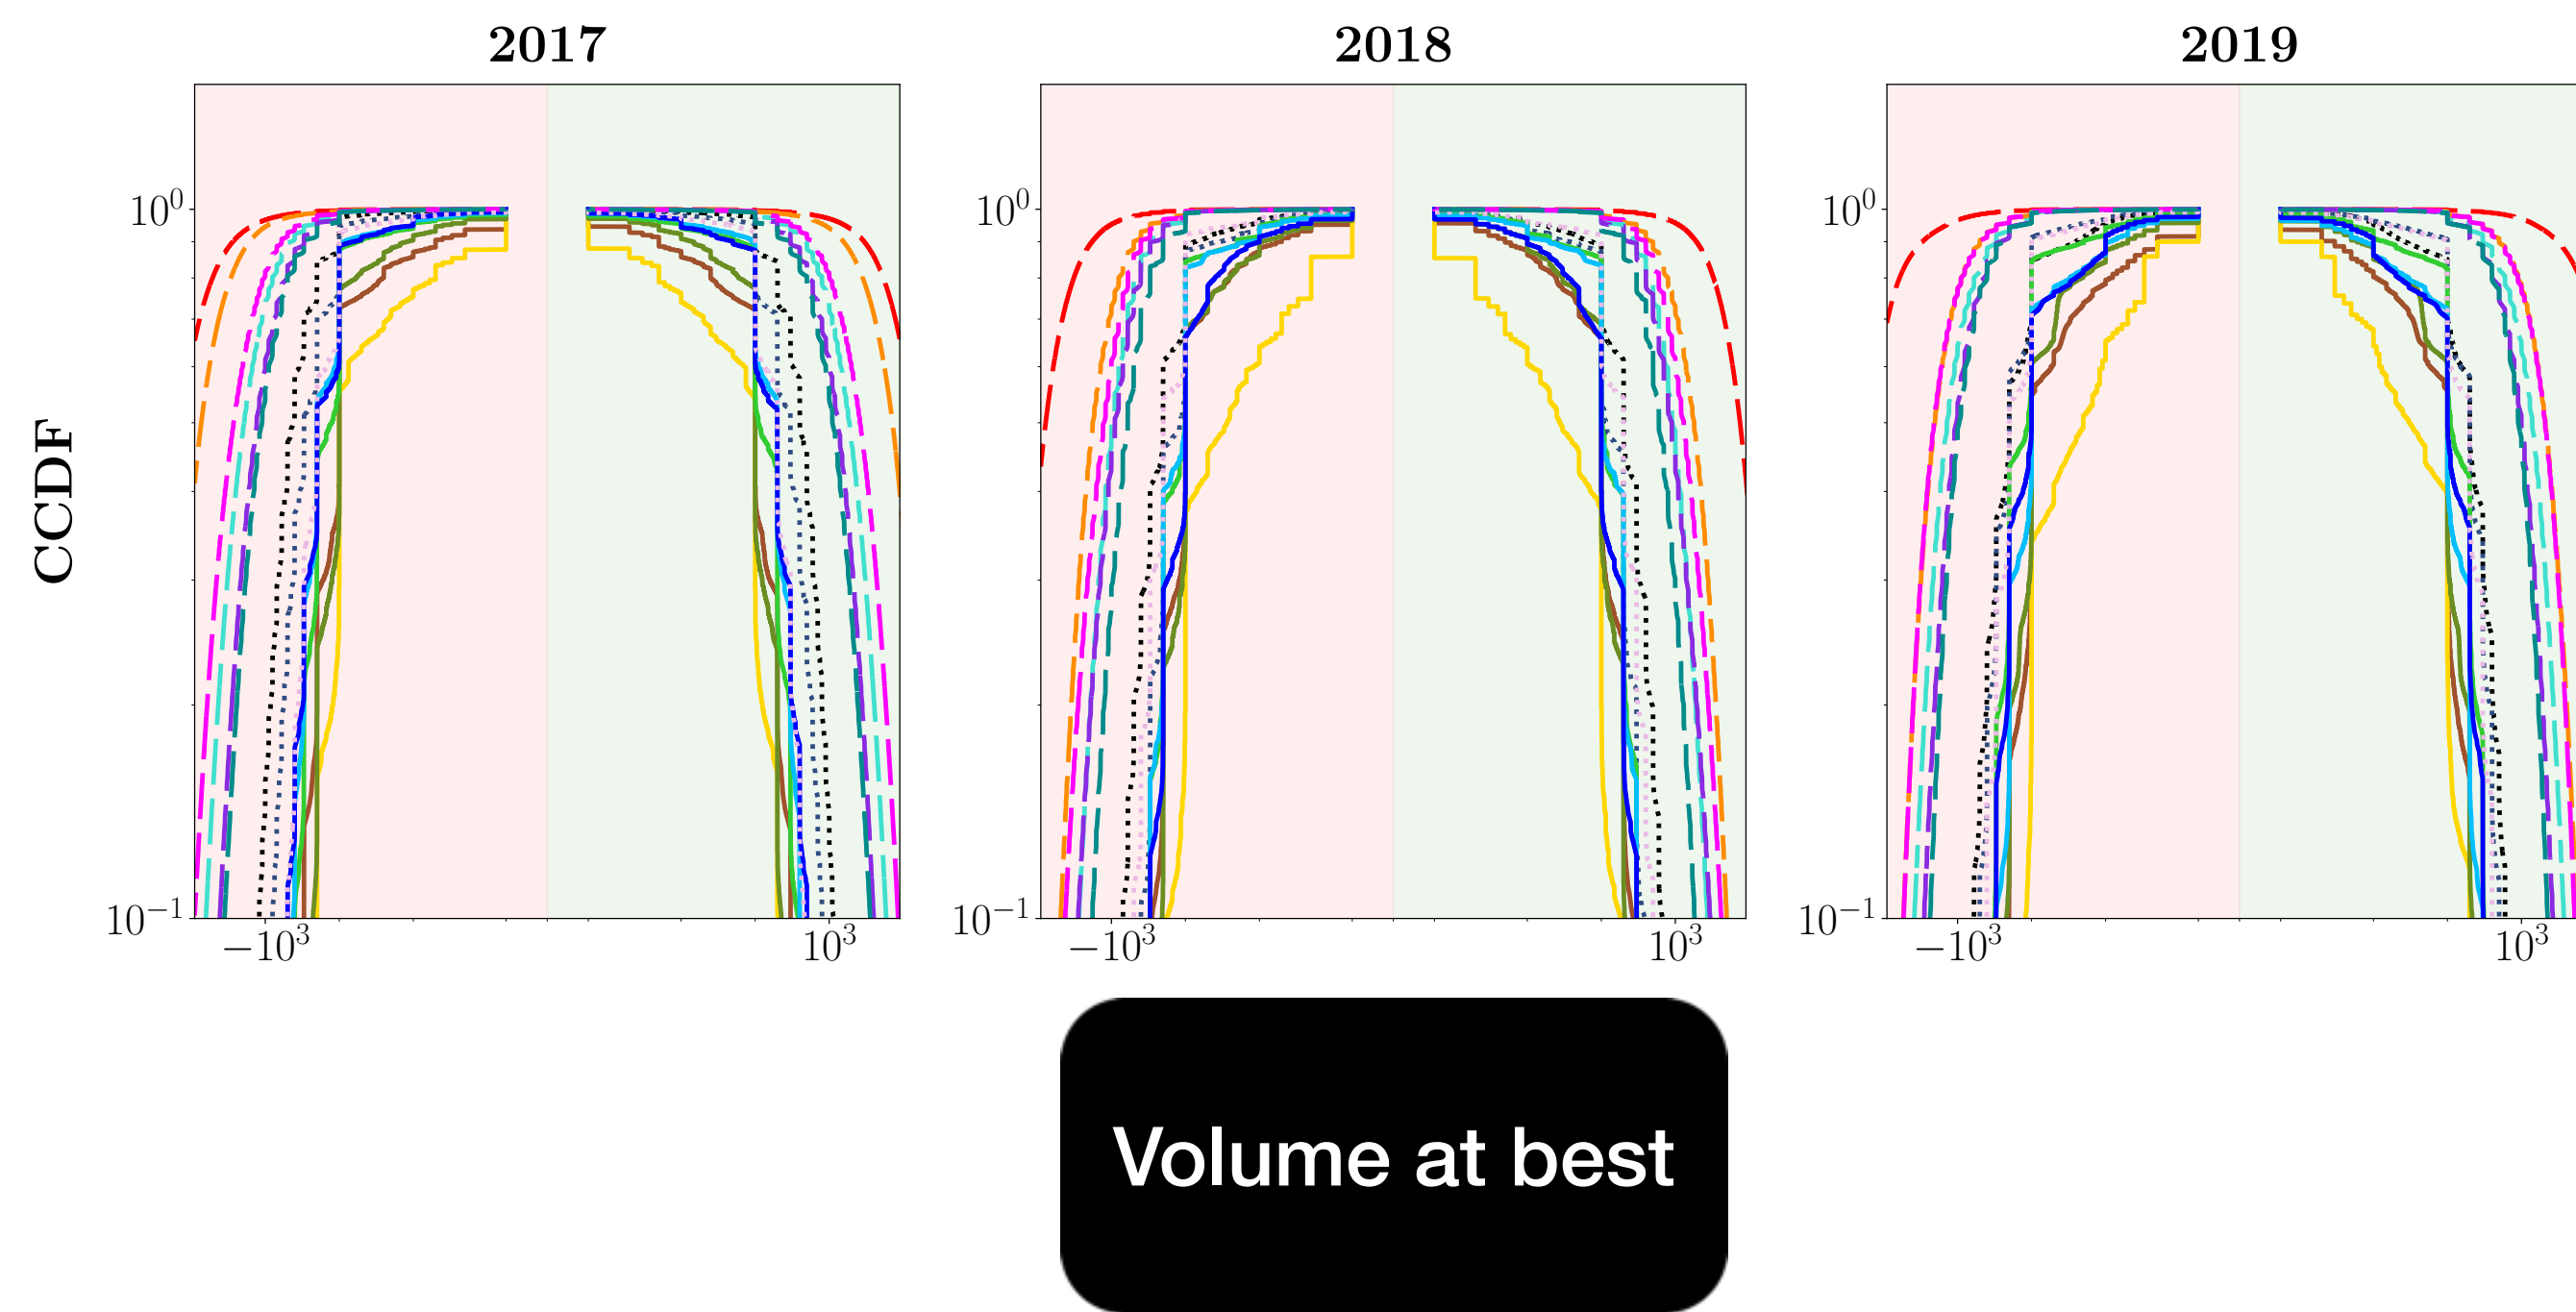
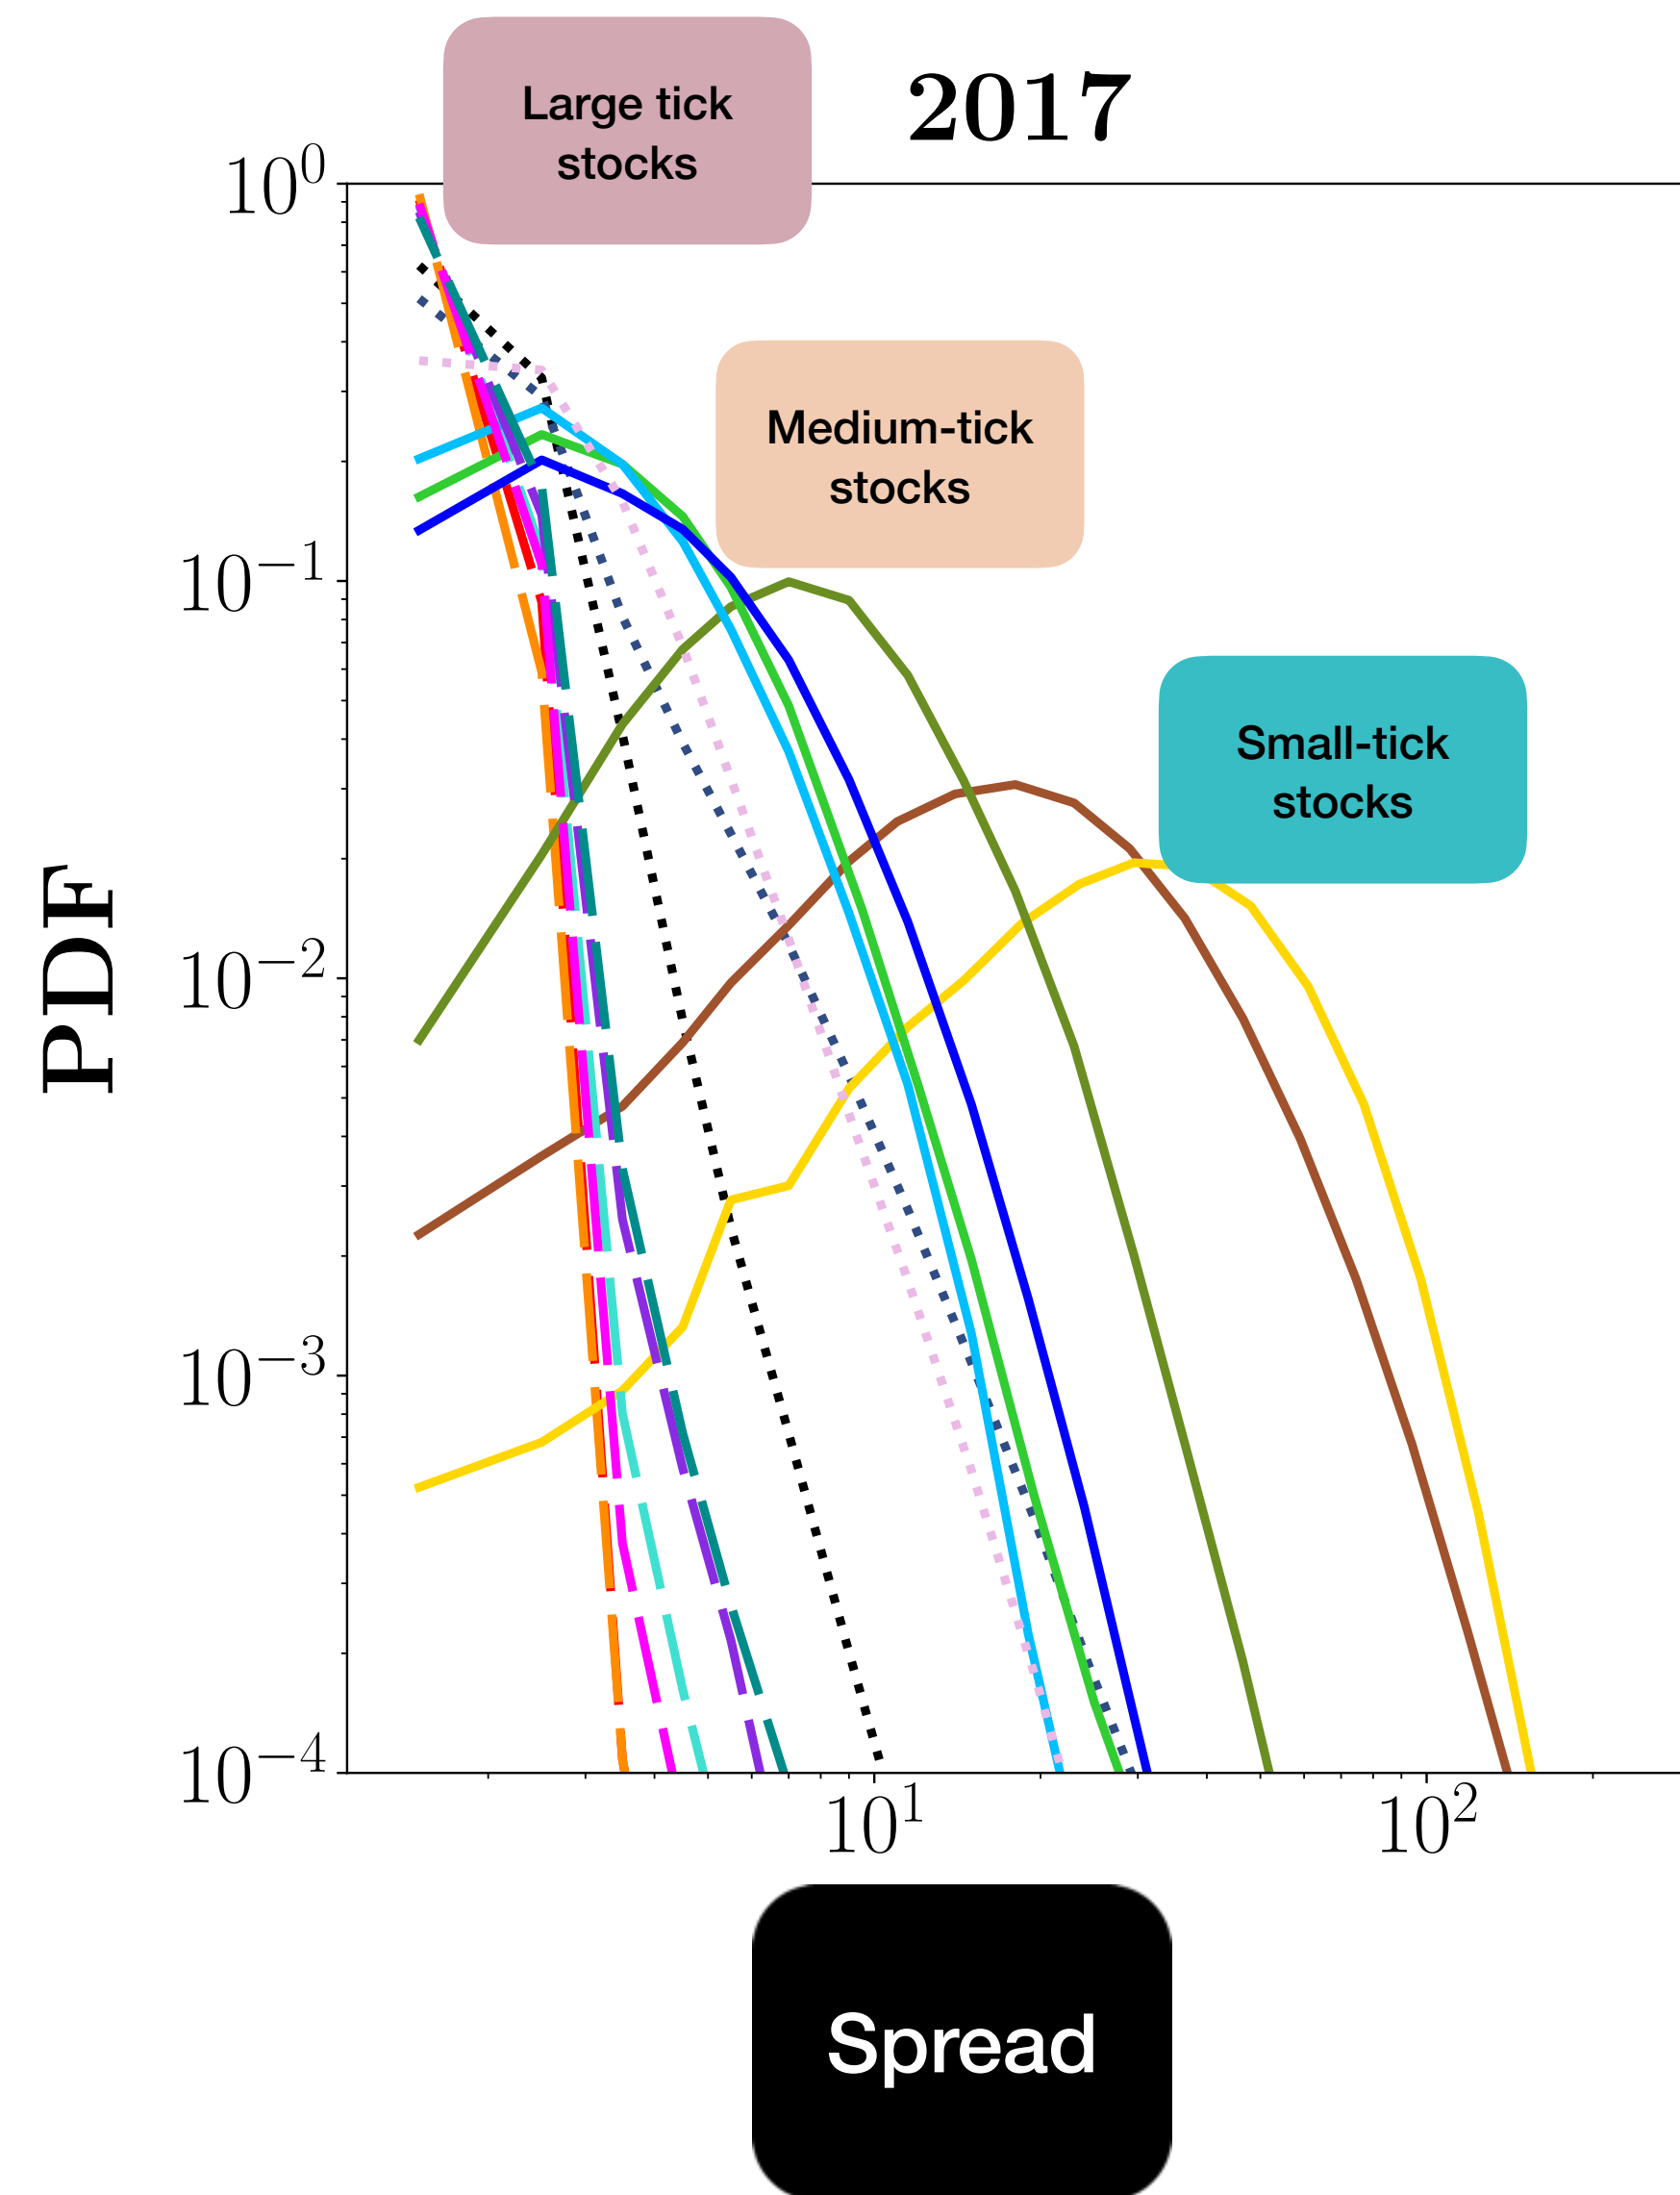
- High liquidity (at best)
- All price levels populated
- Tight spread



- Low liquidity
- Empty price levels
- Large spread

Microstructural properties

Stocks legend



Microstructural properties



$$\mathbb{L}(\tau) = \{p_\ell^{\text{ask}}(\tau), v_\ell^{\text{ask}}(\tau), p_\ell^{\text{bid}}(\tau), v_\ell^{\text{bid}}(\tau)\}_{\ell=1}^L$$

The spatial structure of the ‘compressed’ LOB representation is not homogeneous.

