ZANISTA

Monthly Review of Seminal Quant Finance Papers

15th Edition
September 2025

Available at:
zanista.ai/paperpal

For more info:
arman.khaledian@zanista.ai

ZANISTA

Table of Contetns (papers in August 2025)

Click on the topic card to jump to the first paper on that topic

Sign Up to Get This Newsletter Every Month

Click here to stay ahead of the curve in Quant Finance & AI/LLM research, and subscribe to our monthly newsletter: https://forms.gle/a7KpFxZN4xiz3ZaMA

ZANISTA

Kronos: A Foundation Model for the Language of Financial Markets

Yu Shi|Zongliang Fu|Shuo Chen|Bohan Zhao|Wei Xu|Changshui Zhang|Jian Li
Tsinghua University

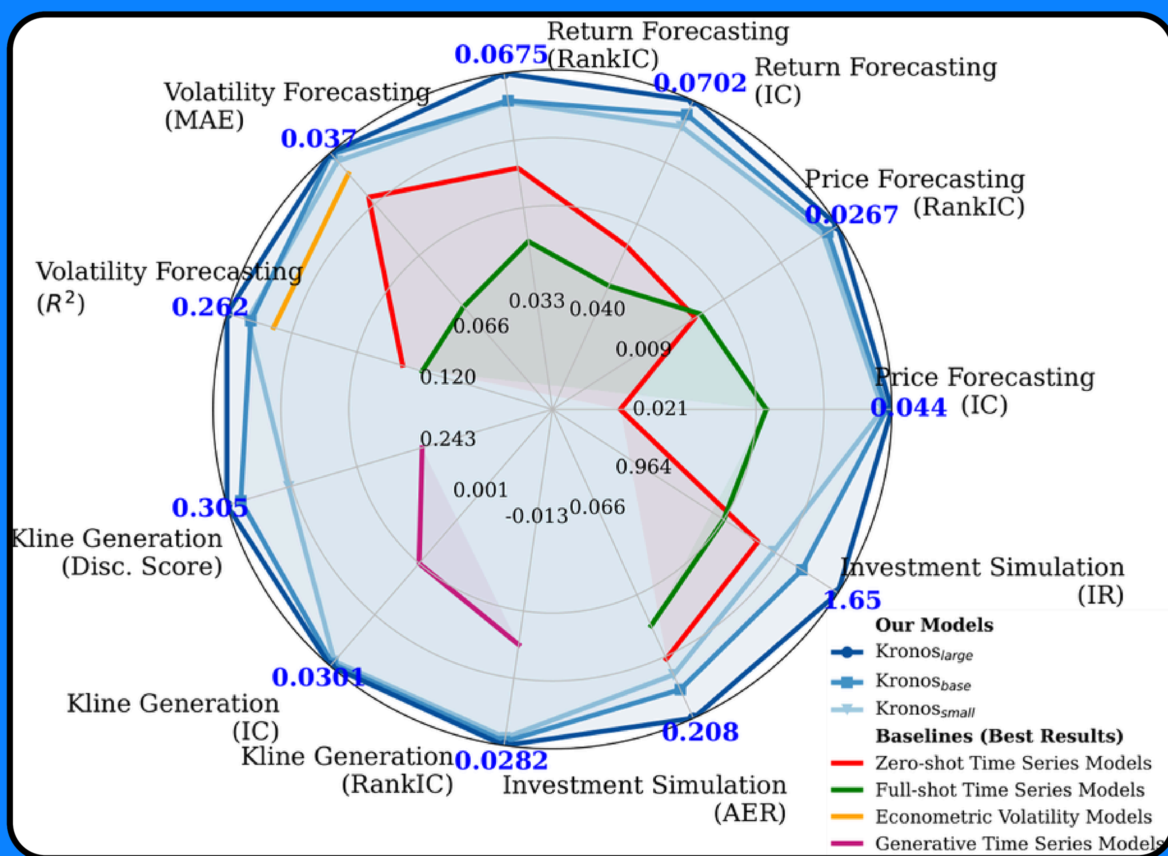

Kronos recasts candlesticks as a discrete language with strong results. Trained on 12 billion K lines from 45 exchanges and seven timeframes, it lifts RankIC 93 percent over the best foundation model, trims volatility MAE 9 percent, and raises synthetic data fidelity 22 percent. A Binary Spherical Quantization tokenizer and decoder Transformer drive coarse to fine autoregression. It leads backtests on excess return and information ratio. Limits include 512 context, OHLCVA, and no live trading.

AI, LLMs, or Machine Learning for Investment Management or Trading

Generative AI for Financial Data Synthesis

FinCast: A Foundation Model for Financial Time-Series Forecasting

Zhuohang Zhu, Haodong Chen, Qiang Qu, Vera Chung
The University of Sydney

FinCast is a 1B parameter decoder only transformer for financial forecasting, trained on 20B time points across crypto, forex, futures, stocks, and macro data. To handle non stationarity and cross domain, multi resolution shifts, it uses token level sparse Mixture of Experts, learnable frequency embeddings, and a point quantile loss that models uncertainty. It achieves strong zero shot accuracy without fine tuning, surpassing leading models on benchmarks and running efficiently.