This representation is prone to dramatic changes due to occasional price-level shifts, significantly impacting predictability when treated as input for deep learning models.

Wu, Y., Mahfouz, M., Magazzeni, D. and Veloso, M., Towards Robust Representation of Limit Orders Books for Deep Learning Models. arXiv preprint arXiv:2110.05479, 2021

Small-tick stocks

	2018			
	Mean		Median	
	Ask	Bid	Ask	Bid
CHTR	71.67	68.84	55.00	54.00
GOOG	93.67	91.83	82.00	81.00
GS	19.07	20.33	15.00	16.00
IBM	12.47	12.70	11.00	11.00
MCD	12.32	12.73	10.00	11.00
NVDA	16.30	16.12	15.00	15.00

Medium-tick stocks

AAPL	9.52	9.54	9.00	9.00
ABBV	23.27	19.79	11.00	11.00
PM	13.64	13.60	11.00	11.00

Large tick stocks

BAC	9.00	9.00	9.00	9.00
CSCO	9.01	9.01	9.00	9.00
KO	9.04	9.04	9.00	9.00
ORCL	9.05	9.06	9.00	9.00
PFE	9.03	9.03	9.00	9.00
VZ	9.07	9.09	9.00	9.00

$$[\mathbb{H}]_\tau^{\text{Ask}} = \frac{p_{10}^{\text{ask}}(\tau) - p_1^{\text{ask}}(\tau)}{\theta}$$

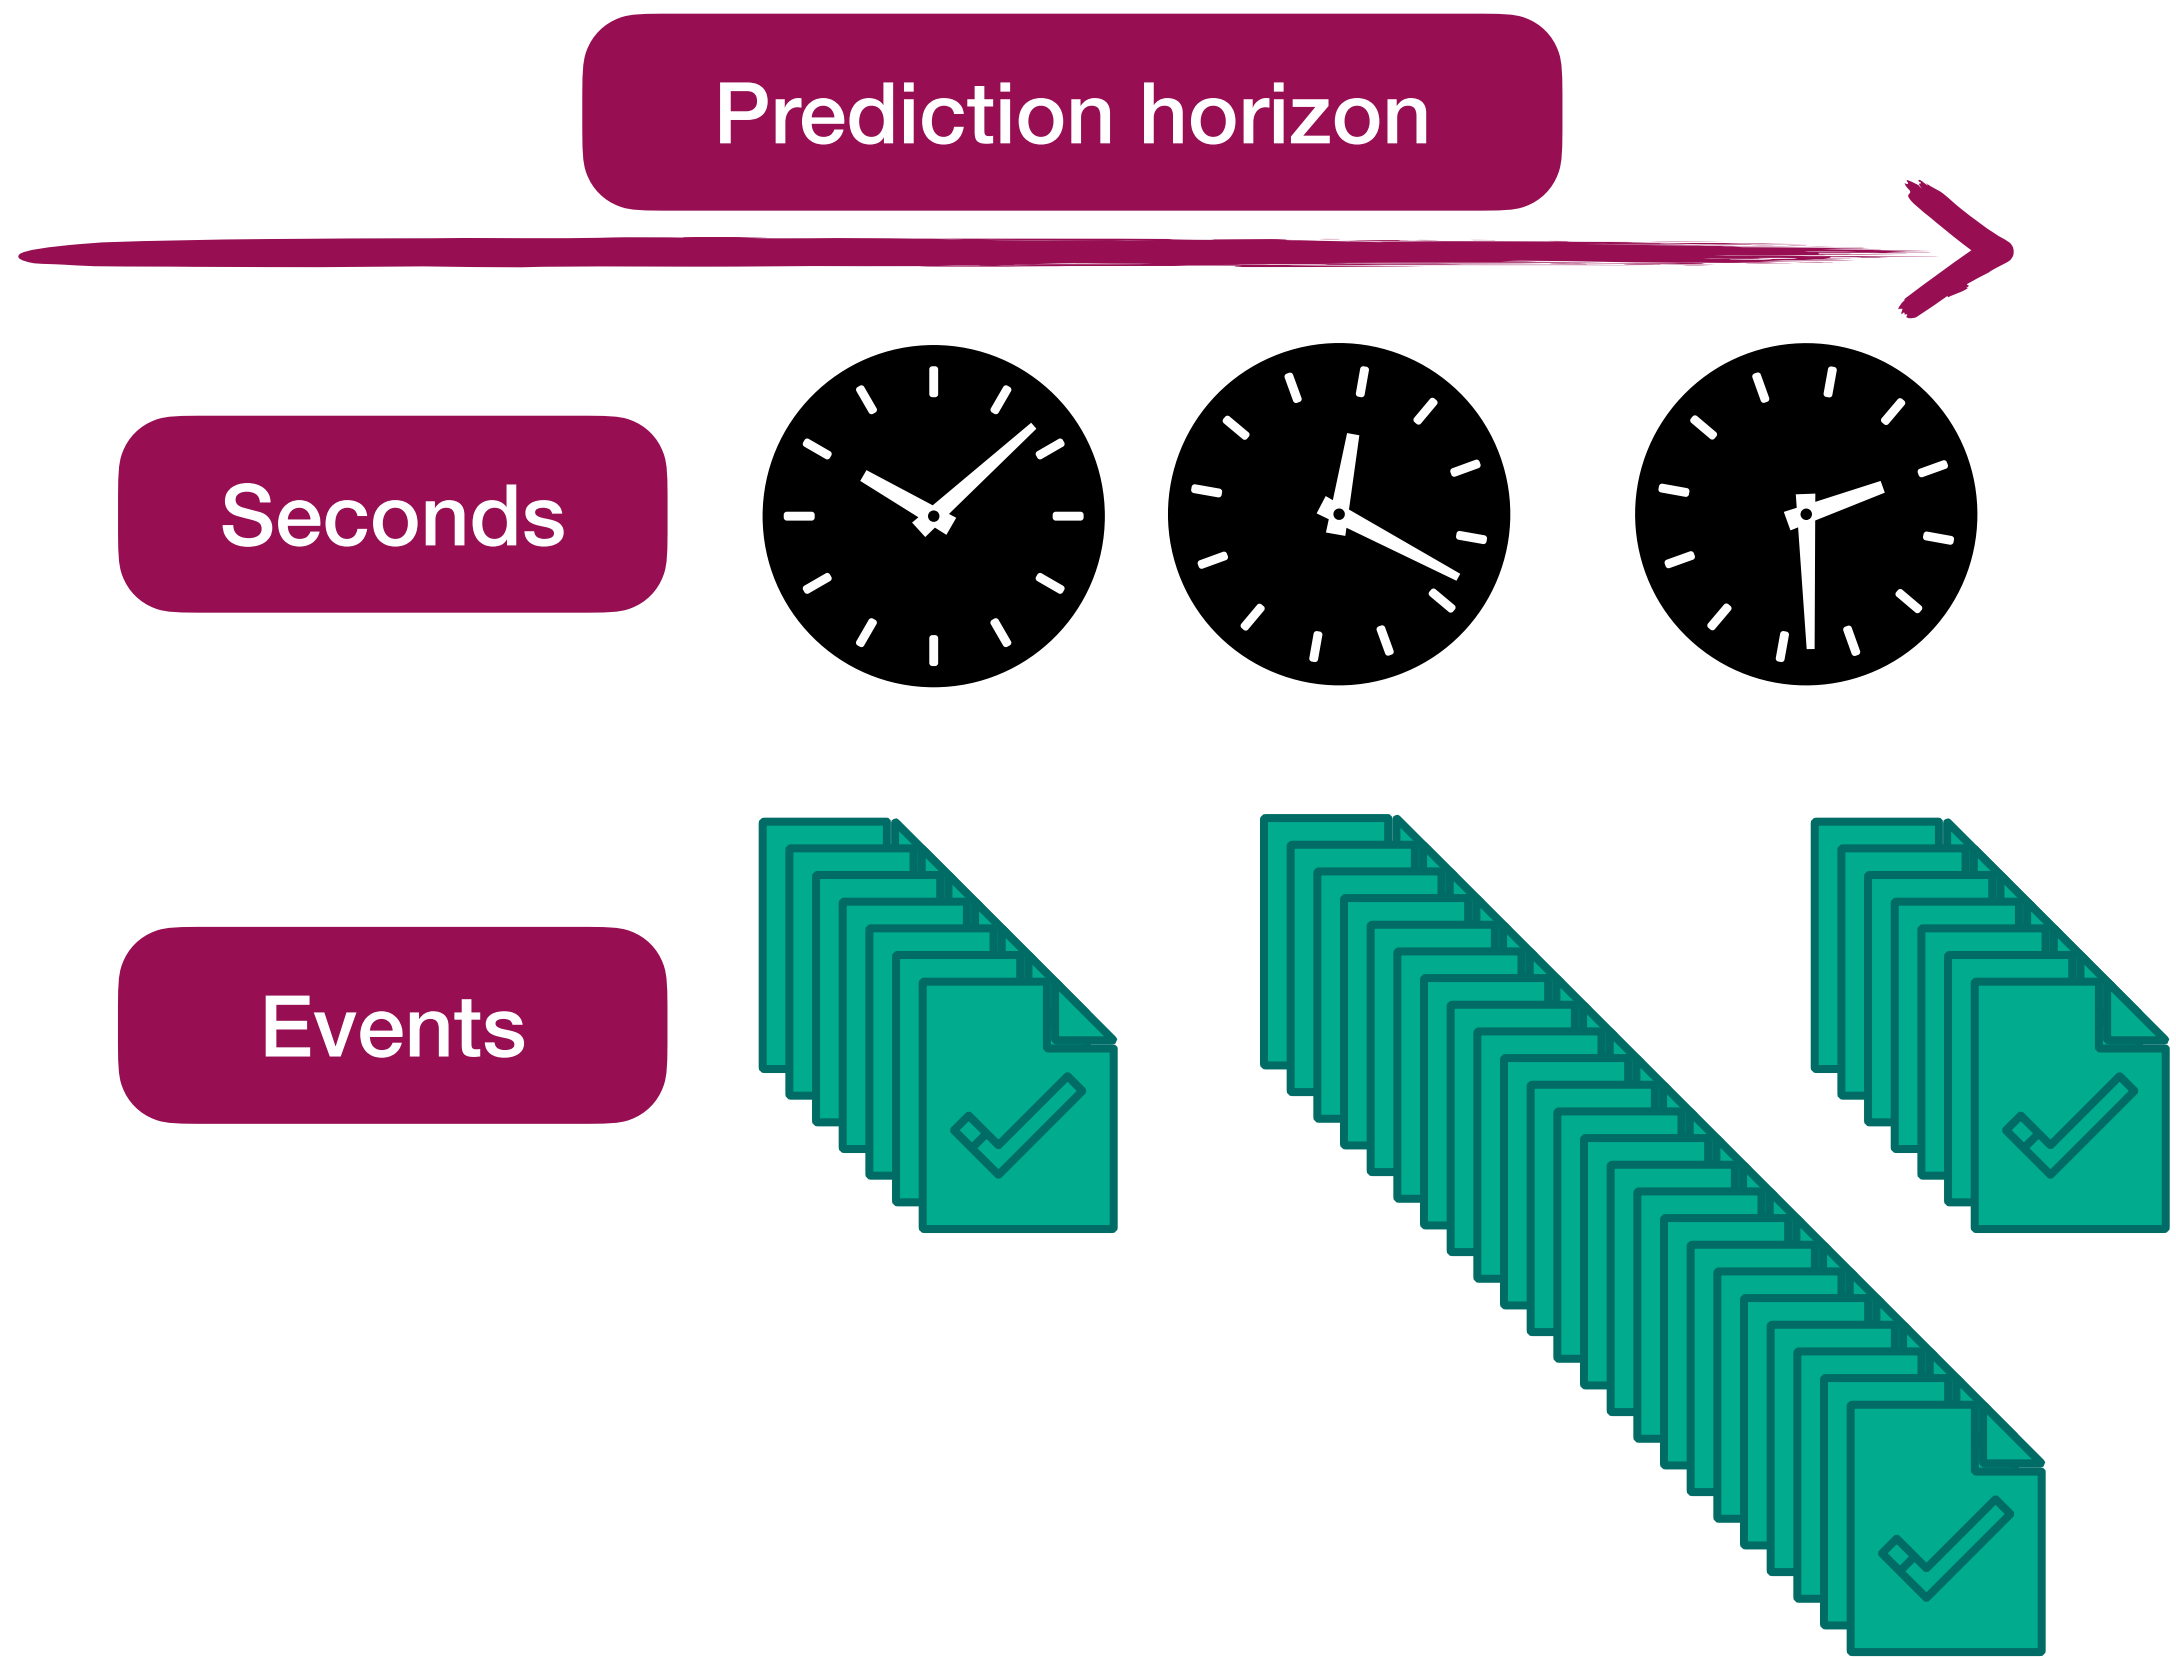
$$[\mathbb{H}]_\tau^{\text{Bid}} = \frac{p_1^{\text{bid}}(\tau) - p_{10}^{\text{bid}}(\tau)}{\theta}$$

Actual LOB depth

About time

Average probability of physical/event time mapping

Ticker	<1s			≥1s & <10s			≥ 10s			
	H10	H50	H100	H10	H50	H100	H10	H50	H100	
Small-tick stocks	CHTR	0.46	0.06	0.01	0.41	0.36	0.18	0.13	0.58	0.81
	GOOG	0.76	0.32	0.14	0.18	0.53	0.56	0.06	0.15	0.30
	GS	0.51	0.05	0.01	0.39	0.56	0.29	0.10	0.39	0.70
	IBM	0.57	0.07	0.01	0.35	0.65	0.40	0.08	0.28	0.59
	MCD	0.00	0.00	0.00	0.00	0.00	0.00	1.00	1.00	1.00
	NVDA	0.72	0.24	0.06	0.18	0.68	0.73	0.10	0.08	0.21
Medium-tick stocks	AAPL	0.92	0.55	0.23	0.04	0.41	0.73	0.04	0.04	0.04
	ABBV	0.00	0.00	0.00	0.31	0.63	0.40	0.69	0.37	0.60
	PM	0.63	0.07	0.01	0.32	0.69	0.42	0.05	0.24	0.57
Large tick stocks	BAC	0.78	0.43	0.30	0.12	0.47	0.55	0.10	0.10	0.15
	CSCO	0.80	0.45	0.24	0.13	0.46	0.58	0.07	0.09	0.18
	KO	0.68	0.33	0.14	0.23	0.49	0.51	0.09	0.18	0.35
	ORCL	0.75	0.39	0.18	0.17	0.45	0.56	0.08	0.16	0.26
	PFE	0.72	0.38	0.17	0.19	0.48	0.54	0.09	0.14	0.29
VZ	0.73	0.38	0.16	0.19	0.48	0.57	0.08	0.14	0.27	



The heterogeneous mapping between physical time (seconds) and event time (market updates) leads to varying practical relevance of forecasts across stocks.

Part 1: Evaluating Prediction Accuracy & Usability

Objective:

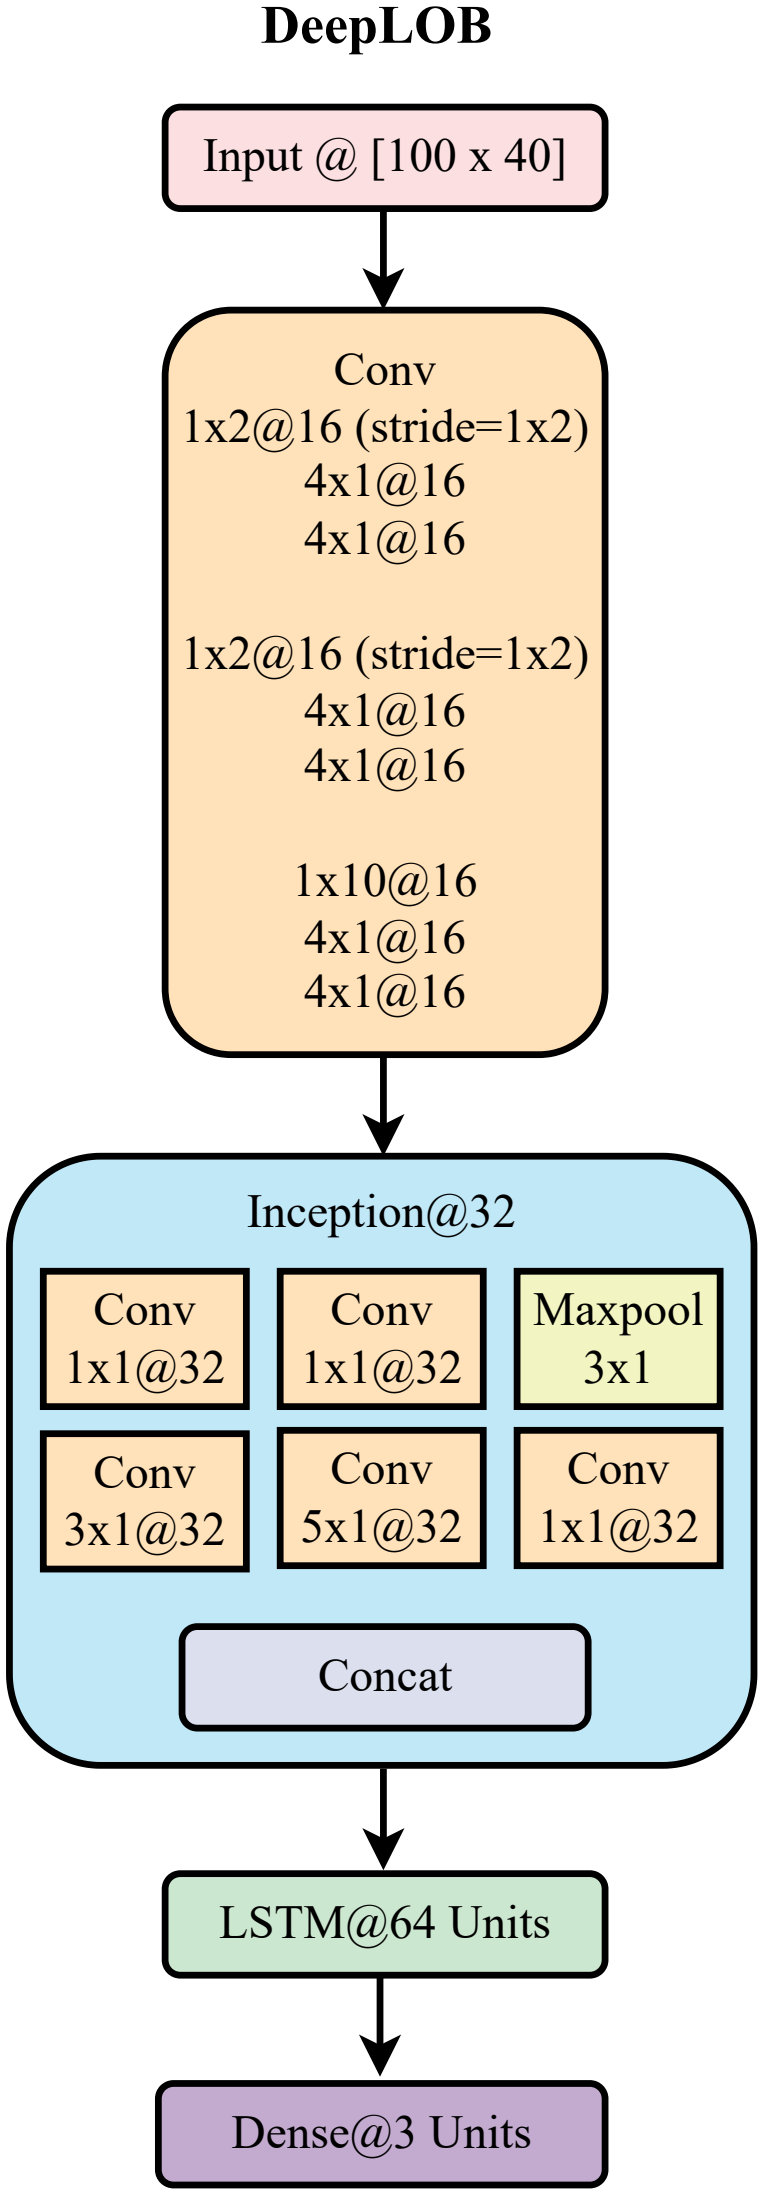


Assess the effectiveness of a state-of-the-art deep learning architecture (DeepLOB) across assets with different microstructural properties.

Focus:

- Evaluate predictive performance at multiple horizons.
- Investigate how asset characteristics (e.g., tick size, liquidity) impact accuracy.
- Assess practical usability in automated trading strategies.

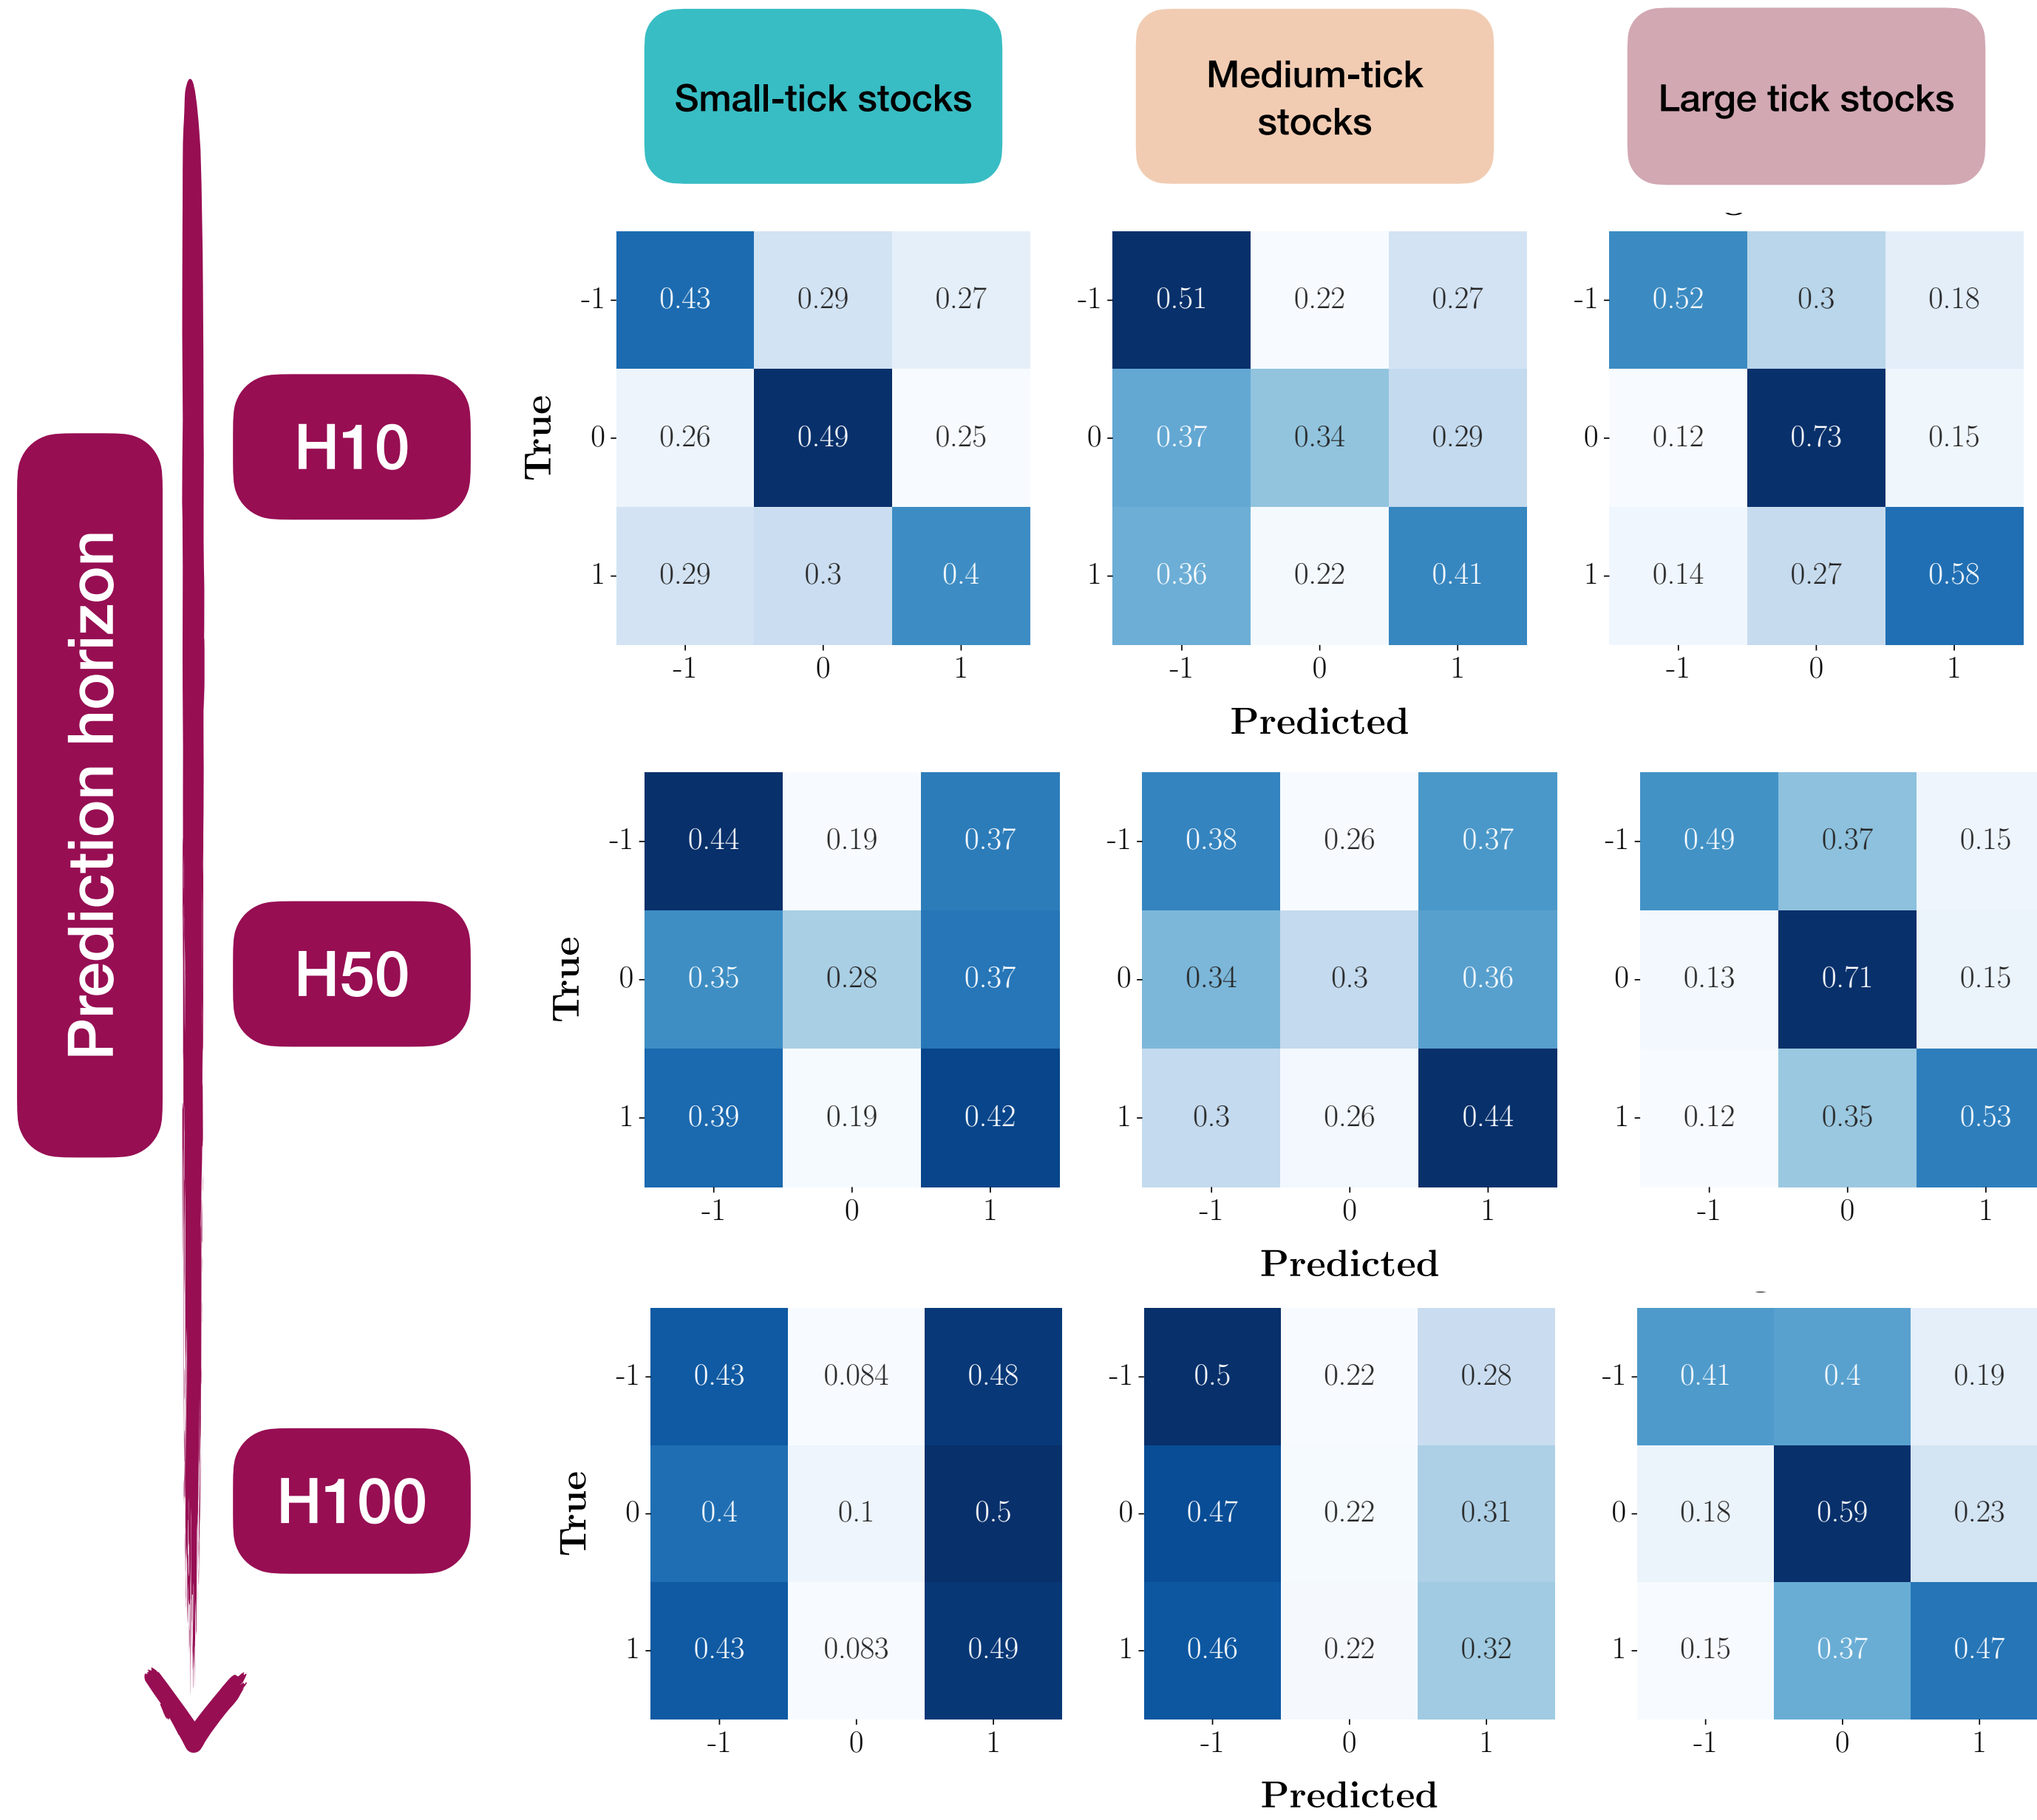
DeepLOB architecture

$$\mathbb{L}(\tau) = \{p_\ell^{\text{ask}}(\tau), v_\ell^{\text{ask}}(\tau), p_\ell^{\text{bid}}(\tau), v_\ell^{\text{bid}}(\tau)\}_{\ell=1}^L$$



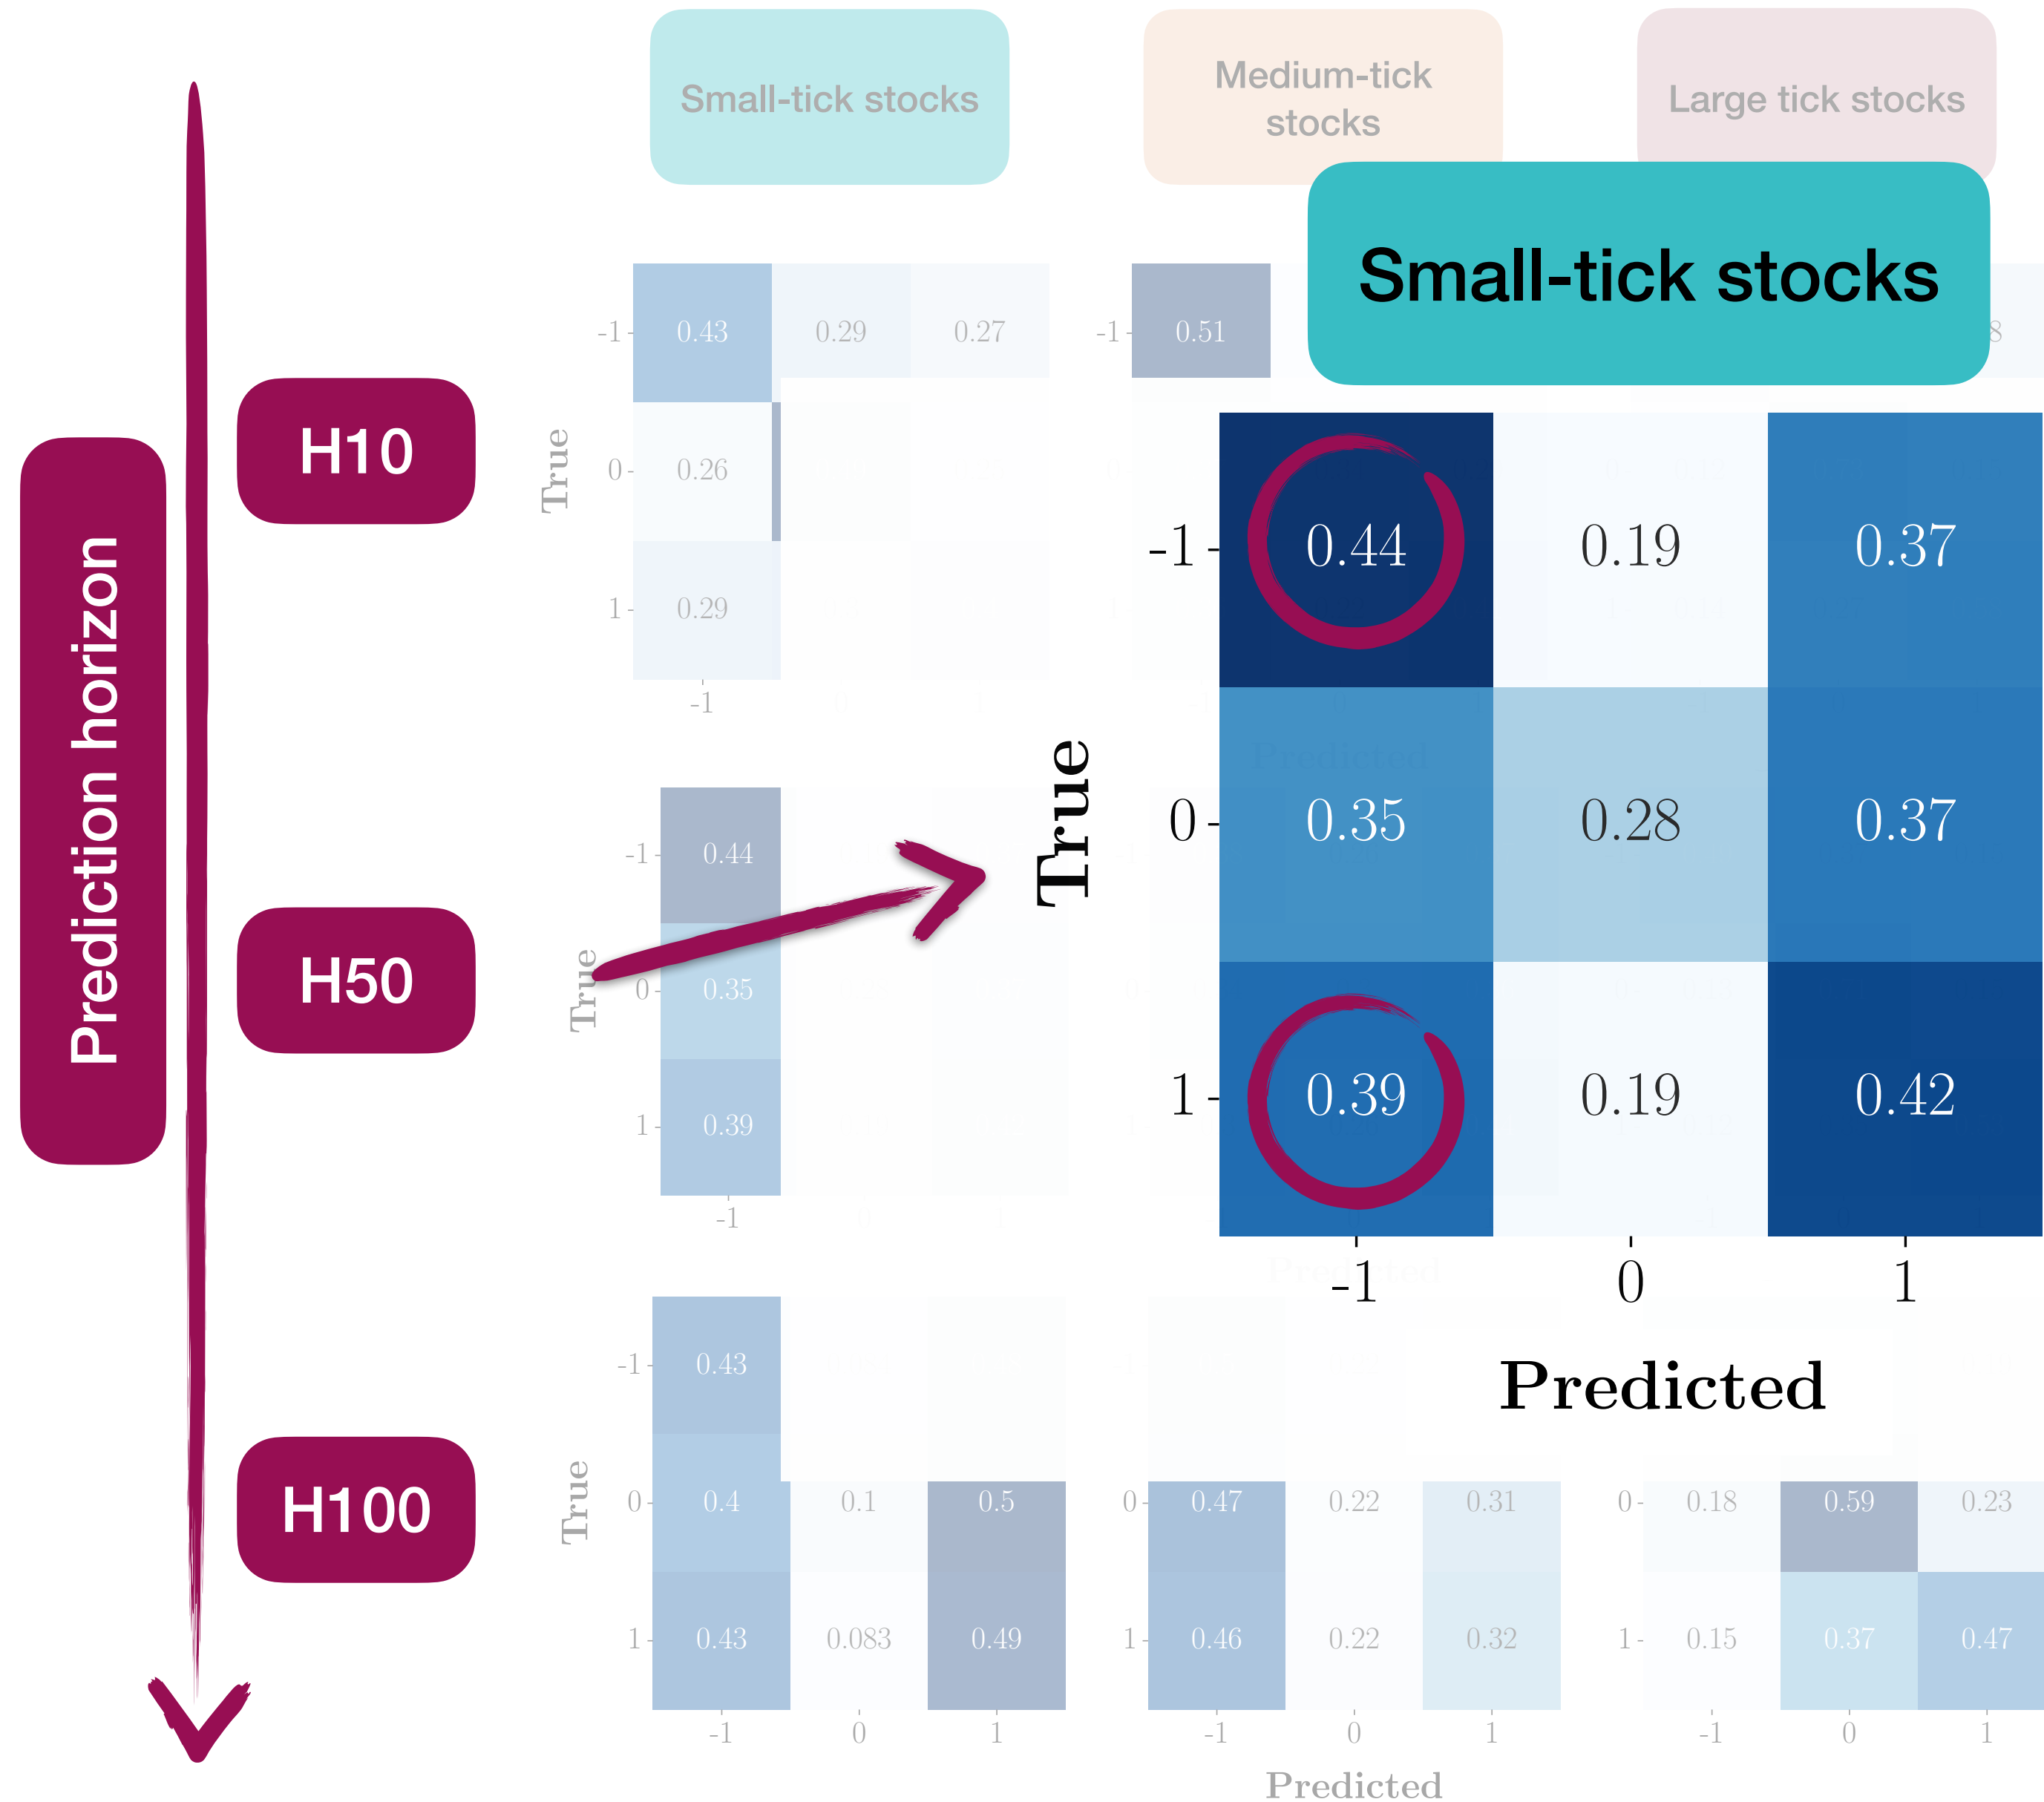
- Spatial features (across price/volumes)
- Local temporal interactions
- Dynamical behaviour over multiple timescales
- Temporal features (LOB time evolution)