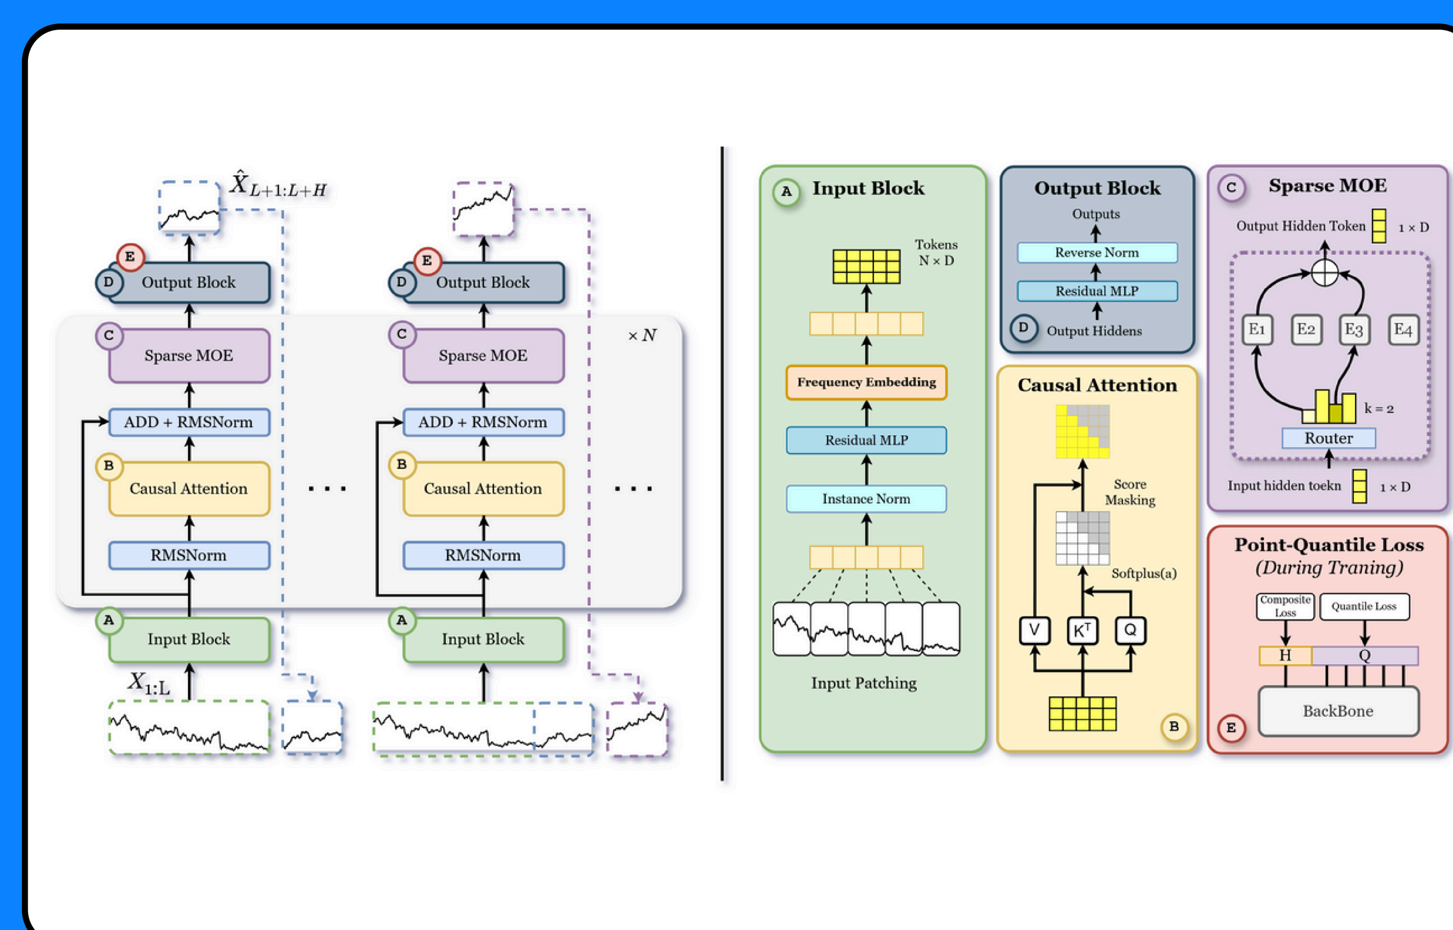

AI, LLMs, or Machine Learning for Investment Management or Trading

Structured Agentic Workflows for Financial Time-Series Modeling with LLMs and Reflective Feedback

Yihao Ang|Yifan Bao|Lei Jiang|Jiajie Tao|Anthony K. H. Tung|Lukasz Szpruch|Hao Ni
National University of Singapore University College London University of Edinburgh

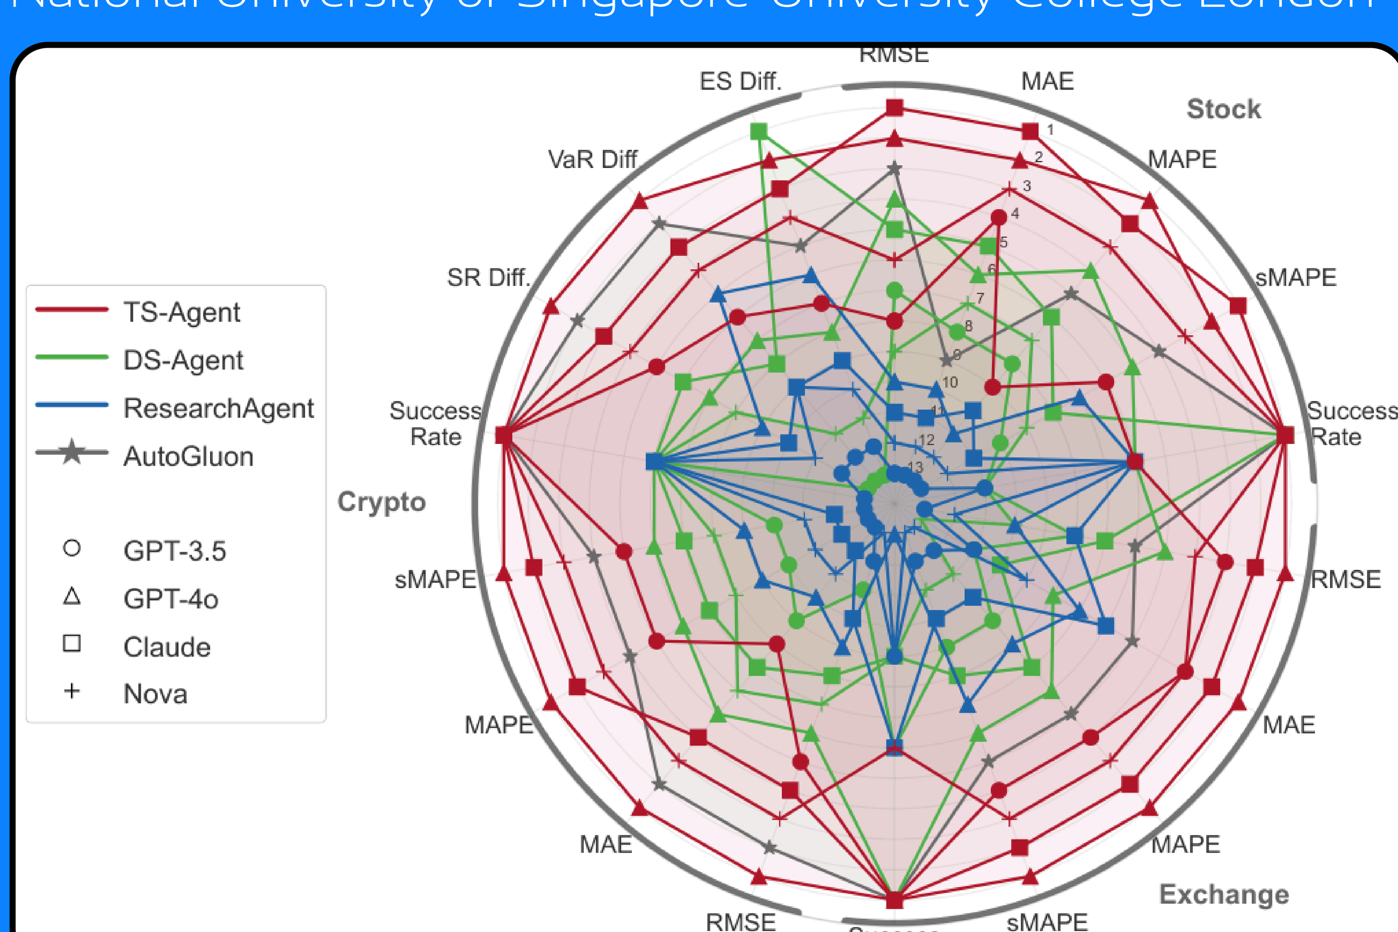

TS-Agent introduces a structured LLM workflow for financial time series that beats AutoML and agentic baselines on forecasting and generation, and remains consistent across LLM backbones. A two-stage loop selects models, then refines code using three knowledge banks and a chain-of-code-edits with experimental feedback. On Crypto, Exchange, and Stock it cuts RMSE 20 percent versus AutoGluon and up to 30 percent versus DS-Agent. Decisions are logged. Limits include dataset scope, bias, compute, and limited backtests.

AI, LLMs, or Machine Learning for Investment Management or Trading

AI Applications in Algorithmic Trading

ZANISTA

AlphaEval: A Comprehensive and Efficient Evaluation Framework for Formula Alpha Mining

Hongjun Ding|Binqi Chen|Jinsheng Huang|Taian Guo|Zhengyang Mao|Guoyi Shao|Lutong Zou|Luchen Liu|Ming Zhang
Peking University City University of New York Zhengren Quant

AlphaEval debuts a backtest free framework to score alphas on predictiveness, stability, noise resilience, financial logic, and diversity. It matches backtests and runs 25 percent faster. Contributions include a benchmark, three metrics for rank stability, shock robustness, and multicollinearity, plus an LLM logic score. On shares and S and P 500, scores mirror portfolios. Stability predicts lower turnover. High robustness cuts drawdowns. RL and GA excel at stability. LLMs lead predictability with robustness trade off.