Types of errors



- **Large-tick** stocks consistently exhibit superior predictive performance, with fewer misclassifications and robust results across all horizons.

Types of errors



- Models trained on **small- and medium-tick** stocks show significant reciprocal misclassification between extreme classes (i.e., -1 and 1), with higher misclassification rates at longer prediction horizons.

Testing performance

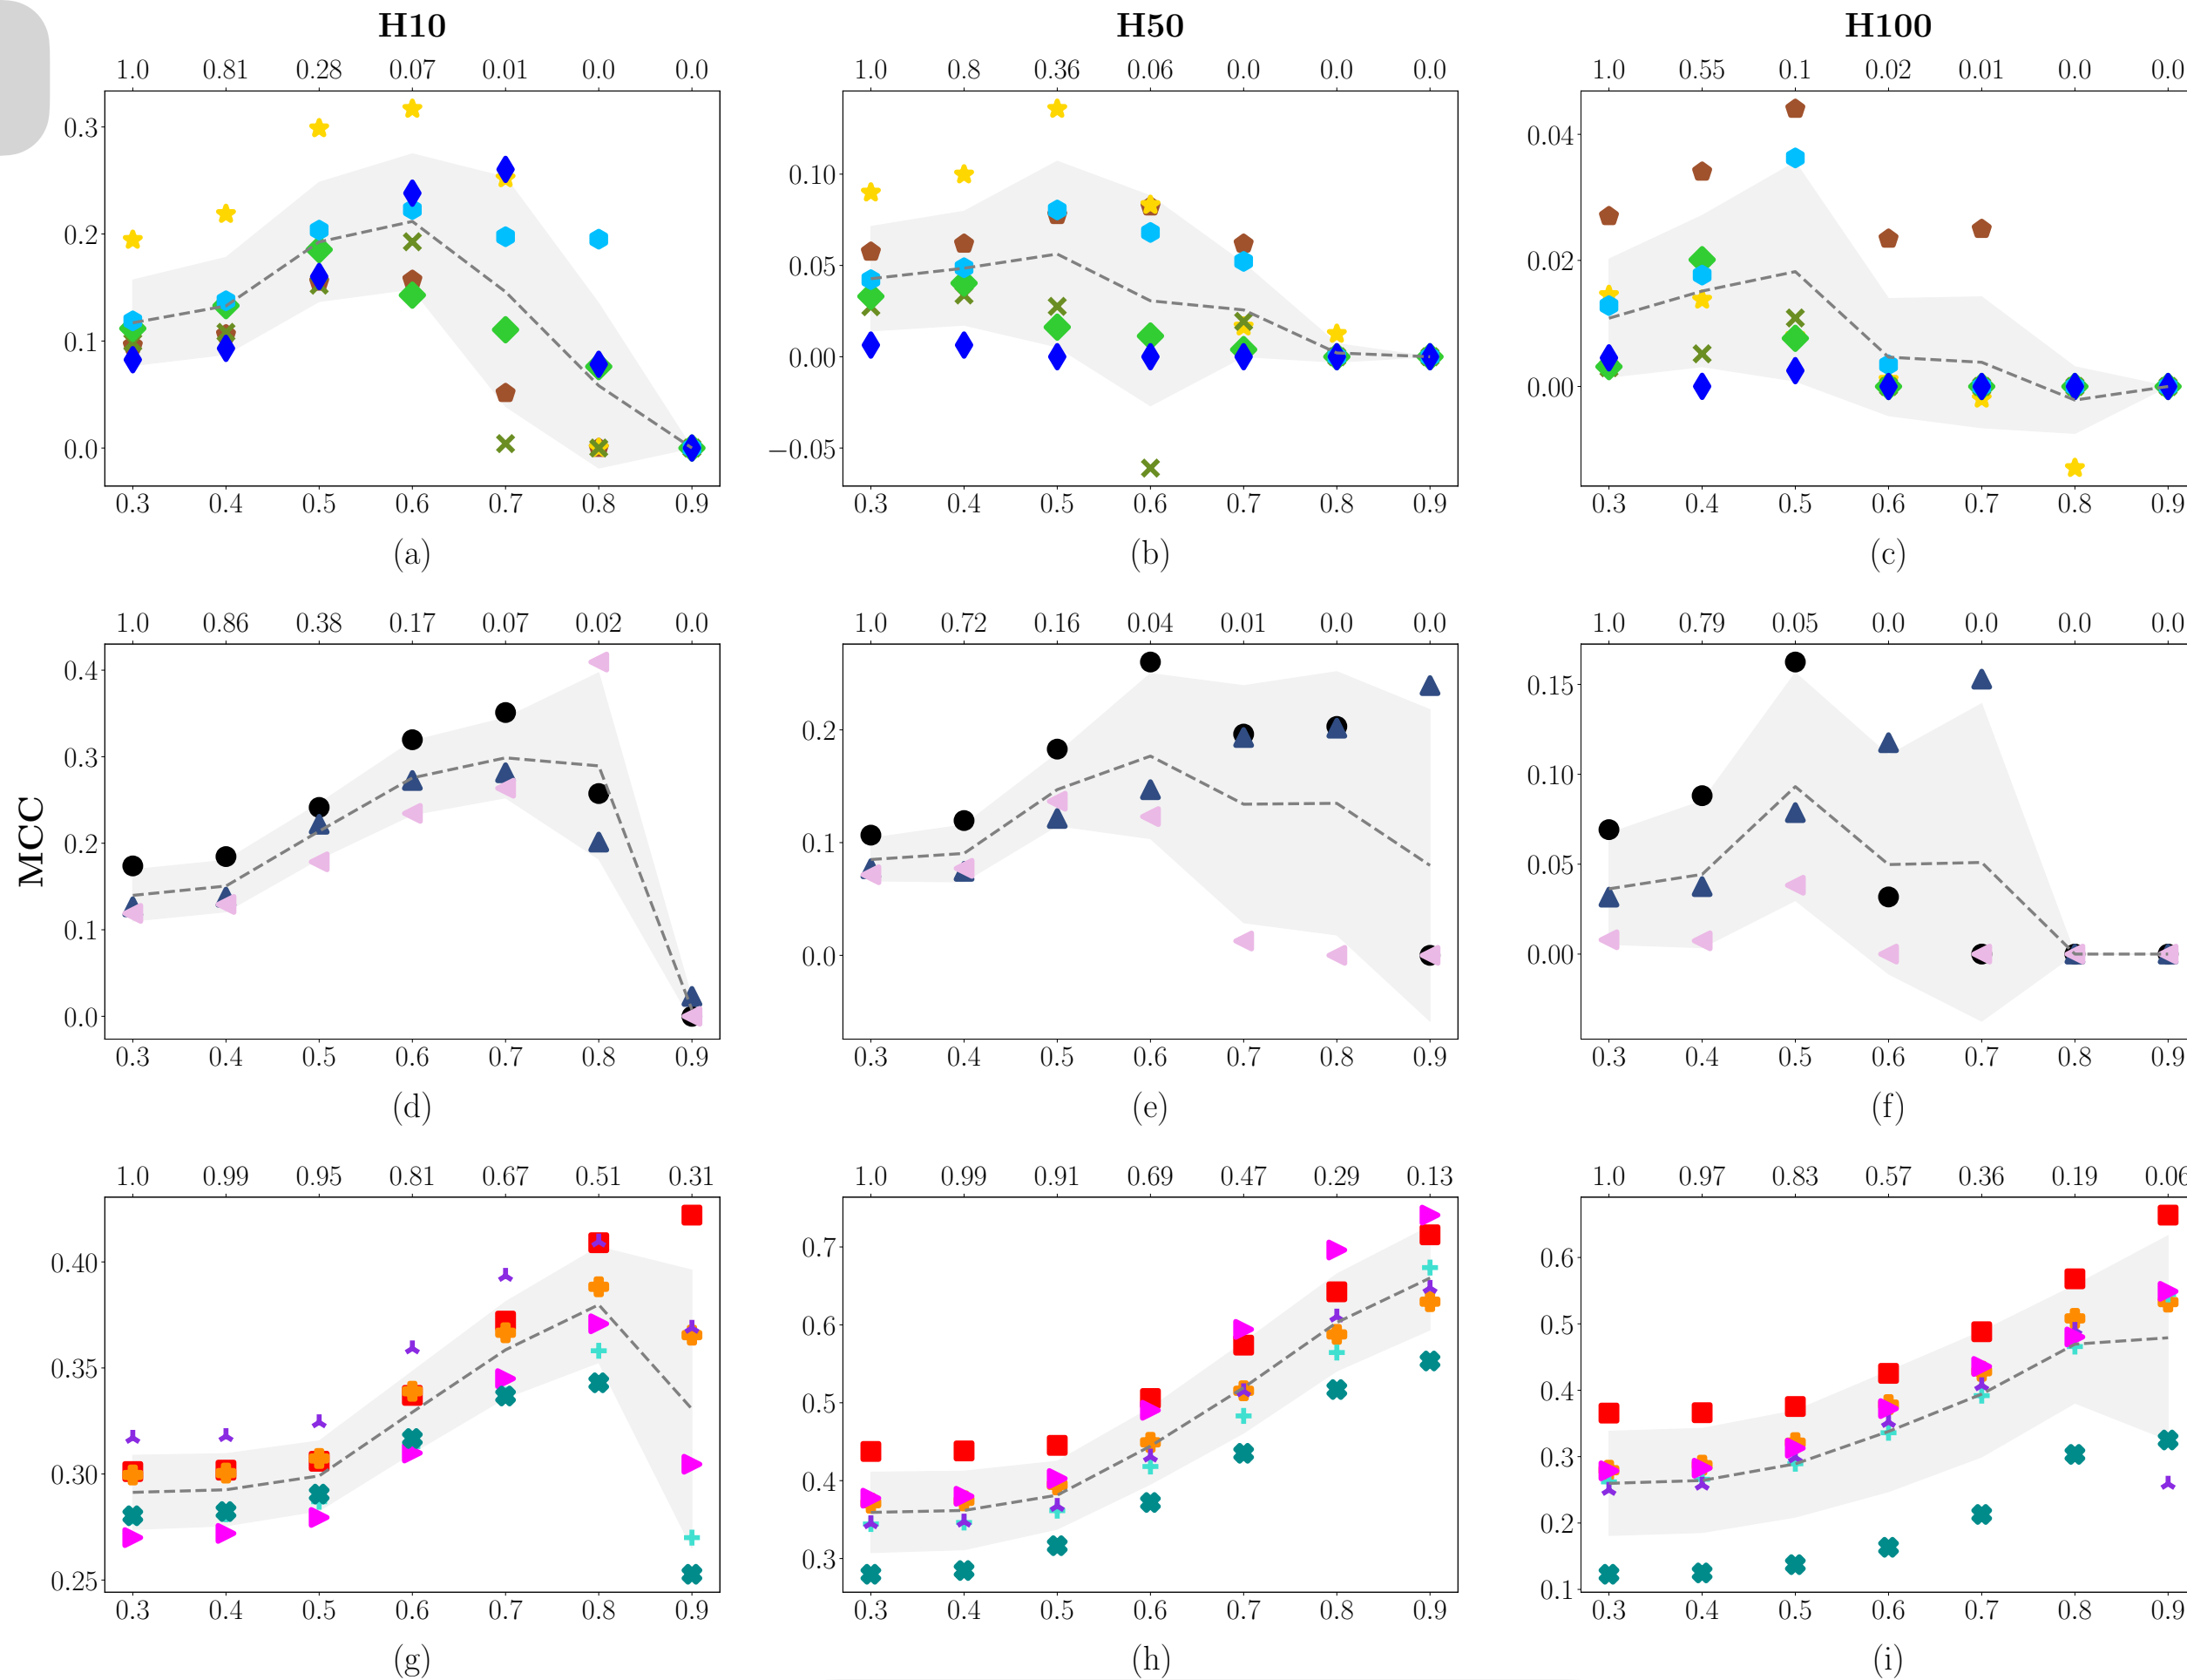
Prediction horizon

Stocks legend

CHTR	IBM	AAPL	BAC	ORCL
GOOG	MCD	ABBV	CSCO	PFE
GS	NVDA	PM	KO	VZ

% Data available

MCC



Small

Medium

Large

- The application of **probability thresholds** enhances predictive accuracy for all classes.
- Large-tick stocks** show both higher accuracy and greater resilience in maintaining a substantial proportion of usable forecasts even at stricter thresholds.

Probability Thresholds

Do we have the right tools?



Standard ML metrics do not necessarily align with trading success, and naïve backtesting can mislead by ignoring real-world frictions like execution risk and costs.




Traditional ML metrics drawbacks

- Resistance to class imbalances
- Difficulty in evaluating actual trading performances

Back-testing drawbacks

- Assumptions on orders' execution
- Accounting of slippage
- Estimation of transaction costs