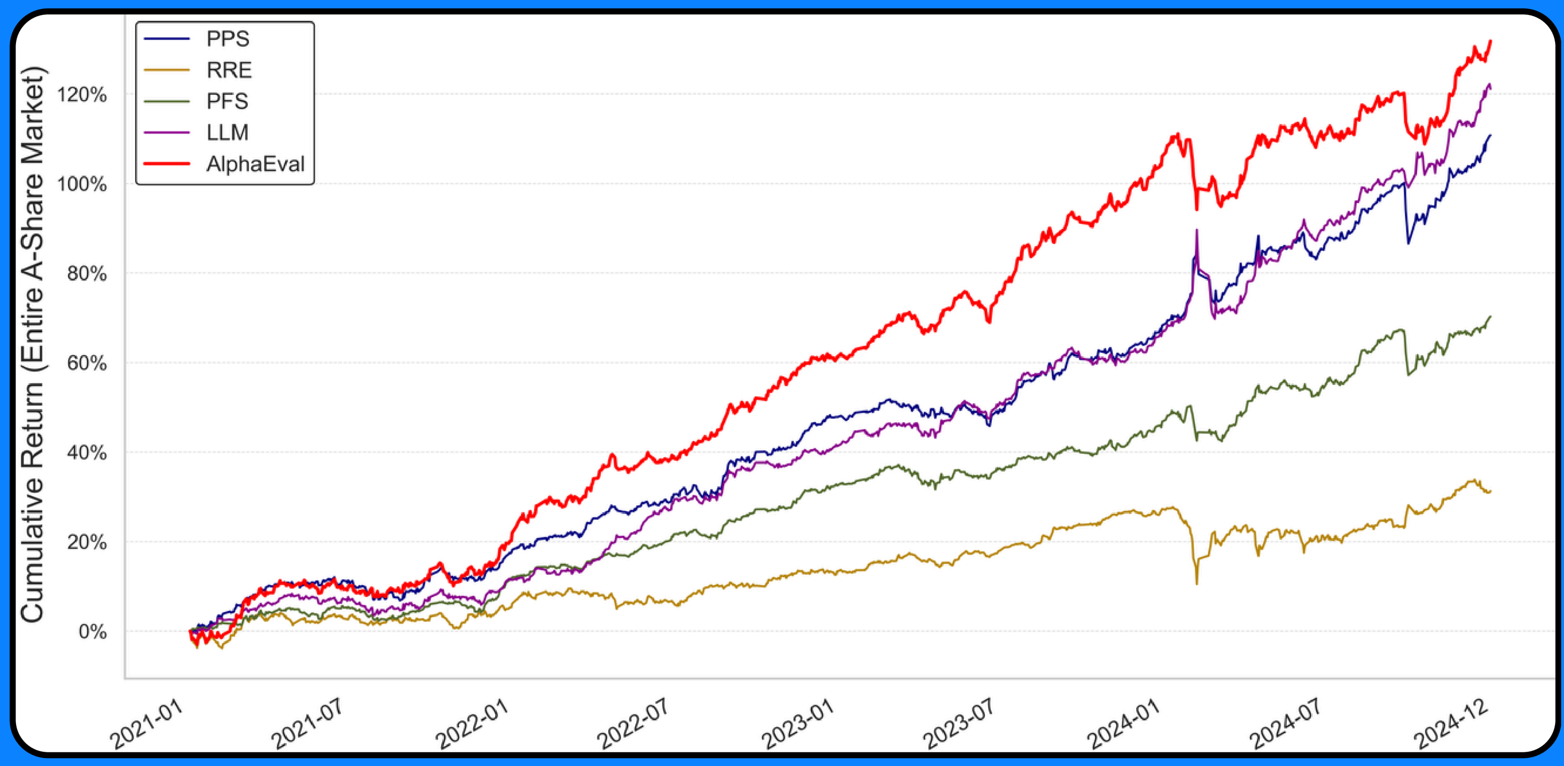

AI, LLMs, or Machine Learning for Investment Management or Trading | Financial Ratios & Quantitative Analysis in AI Models

Does generative AI facilitate investor Trading? Early evidence from ChatGPT outages

Q Cheng|P Lin|Y Zhao
Singapore Management University

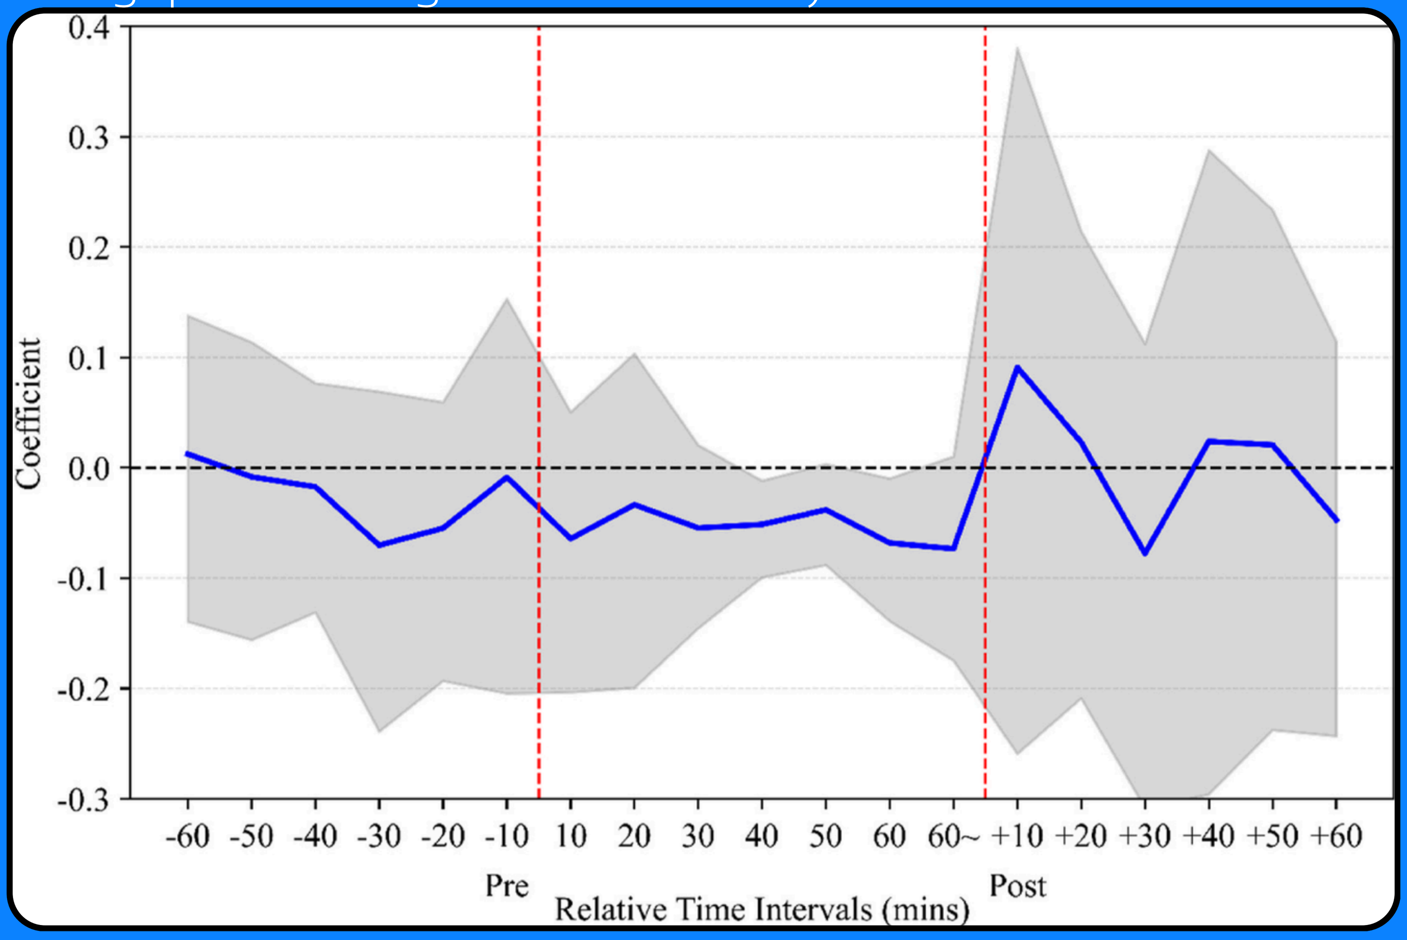

U.S. stock trading slowed during eight ChatGPT outages in 2023. Volume fell 10 percent, or 5.6 percent of a standard deviation, with larger drops around firm news and in stocks held by transient institutions. Retail and nonretail activity declined, more for nonretail. Odd-lot trades fell during OpenAI API outages. Price impact and return variance dropped and spreads narrowed, signaling less informed trading. Effects appeared 30 minutes in. An exposure measure predicts gains in price informativeness.

AI, LLMs, or Machine Learning for Investment Management or Trading | AI in High-Frequency Trading

A Financial Brain Scan of the LLM

Hui Chen|Antoine Didisheim|Luciano Somoza|Hanqing Tian
Massachusetts Institute of Technology National Bureau of Economic Research University of Melbourne

Researchers brain scan an LLM by inserting a sparse autoencoder that maps hidden states to concepts and enables steering. On 3.7 million Reuters news, Gemma-2-9B-IT forecasts next-day returns and builds long short portfolios. Interpretable embeddings beat last-layer baselines, Sharpe 5.51 vs 4.91, improving with more features. Clustering 17 economic groups shows Sentiment, Finance/Markets, Technical Analysis drive gains. Steering fixes optimism bias, lifts Sharpe 3.87 to 4.28. Limits include look-ahead, one model, one source, noisy labels.