Predictions \rightarrow trades




O: open **buy/sell** position

M: maintain **buy/sell** position

C: close **buy/sell** position

Predictions trades

Example 1



Targets

-1 -1 1 1 1 1 -1 -1 0 -1 1 1 1

Predictions

-1 -1 -1 -1 0 0 -1 -1 0 -1 1 1 1

MCC: 0.59

F1: 0.64

PT: 3

TT: 1

CT: 0

p_T : 0.00

Example 2



Targets

-1 -1 1 1 1 1 -1 -1 0 -1 1 1 1

Predictions

-1 0 0 1 0 0 0 -1 -1 0 0 -1 1

MCC: 0.11

F1: 0.29

PT: 3

TT: 3

CT: 3

p_T : 1.00

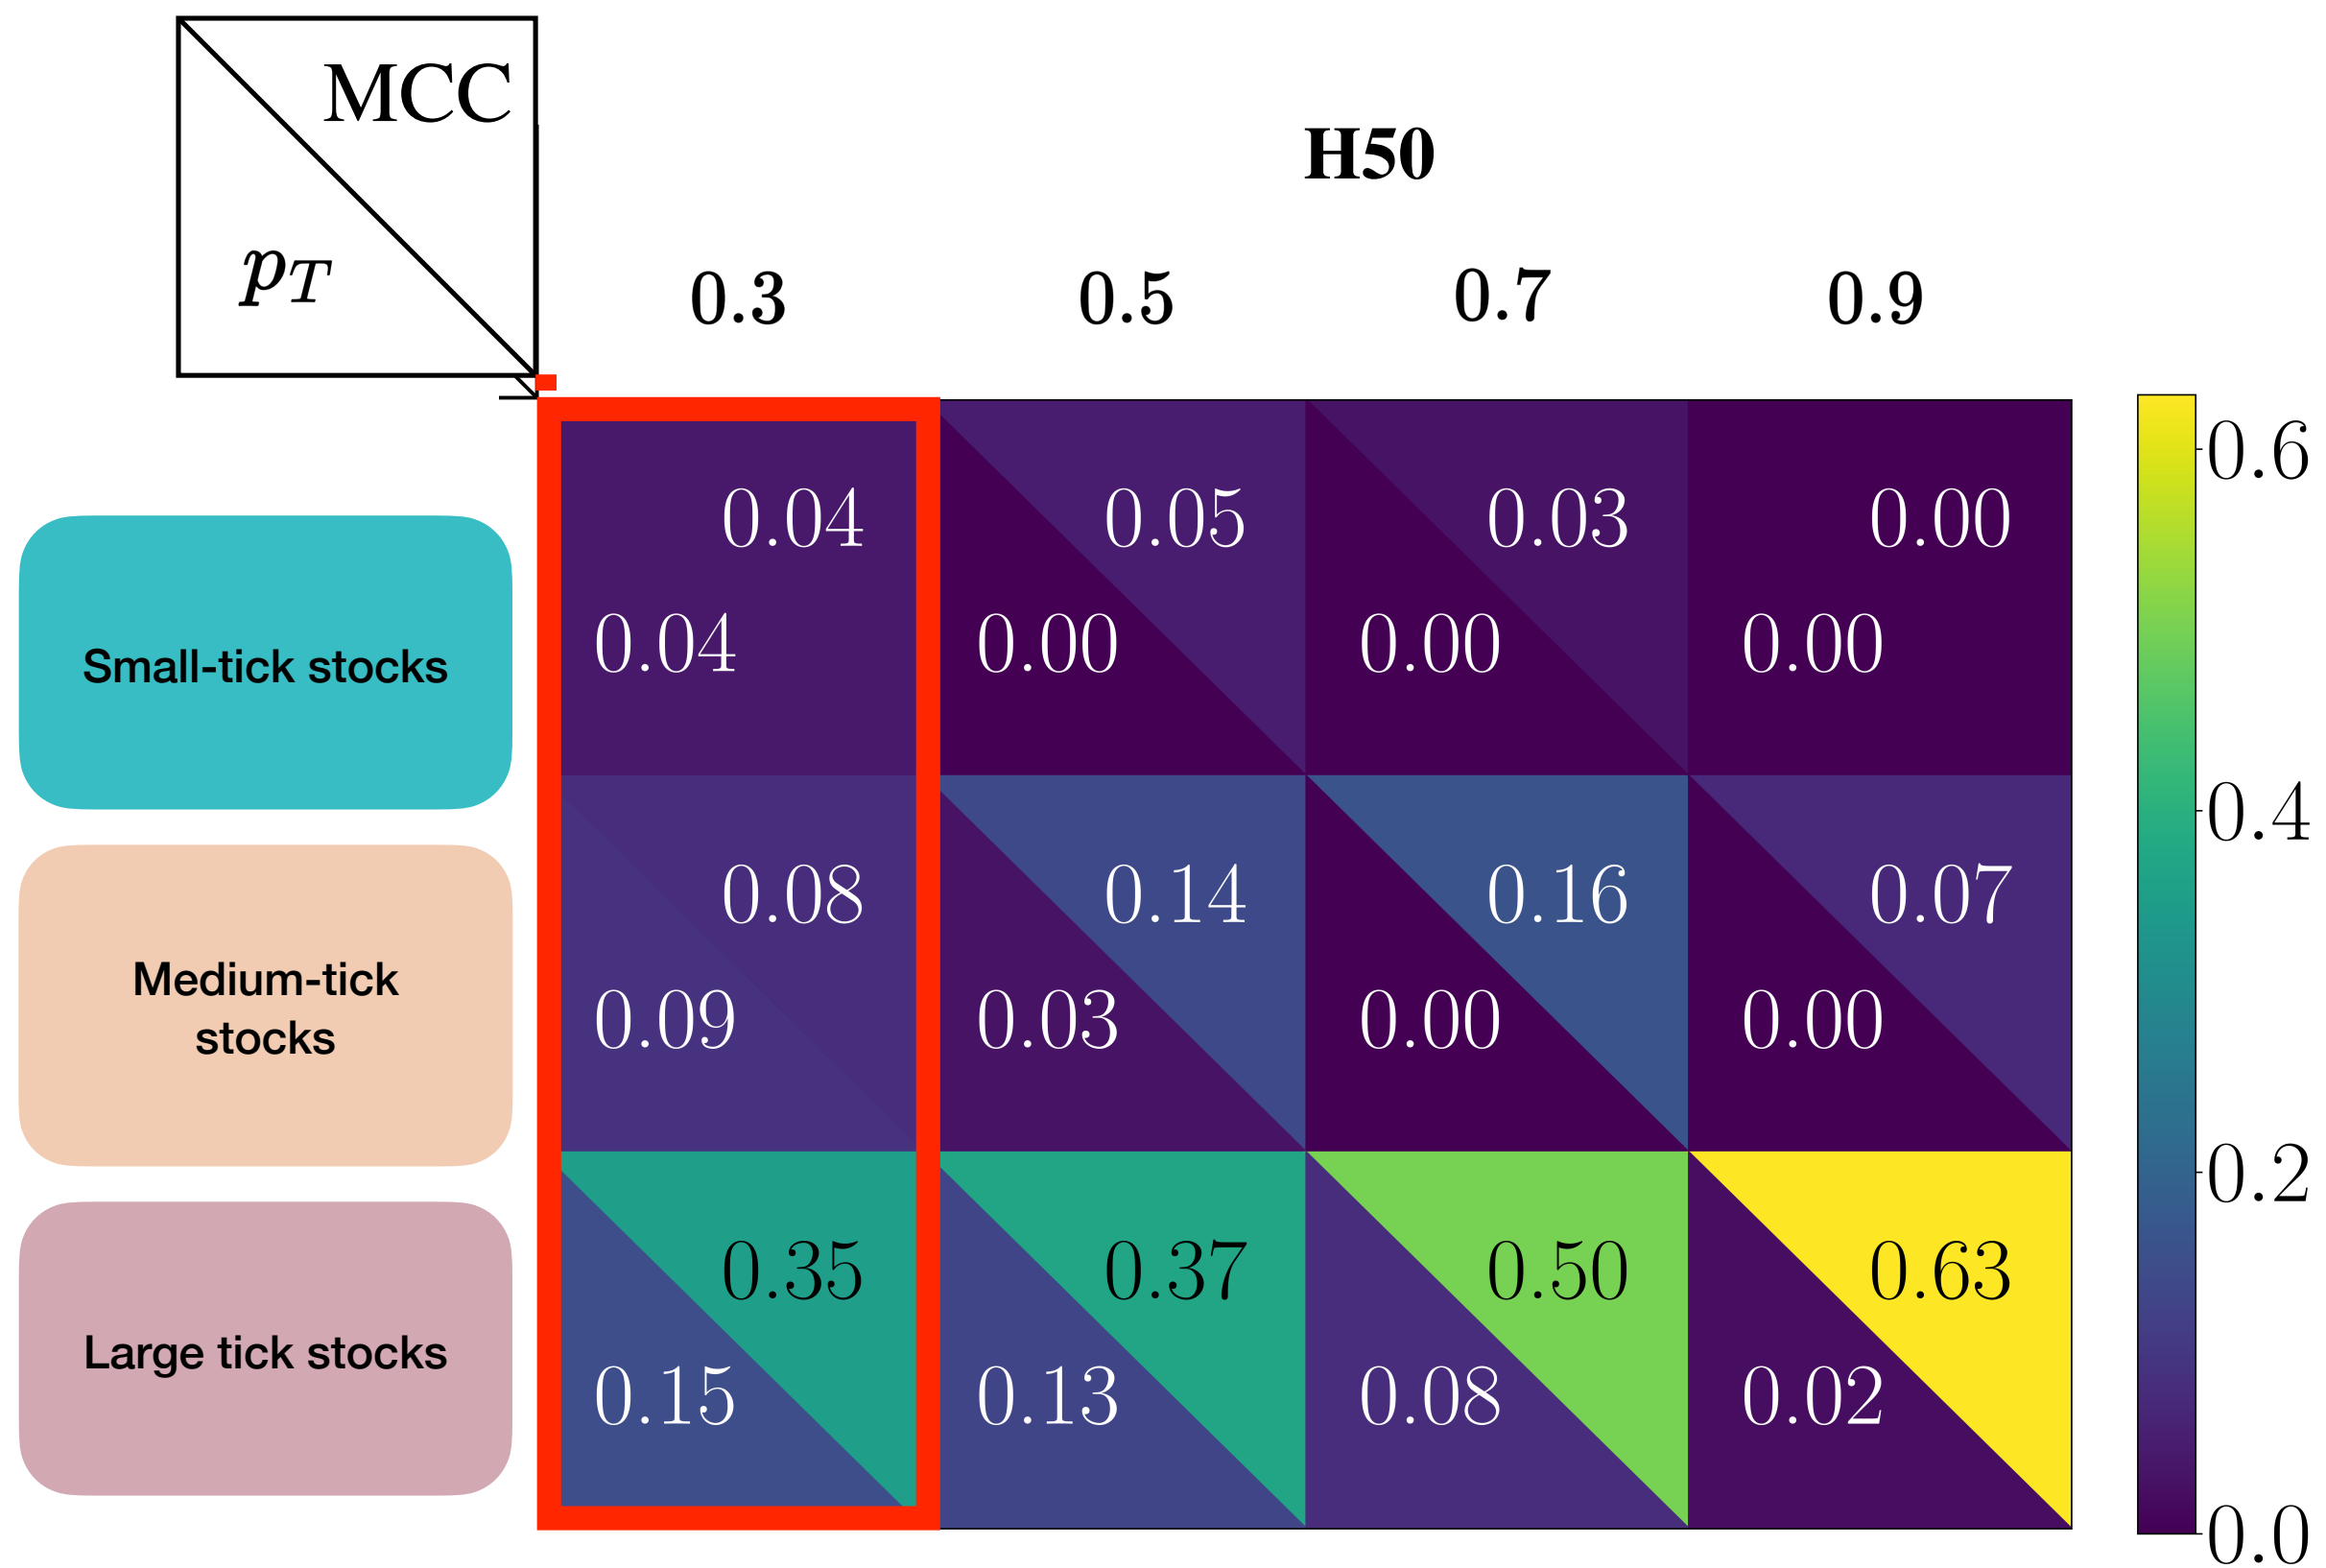
Predictions ➔ trades

When you are right matters more than how often

- In this application domain, we are more interested in the *chronological location of prediction errors* rather than in the number of their occurrences.
- We are interested in (i) having at least one correct prediction for each ‘Down’ or ‘Up’ movement; and, consequently, (ii) in not having any premature closing signal for an existing open position. The nature and the number of other errors are tolerable when these two conditions are satisfied.
- In real-world scenarios, also *probabilities* associated with forecasts should be taken into account: we can decide to enter or exit a position based on the probability associated with the forecast (i.e., the strength of the signal).



Predictions \rightarrow trades



- p_T decreases for increasing probability thresholds and increases moving from small-tick stocks to large-tick stocks;
- MCC increases for increasing probability thresholds (this is more evident moving to longer prediction horizons) and increases also moving from small-tick stocks to large-tick stocks.

Main take aways - part 1

**Microstructural
properties
matter**

Large tick stocks are more predictable across all prediction horizons

**Rethinking
performance
metrics**

ML metrics may not be suitable in assessing the goodness of predictions in a financial context.

Part 2: Enhancing DeepLOB with Complexity Science Tools

Objective:


Augment DeepLOB using Information Filtering Network (IFN) techniques to capture structured dependencies in limit order books.

Focus:


- Preprocessing LOB data with complexity science tools.
- Integrating filtered information into deep learning models.
- Demonstrate improved learning and interpretability of volume-flow interactions.

Information Filtering


Information Filtering Networks are an effective tool to *model dependency structures* among variables characterizing complex systems via network science, while imposing *topological constraints* (e.g., being a tree or a planar graph) and optimizing global properties (e.g., the model's likelihood).



Time series



Metric



Constraint & Network

Information Filtering + LOB data

We show the process to construct the **Triangulated Maximally Filtered Graph (TMFG)** from Limit Order Book (LOB) data.

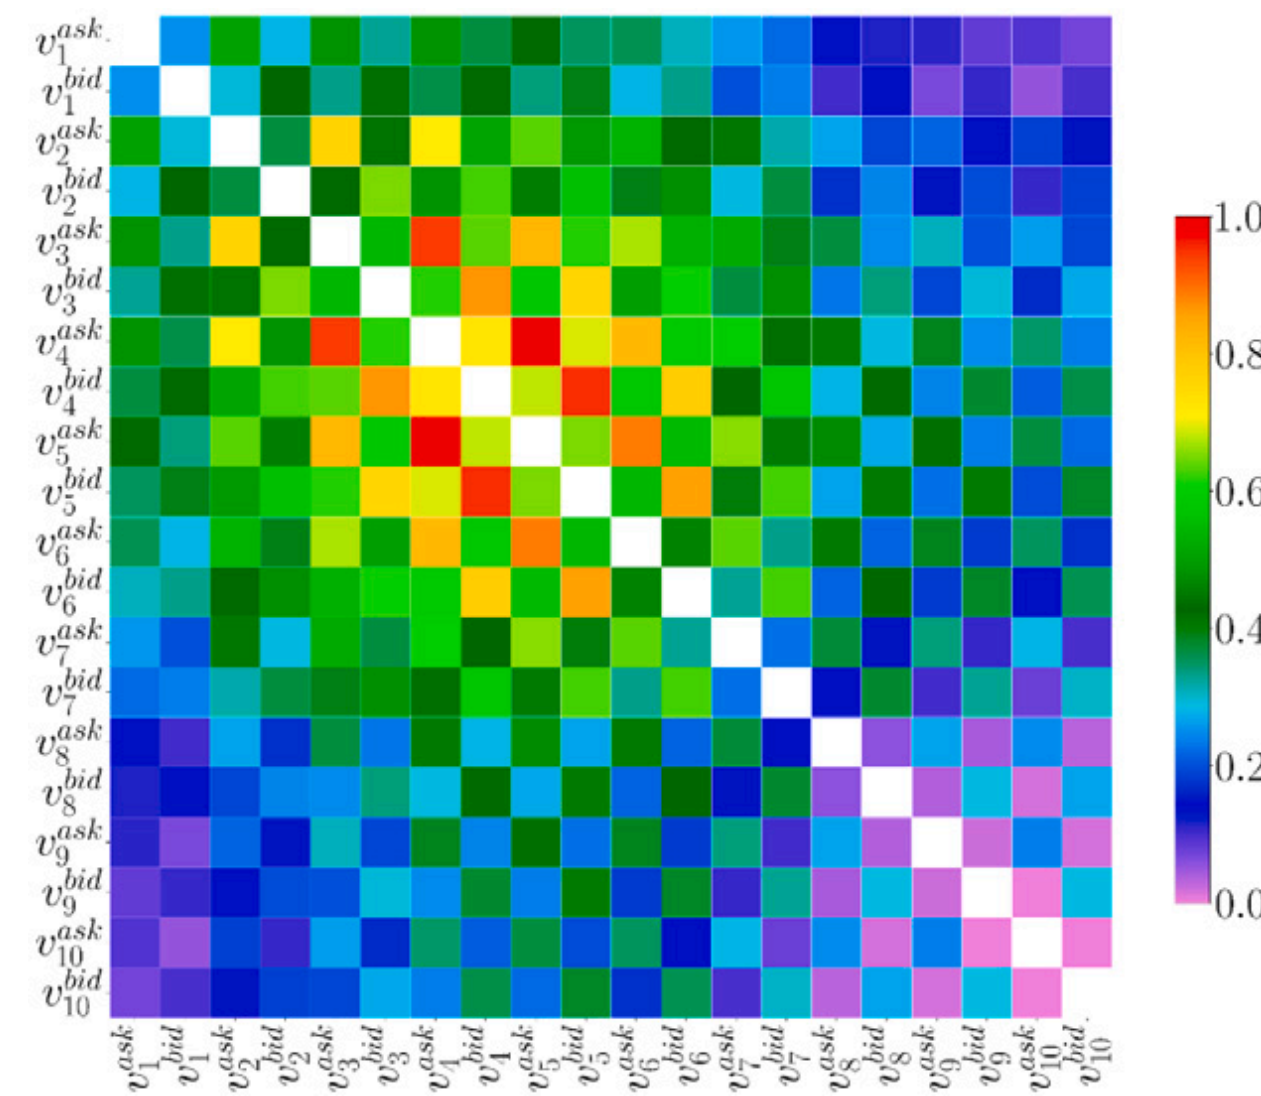
(a) In the first preprocessing step, only volume information is retained, and volume values are grouped into **uniformly spaced bins** to reduce noise and discretize the data.

(b) Pairwise **mutual information (MI)** is then computed between all volume levels to capture dependencies—both within and across bid and ask sides. The resulting MI matrix reflects the **strength of informational connections** between volume levels.


(c) Finally, the TMFG is constructed using the MI matrix as input, yielding a sparse and interpretable network. In the TMFG, **nodes represent volume levels**, and **edges capture significant dependencies**.

v_1^{ask}	v_1^{bid}	v_{10}^{ask}	v_{10}^{bid}
138	135	89	201
61	135	89	201
61	135	89	203
61	135	289	203
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮
⋮	⋮	⋮	⋮

LOB volumes



Mutual information



IFN (TMFG)

Information Filtering + Deep Learning

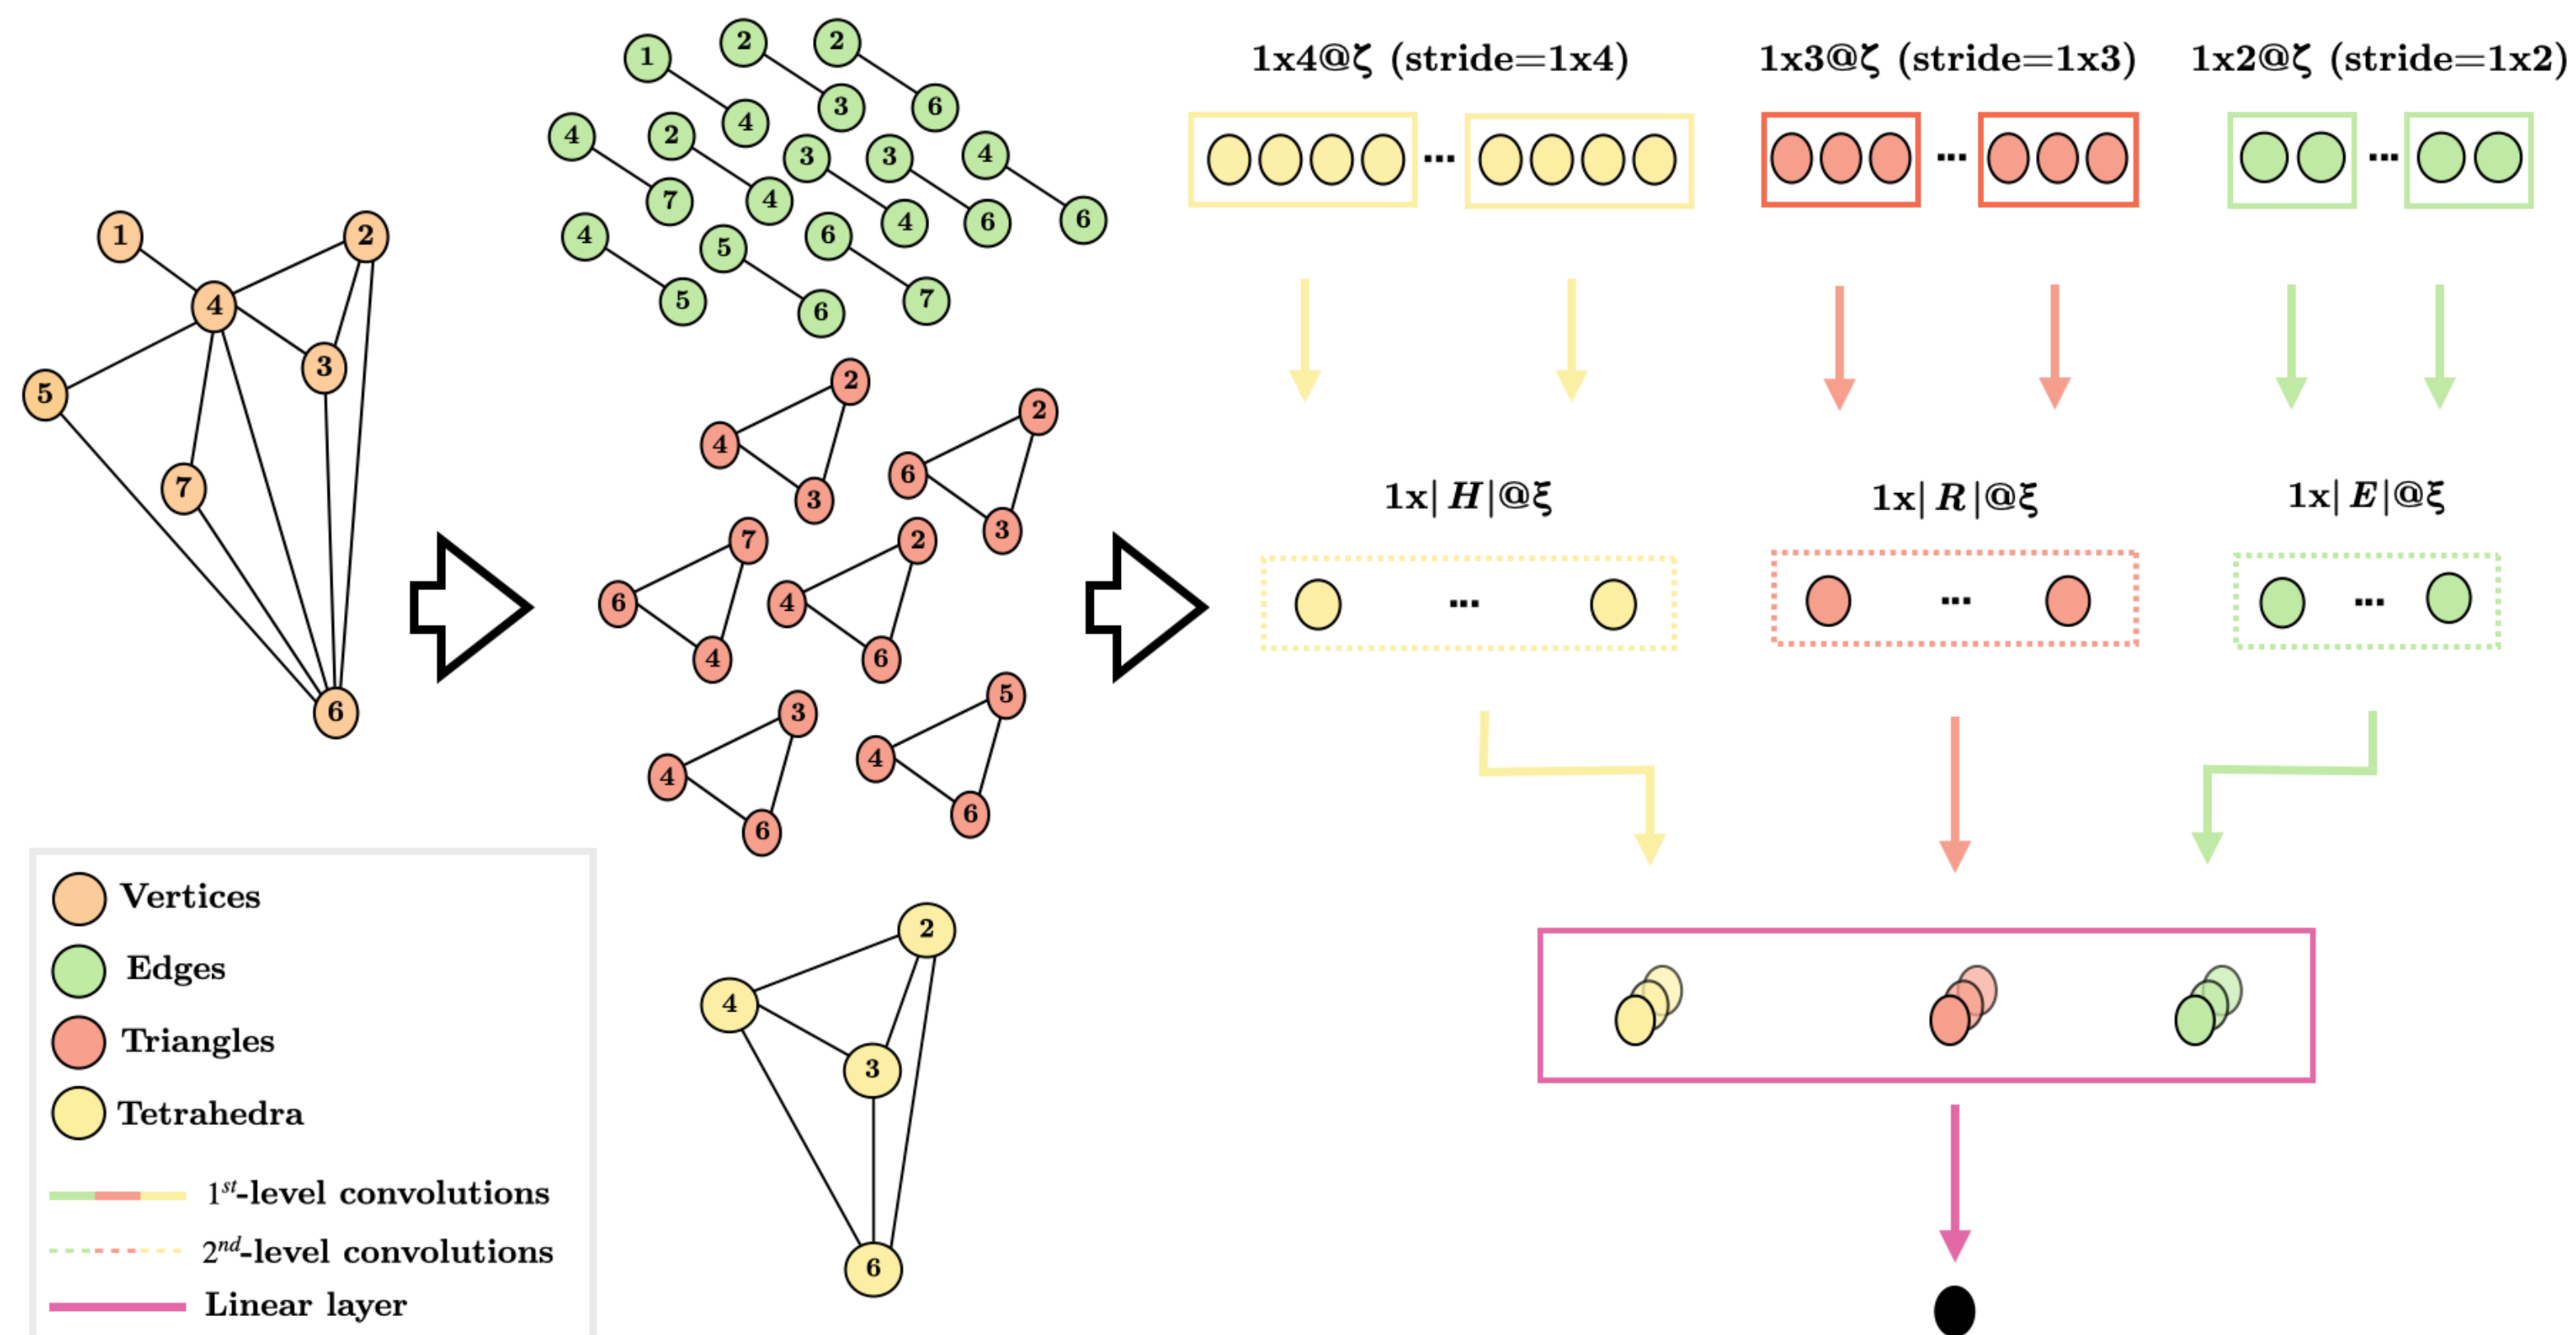


Wang, Y., Briola, A., & Aste, T.,
Homological Neural Networks: A Sparse
Architecture for Multivariate Complexity. In
*Topological, Algebraic and Geometric Learning
Workshops 2023* (pp. 228-241). PMLR.

HNN

(i) We start from a graph representing the dependency structures of features; (ii) we re-arrange the network's representation to highlight the underlying simplicial structures (i.e. edges, triangles, tetrahedra); and (iii) we end up with a layered representation, which takes into account higher order sub-structures and their interconnections, and can be converted into a computational unit (i.e. a sparse MLP).

Information Filtering + Deep Learning




Briola A., Wang Y., Bartolucci S., Aste T.
 Homological Convolutional Neural
 Networks, *Proceedings of the 2nd NeurIPS
 Workshop on Symmetry and Geometry in
 Neural Representations*, PMLR
 228:206-231, 2024.

HCNN

The model uses a chordal graph to capture feature dependencies, identifies simplicial structures (edges, triangles, tetrahedra), and processes them through two levels of convolutions — first to extract local non-linear relationships, then to uncover higher-order patterns — before applying a final linear layer to generate predictions.

Information Filtering + Deep Learning + LOB data



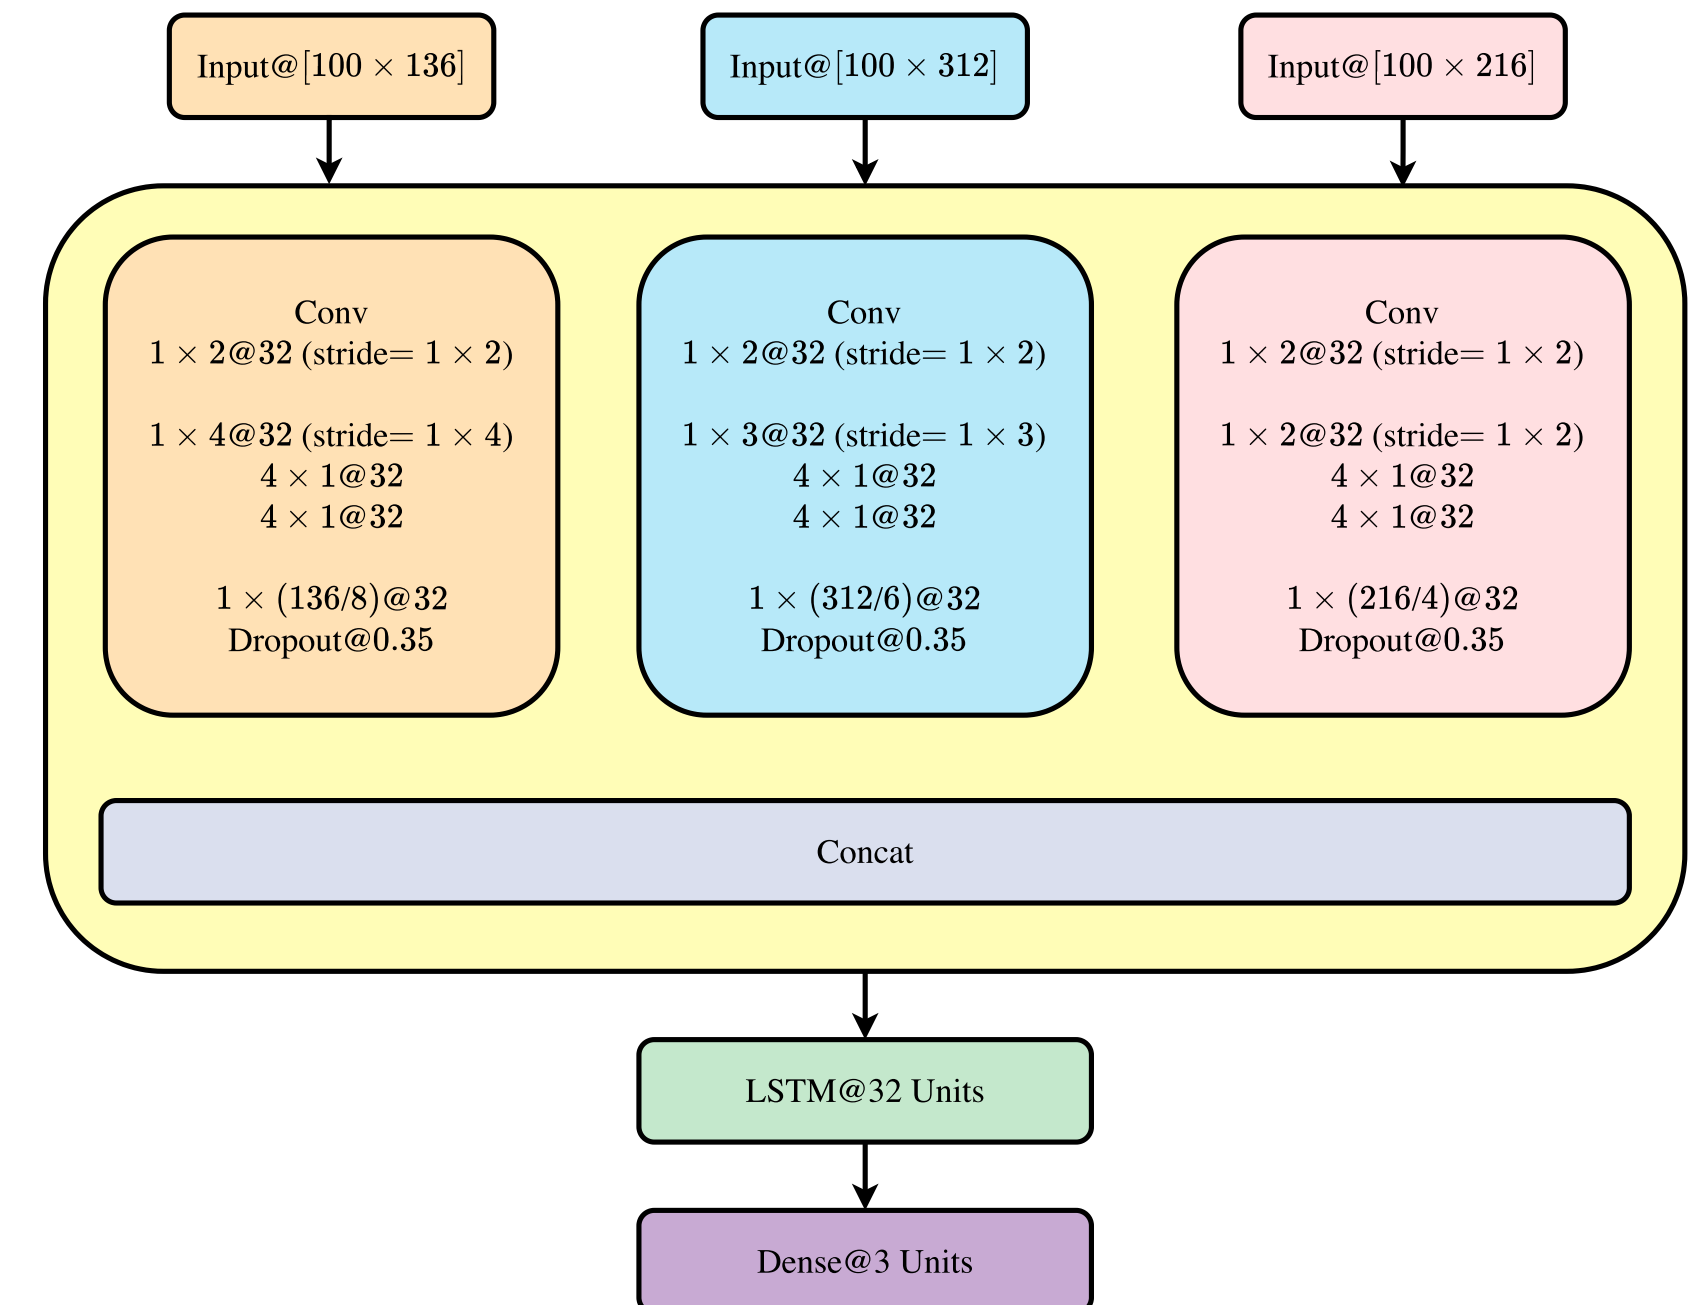
a)

$$\begin{bmatrix} p_4^{\text{ask}}(\tau_{-99}), v_4^{\text{ask}}(\tau_{-99}), p_3^{\text{bid}}(\tau_{-99}), v_3^{\text{bid}}(\tau_{-99}), p_4^{\text{bid}}(\tau_{-99}), v_4^{\text{bid}}(\tau_{-99}), p_5^{\text{bid}}(\tau_{-99}), v_5^{\text{bid}}(\tau_{-99}), \dots, p_3^{\text{ask}}(\tau_{-99}), v_3^{\text{ask}}(\tau_{-99}), p_4^{\text{ask}}(\tau_{-99}), v_4^{\text{ask}}(\tau_{-99}), p_5^{\text{ask}}(\tau_{-99}), v_5^{\text{ask}}(\tau_{-99}), p_4^{\text{bid}}(\tau_{-99}), v_4^{\text{bid}}(\tau_{-99}) \\ \vdots \\ p_4^{\text{ask}}(\tau_0), v_4^{\text{ask}}(\tau_0), p_3^{\text{bid}}(\tau_0), v_3^{\text{bid}}(\tau_0), p_4^{\text{bid}}(\tau_0), v_4^{\text{bid}}(\tau_0), p_5^{\text{bid}}(\tau_0), v_5^{\text{bid}}(\tau_0), \dots, p_3^{\text{ask}}(\tau_0), v_3^{\text{ask}}(\tau_0), p_4^{\text{ask}}(\tau_0), v_4^{\text{ask}}(\tau_0), p_5^{\text{ask}}(\tau_0), v_5^{\text{ask}}(\tau_0), p_4^{\text{bid}}(\tau_0), v_4^{\text{bid}}(\tau_0) \end{bmatrix}$$

$$\begin{bmatrix} p_4^{\text{ask}}(\tau_{-99}), v_4^{\text{ask}}(\tau_{-99}), p_2^{\text{bid}}(\tau_{-99}), v_2^{\text{bid}}(\tau_{-99}), p_3^{\text{bid}}(\tau_{-99}), v_3^{\text{bid}}(\tau_{-99}), \dots, p_3^{\text{ask}}(\tau_{-99}), v_3^{\text{ask}}(\tau_{-99}), p_4^{\text{ask}}(\tau_{-99}), v_4^{\text{ask}}(\tau_{-99}), p_4^{\text{ask}}(\tau_{-99}), v_4^{\text{ask}}(\tau_{-99}) \\ \vdots \\ p_4^{\text{ask}}(\tau_0), v_4^{\text{ask}}(\tau_0), p_2^{\text{bid}}(\tau_0), v_2^{\text{bid}}(\tau_0), p_3^{\text{bid}}(\tau_0), v_3^{\text{bid}}(\tau_0), \dots, p_3^{\text{ask}}(\tau_0), v_3^{\text{ask}}(\tau_0), p_4^{\text{ask}}(\tau_0), v_4^{\text{ask}}(\tau_0), p_4^{\text{ask}}(\tau_0), v_4^{\text{ask}}(\tau_0) \end{bmatrix}$$

$$\begin{bmatrix} p_4^{\text{ask}}(\tau_{-99}), v_4^{\text{ask}}(\tau_{-99}), p_2^{\text{bid}}(\tau_{-99}), v_2^{\text{bid}}(\tau_{-99}), \dots, p_3^{\text{bid}}(\tau_{-99}), v_3^{\text{bid}}(\tau_{-99}), p_4^{\text{bid}}(\tau_{-99}), v_4^{\text{bid}}(\tau_{-99}) \\ \vdots \\ p_4^{\text{ask}}(\tau_0), v_4^{\text{ask}}(\tau_0), p_2^{\text{bid}}(\tau_0), v_2^{\text{bid}}(\tau_0), \dots, p_3^{\text{bid}}(\tau_0), v_3^{\text{bid}}(\tau_0), p_4^{\text{bid}}(\tau_0), v_4^{\text{bid}}(\tau_0) \end{bmatrix}$$

b)



HLOB: To prepare the inputs for the model: (i) for each timestamp in the input's temporal dimension, we flatten each of the sets; (ii) we incorporate the corresponding price levels' data into each representative of these three new input sets.

Testing against SOTA models

	Tsantekidis et al. (2017a) CNN1 (2017)	Tsantekidis et al. (2020) CNN2 (2020)	Guo and Chen (2023) DLA (2022)	Vaswani et al. (2017) Transformer (2017)	Liu et al. (2023) iTransformer (2023)	Briola et al. LobTransformer (2024)	Tran et al. (2021) BinTabl (2021)	Tran et al. (2021) BinCTabl (2021)	Zhang, Zohren, and Roberts (2019) DeepLOB (2019)	Briola et al. HLOB (2024)
original code availability	✗	✗	✗	-	-	-	✗	✗	✓	-
n. trainable parameters	3.5×10^4	2.8×10^5	2.2×10^5	1.1×10^5	1.1×10^5	2.0×10^6	6.6×10^3	2.2×10^4	1.4×10^5	1.8×10^5
inference time (ms)	0.07	0.14	0.15	0.16	0.15	0.29	0.19	0.13	0.16	0.16

10 State-Of-The-Art models

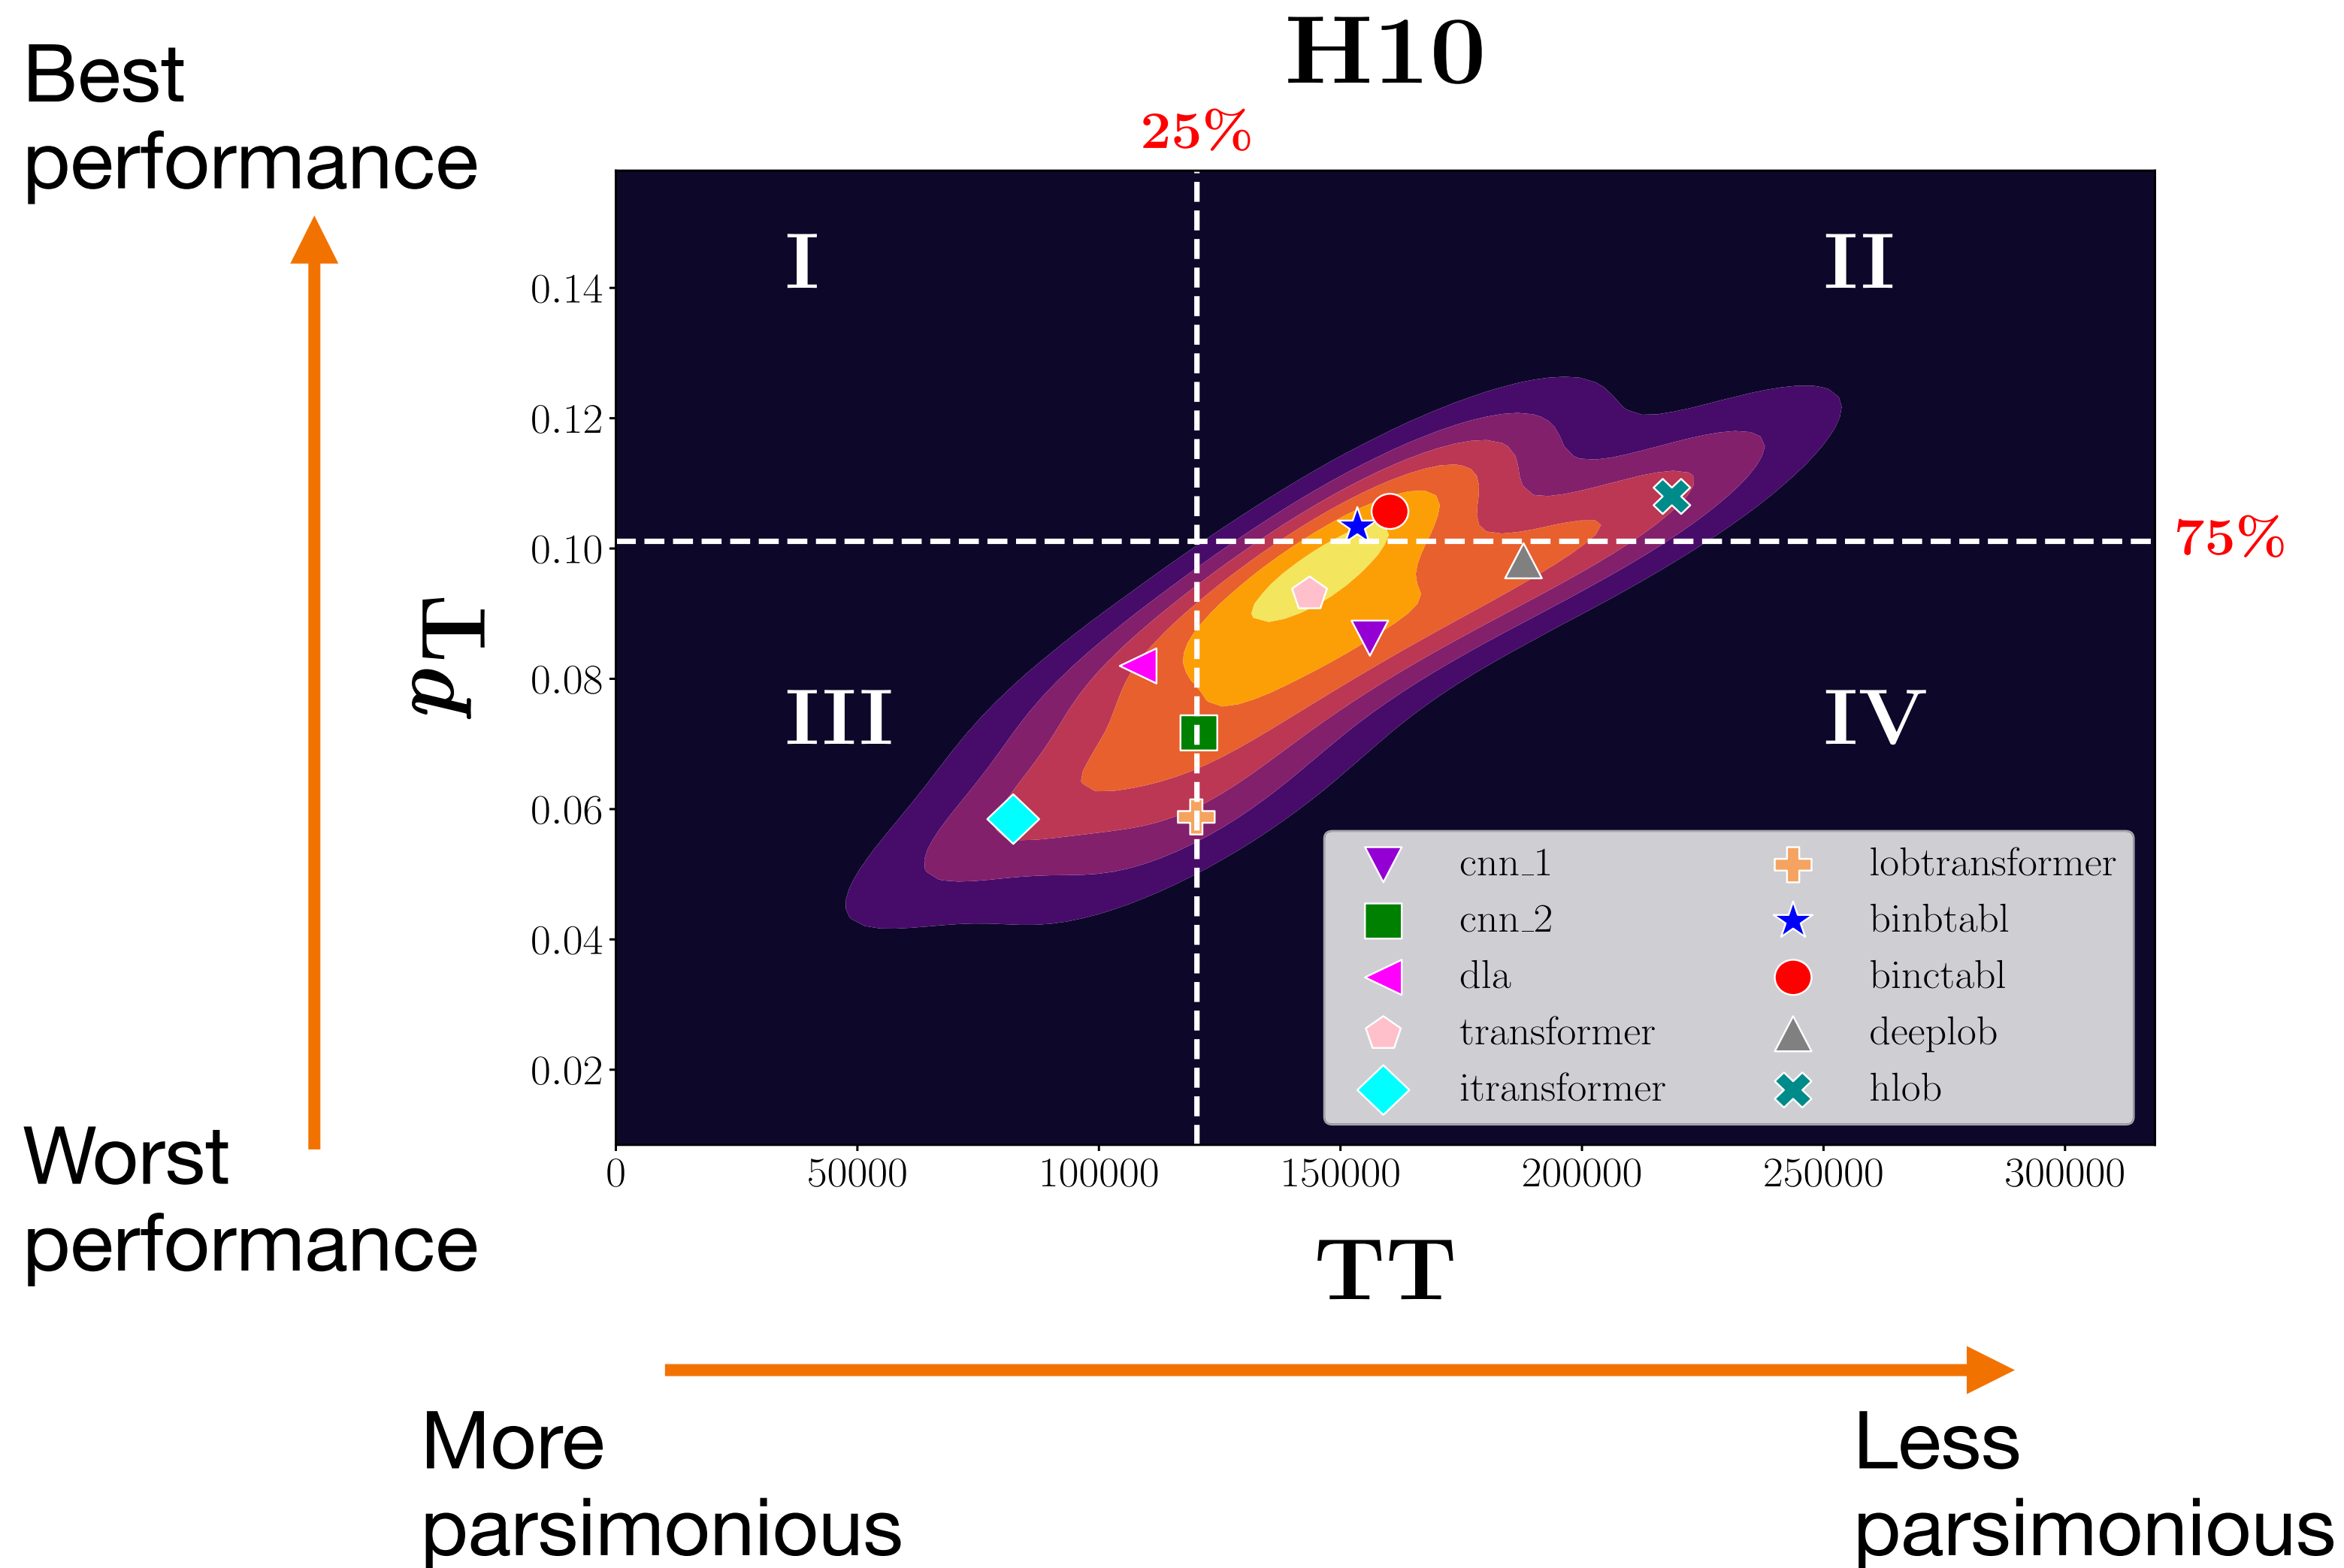
Small-tick

Medium-tick

Large tick

	H10																					Deep LOB			HLOB					
	cnn1			cnn2			dla			transformer			itransformer			lobtransformer			binbtabl			binctabl			F1	MCC	p_T	F1	MCC	p_T
	F1	MCC	p_T	F1	MCC	p_T	F1	MCC	p_T	F1	MCC	p_T	F1	MCC	p_T	F1	MCC	p_T	F1	MCC	p_T	F1	MCC	p_T						
CHTR	0.39	0.11	0.06	0.38	0.09	0.05	0.39	0.11	0.04	0.40	0.12	0.06	0.35	0.05	0.04	0.31	0.07	0.04	0.42	0.15	0.06	0.43	0.16	0.06	0.39	0.10	0.05	0.43	0.17	0.06
GOOG	0.42	0.16	0.04	0.42	0.16	0.04	0.39	0.13	0.03	0.41	0.15	0.04	0.27	0.04	0.04	0.44	0.18	0.05	0.45	0.18	0.08	0.46	0.20	0.08	0.45	0.19	0.04	0.46	0.21	0.05
GS	0.36	0.10	0.09	0.29	0.06	0.06	0.38	0.09	0.08	0.38	0.12	0.10	0.31	0.05	0.02	0.16	0.00	0.00	0.40	0.15	0.09	0.41	0.15	0.10	0.34	0.10	0.08	0.41	0.17	0.12
IBM	0.36	0.09	0.12	0.36	0.08	0.08	0.35	0.08	0.11	0.35	0.11	0.11	0.30	0.06	0.02	0.30	0.05	0.05	0.36	0.10	0.14	0.37	0.10	0.14	0.38	0.11	0.13	0.40	0.13	0.14
MCD	0.37	0.08	0.10	0.35	0.08	0.08	0.38	0.09	0.10	0.38	0.10	0.11	0.31	0.04	0.02	0.28	0.04	0.01	0.39	0.11	0.12	0.40	0.11	0.12	0.41	0.12	0.11	0.41	0.13	0.13
NVDA	0.31	0.06	0.07	0.24	0.00	0.00	0.33	0.07	0.06	0.36	0.08	0.10	0.24	0.02	0.03	0.22	0.02	0.00	0.41	0.13	0.13	0.41	0.13	0.14	0.34	0.08	0.09	0.40	0.12	0.14
AAPL	0.42	0.16	0.13	0.39	0.14	0.09	0.41	0.14	0.12	0.41	0.16	0.13	0.35	0.09	0.07	0.39	0.17	0.10	0.41	0.15	0.15	0.42	0.16	0.16	0.43	0.17	0.15	0.42	0.18	0.15
ABBV	0.38	0.13	0.12	0.39	0.12	0.10	0.39	0.13	0.11	0.40	0.15	0.13	0.31	0.06	0.03	0.33	0.10	0.04	0.36	0.13	0.13	0.37	0.14	0.13	0.39	0.13	0.13	0.42	0.18	0.14
PM	0.35	0.08	0.09	0.39	0.10	0.09	0.37	0.08	0.08	0.36	0.09	0.09	0.29	0.02	0.02	0.28	0.04	0.04	0.36	0.10	0.12	0.36	0.11	0.13	0.36	0.12	0.13	0.39	0.13	0.13
BAC	0.43	0.23	0.04	0.38	0.21	0.06	0.38	0.23	0.04	0.44	0.28	0.05	0.36	0.18	0.09	0.46	0.29	0.05	0.45	0.27	0.06	0.45	0.28	0.07	0.46	0.30	0.07	0.47	0.32	0.06
CSCO	0.47	0.29	0.08	0.50	0.29	0.09	0.47	0.28	0.07	0.48	0.29	0.08	0.41	0.19	0.11	0.47	0.27	0.08	0.45	0.28	0.08	0.44	0.27	0.07	0.49	0.30	0.08	0.50	0.33	0.08
KO	0.47	0.26	0.08	0.47	0.27	0.10	0.46	0.28	0.09	0.47	0.28	0.09	0.39	0.17	0.10	0.48	0.30	0.10	0.45	0.28	0.10	0.43	0.27	0.10	0.48	0.28	0.10	0.49	0.31	0.10
ORCL	0.47	0.27	0.10	0.45	0.26	0.09	0.45	0.26	0.10	0.48	0.30	0.11	0.38	0.16	0.07	0.48	0.31	0.11	0.46	0.27	0.10	0.44	0.26	0.10	0.49	0.32	0.11	0.48	0.32	0.11
PFE	0.43	0.24	0.09	0.42	0.24	0.09	0.43	0.25	0.09	0.44	0.25	0.09	0.36	0.17	0.11	0.45	0.27	0.09	0.47	0.29	0.09	0.46	0.28	0.10	0.46	0.27	0.09	0.49	0.32	0.10
VZ	0.47	0.23	0.08	0.42	0.17	0.07	0.47	0.26	0.09	0.47	0.27	0.10	0.39	0.16	0.10	0.46	0.25	0.10	0.45	0.26	0.10	0.41	0.24	0.10	0.49	0.28	0.11	0.46	0.28	0.10

Testing against SOTA models



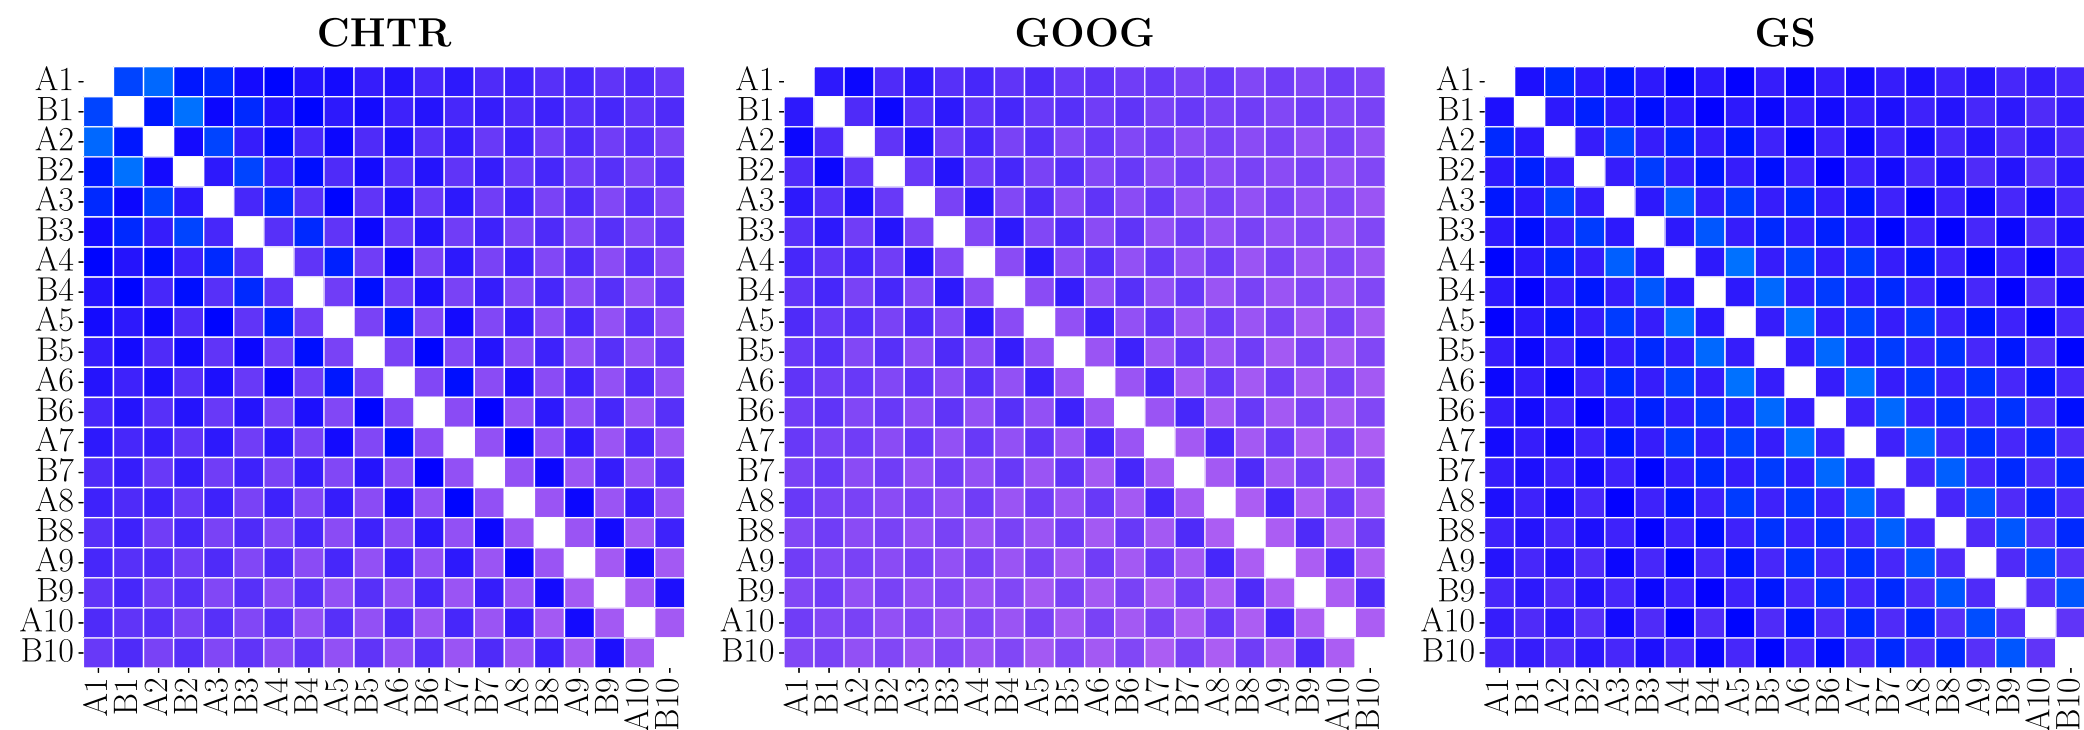
HLOB is not parsimonious, but achieves good performances for all stocks at H10, and for large tick stocks at longer time horizons.

Binbtbl better at long time horizon thanks to dual attention

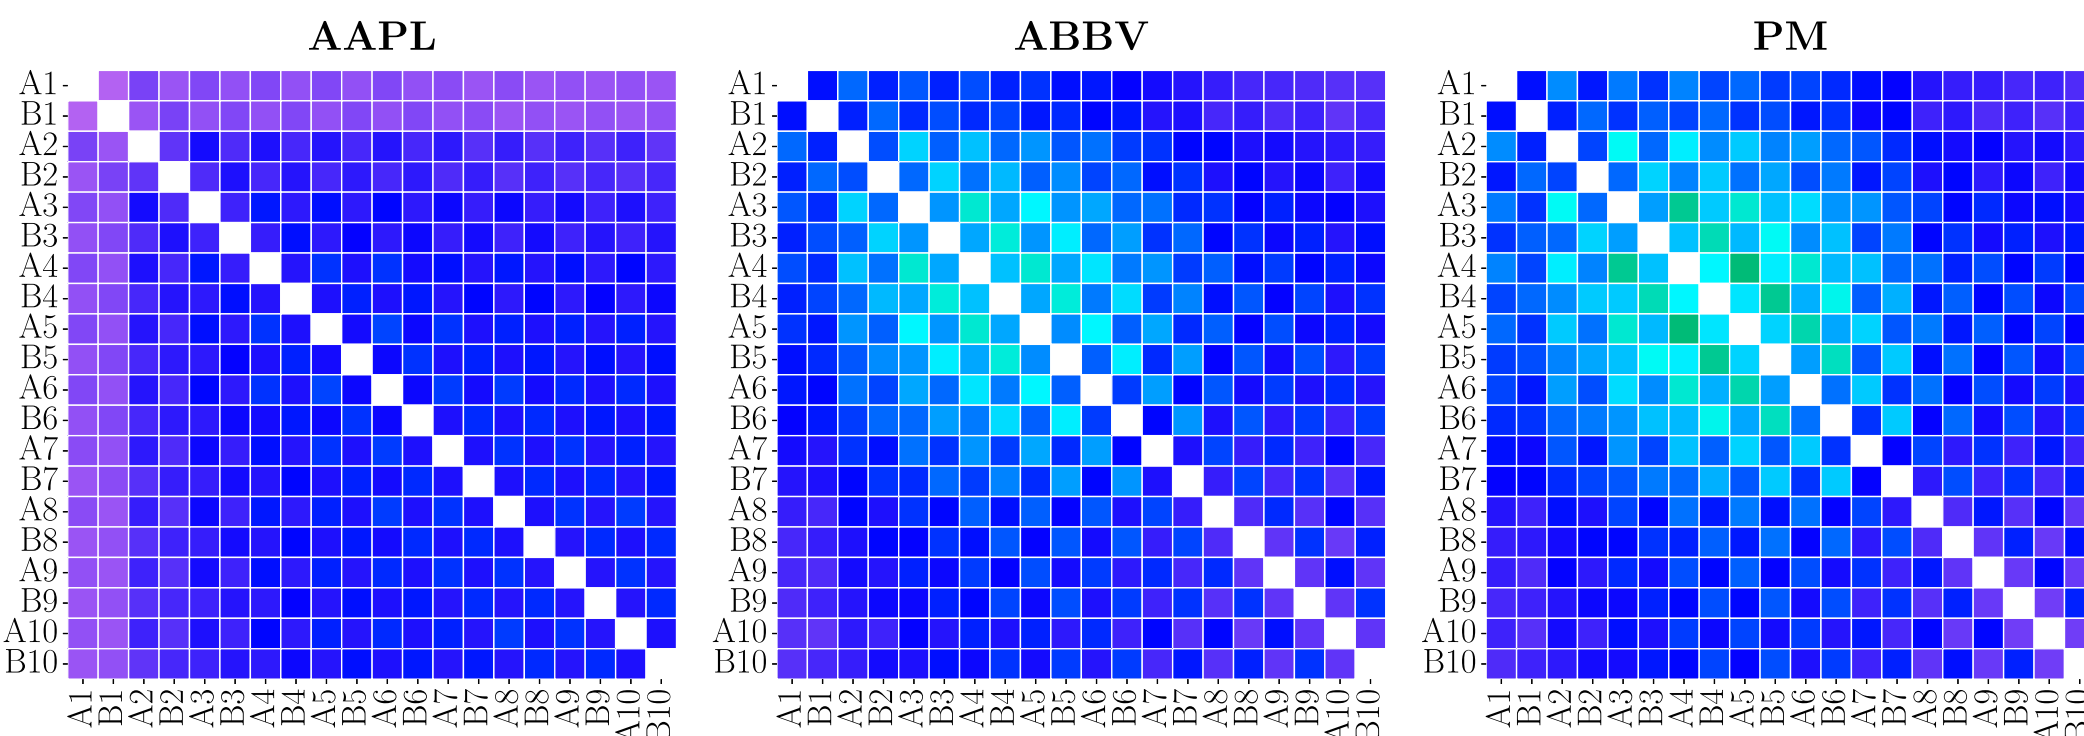
Transformers have consistently worse performance

Spatial information


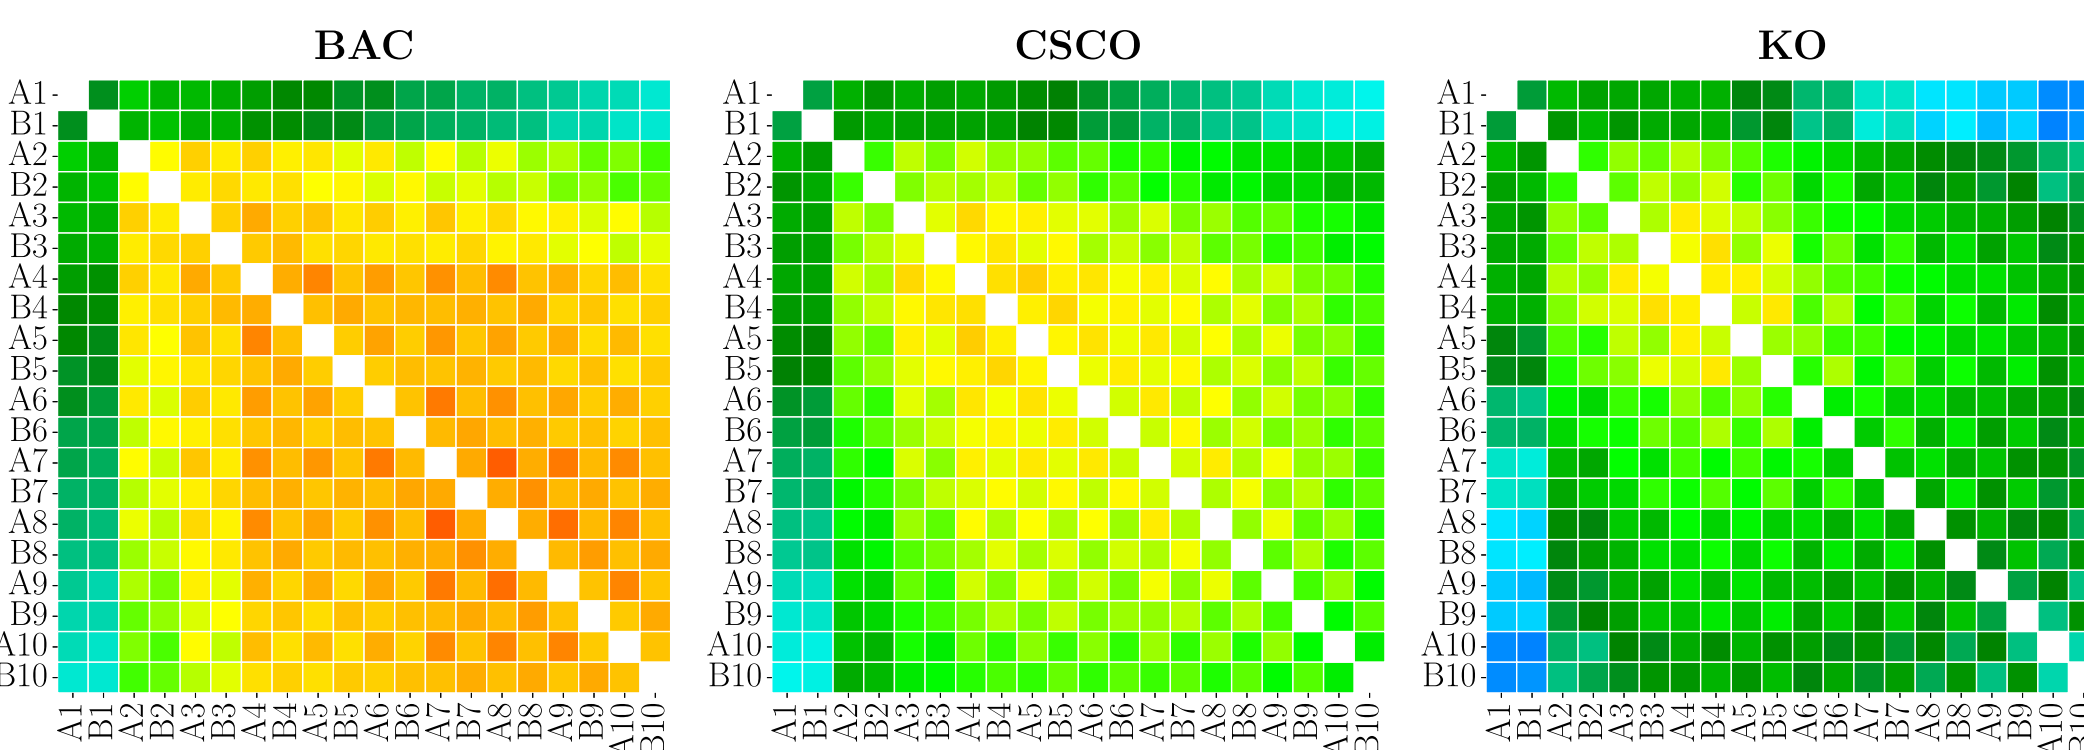
Small-tick stocks



Medium-tick stocks



Large tick stocks



Average MI matrices over the 3-year analysis period

- Different spatial dependency structures emerge depending on the asset's tick size.
- Assets characterised by a clear hierarchical structure (i.e., large-tick stocks) have more chances to be correctly forecast by deep learning models.

Main take aways - part 2

Design of microstructurally-informed models

HLOB leverages topological structure and temporal dynamics to significantly improve LOB forecasting accuracy, especially for large-tick stocks where microstructure is more stable and informative.

Informational Content

The structural information content of the Limit Order Book varies across stock types and time horizons, making deep learning models like HLOB more or less effective depending on the underlying microstructure.

Outlooks

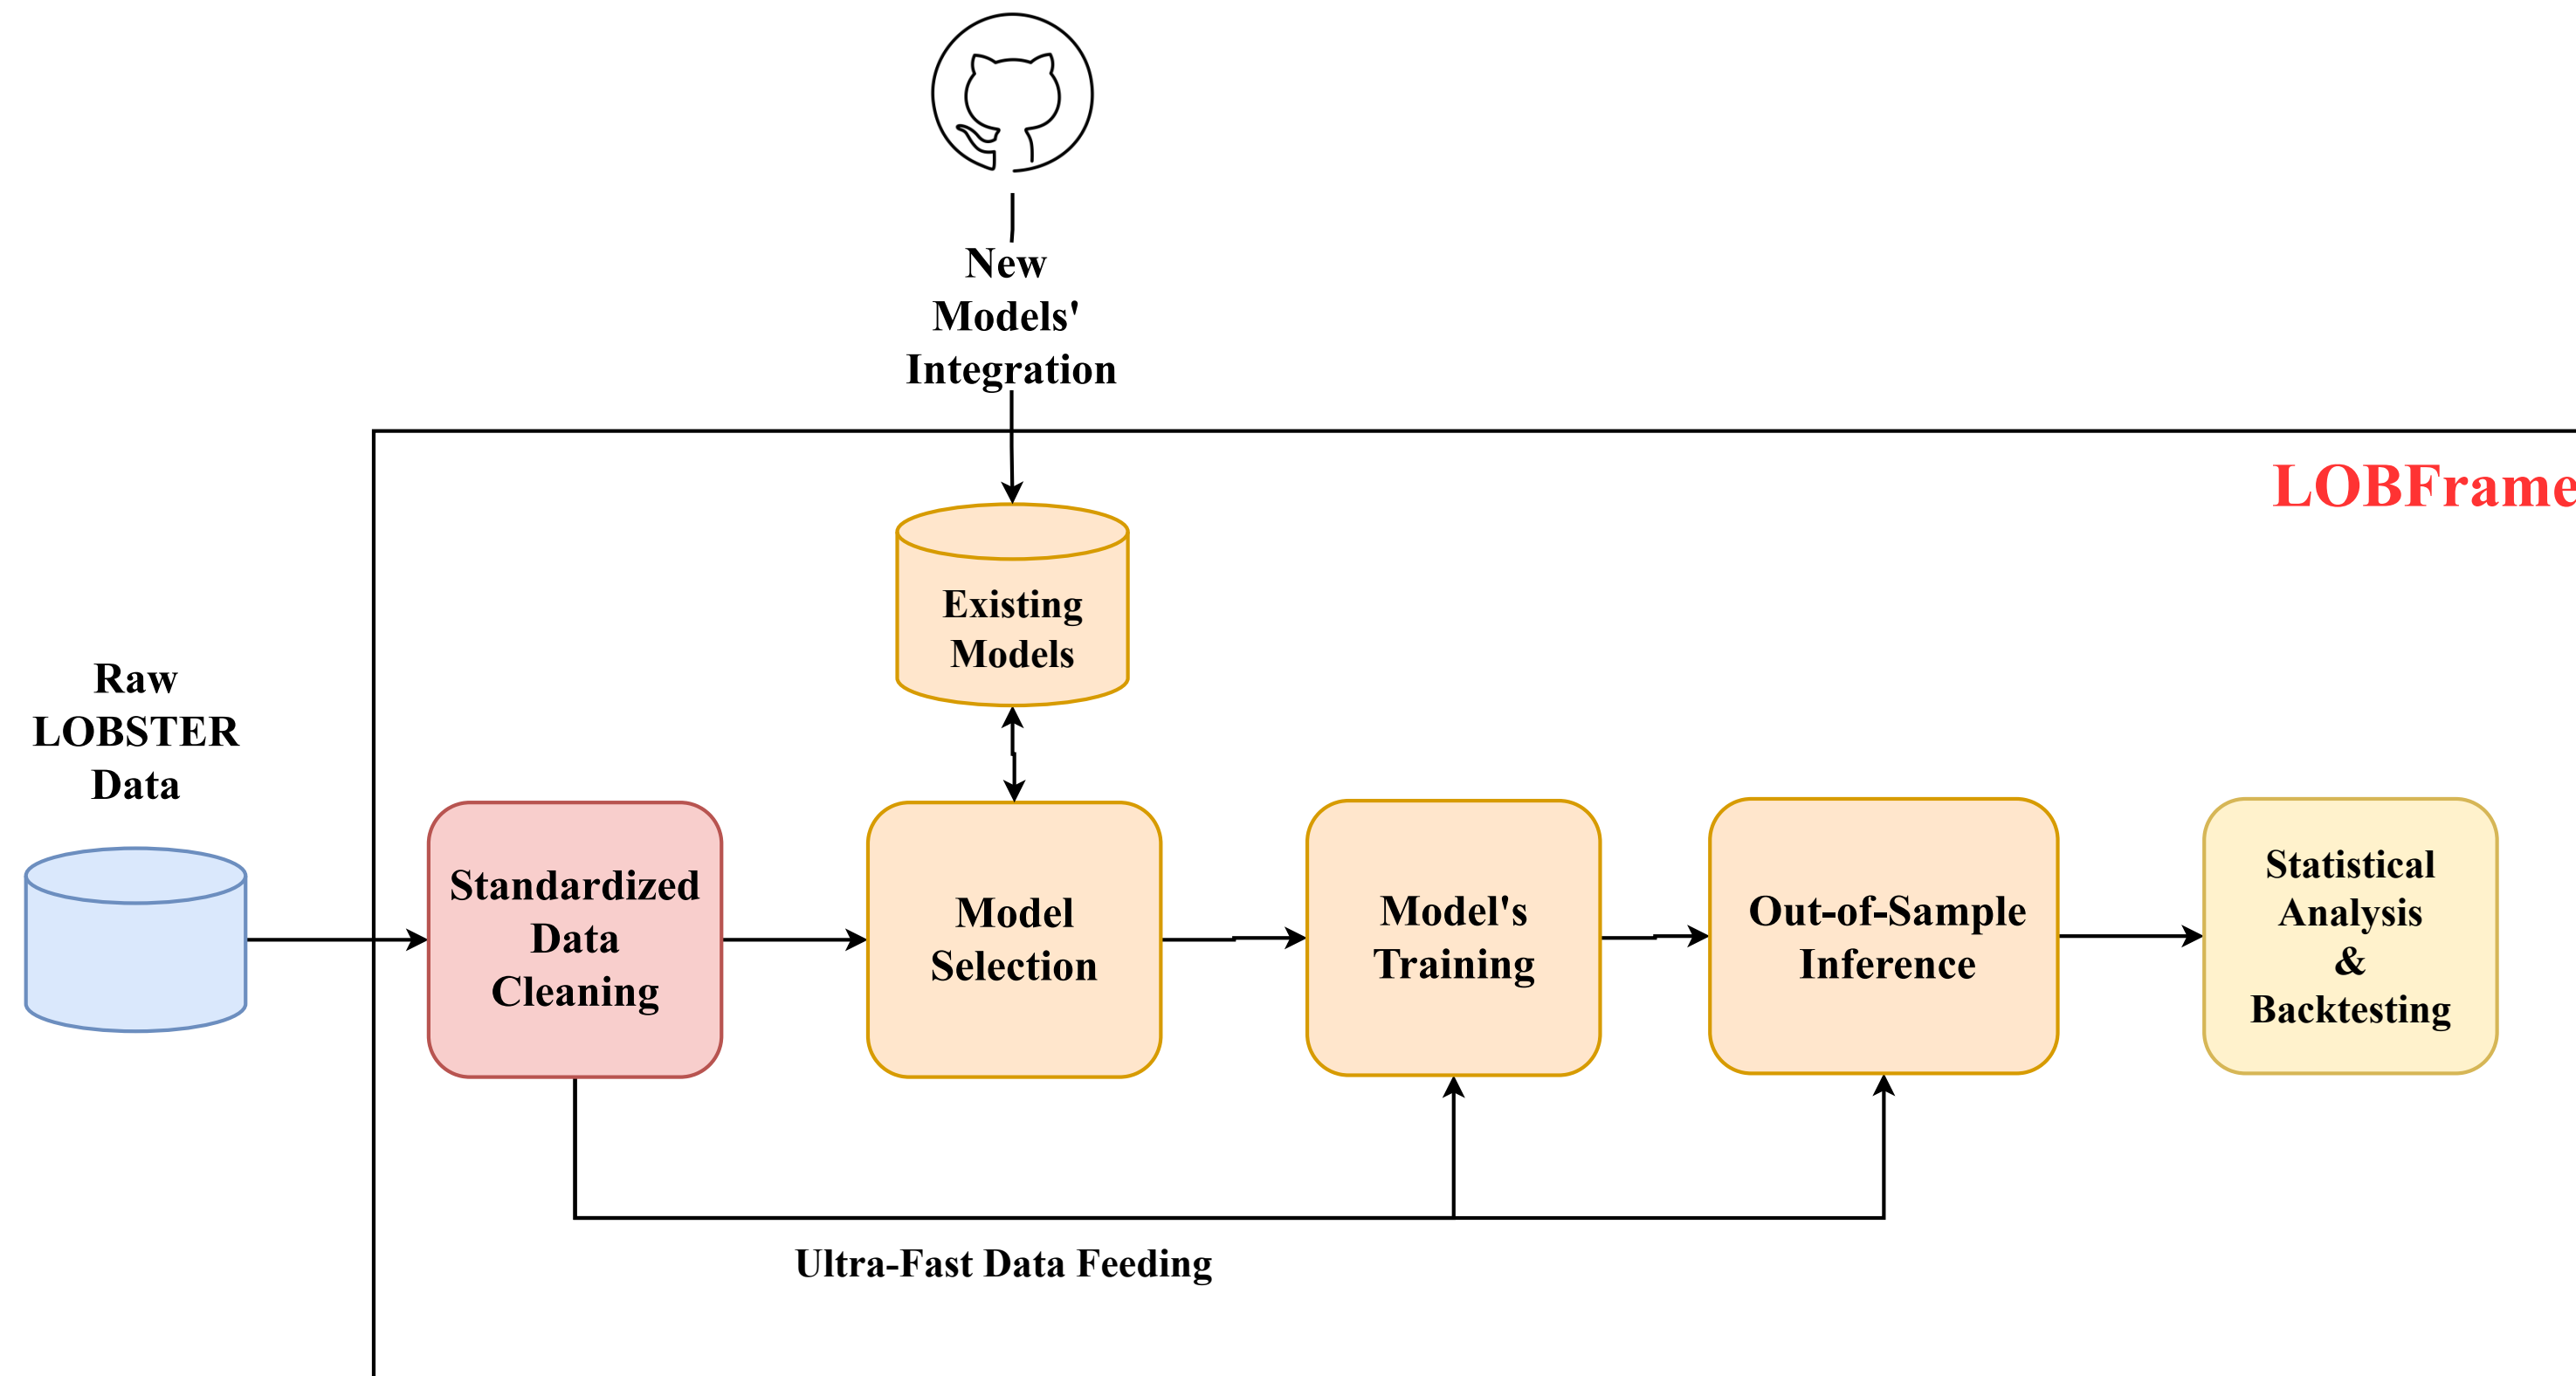
Research Next Steps

- Development of ad-hoc microstructurally-informed architectures to handle sparse LOB structures characterising small- to medium-tick stocks.
- Structured testing of different deep learning models designed to efficiently handle heterogeneous classes of stock.
- Cross-exchanges validation of results.
- Refining metrics to assess predictions' usability in real trading environments.

AI/Finance Community

- Go beyond benchmark datasets.
- Prioritise microstructural awareness in model design.
- Adopt trading-aware evaluation metrics.
- Ensure reproducibility and transparency.

LOBframe pipeline



LOBFrame is a modular, end-to-end framework designed to streamline the **Limit Order Book (LOB)** forecasting workflow.

Data Cleaning & Preprocessing

Standardizes raw LOB data for consistency and usability.

Model Selection & Integration

Supports plug-and-play with both state-of-the-art and custom forecasting models.

Ultra-fast Data Feeding

Ensures efficient model training and real-time or batch inference.

Comprehensive Evaluation

Enables robust statistical analysis, out-of-sample testing, and backtesting.

Code available on Github



Briola, A., Bartolucci, S., & Aste, T. (2025). Deep limit order book forecasting: a microstructural guide. *Quantitative Finance*, **25(7)**, 1101-1131.

Briola, A., Bartolucci, S., & Aste, T. (2025). HLOB–Information persistence and structure in limit order books. *Expert Systems with Applications*, **266**, 126078.

Briola A., Wang Y. , Bartolucci S., Aste T. Homological Convolutional Neural Networks Proceedings of the *2nd NeurIPS Workshop on Symmetry and Geometry in Neural Representations*, PMLR 228:206-231, 2024.

Aste, T. (2025). Information Filtering Networks: Theoretical Foundations, Generative Methodologies, and Real-World Applications. ArXiv preprint arXiv:2505.03812.

LOBframe



Thank you!