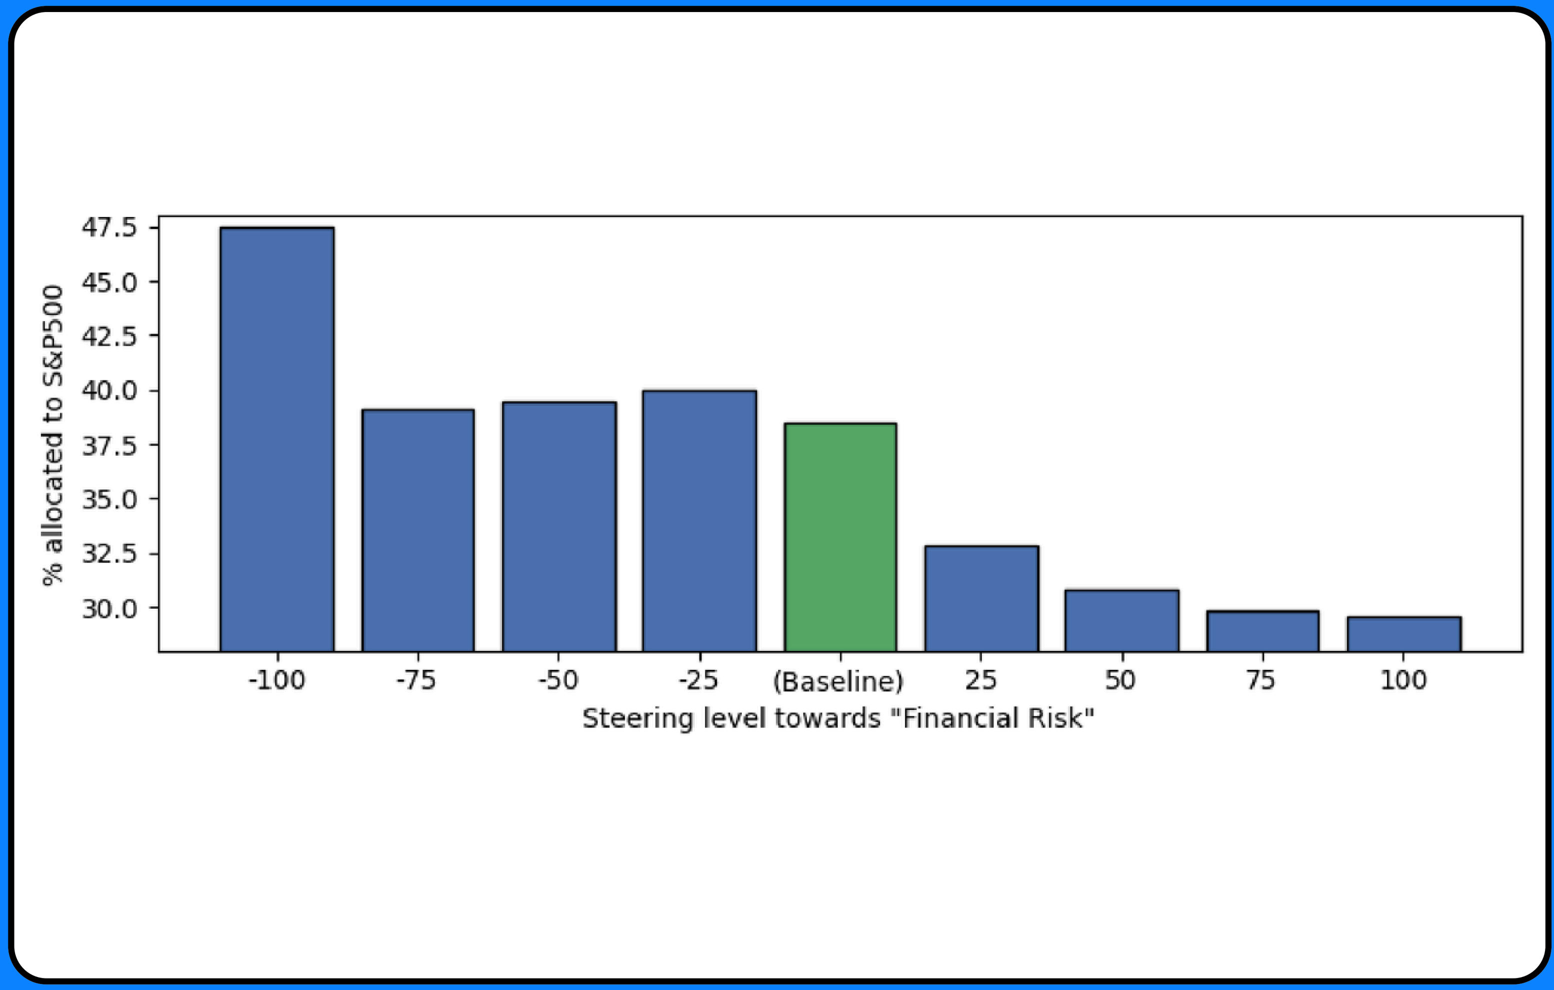

AI, LLMs, or Machine Learning for Investment Management or Trading | Natural Language Processing in Financial Markets

A Review of LLM Agent Applications in Finance and Banking

Devesh Batra|Conor B. Hamill|John Hartley|Ramin Okhrati|Dale Seddon|Harvey Miller|Raad Khraishi|Greig A. Cowan
NatWest University College London

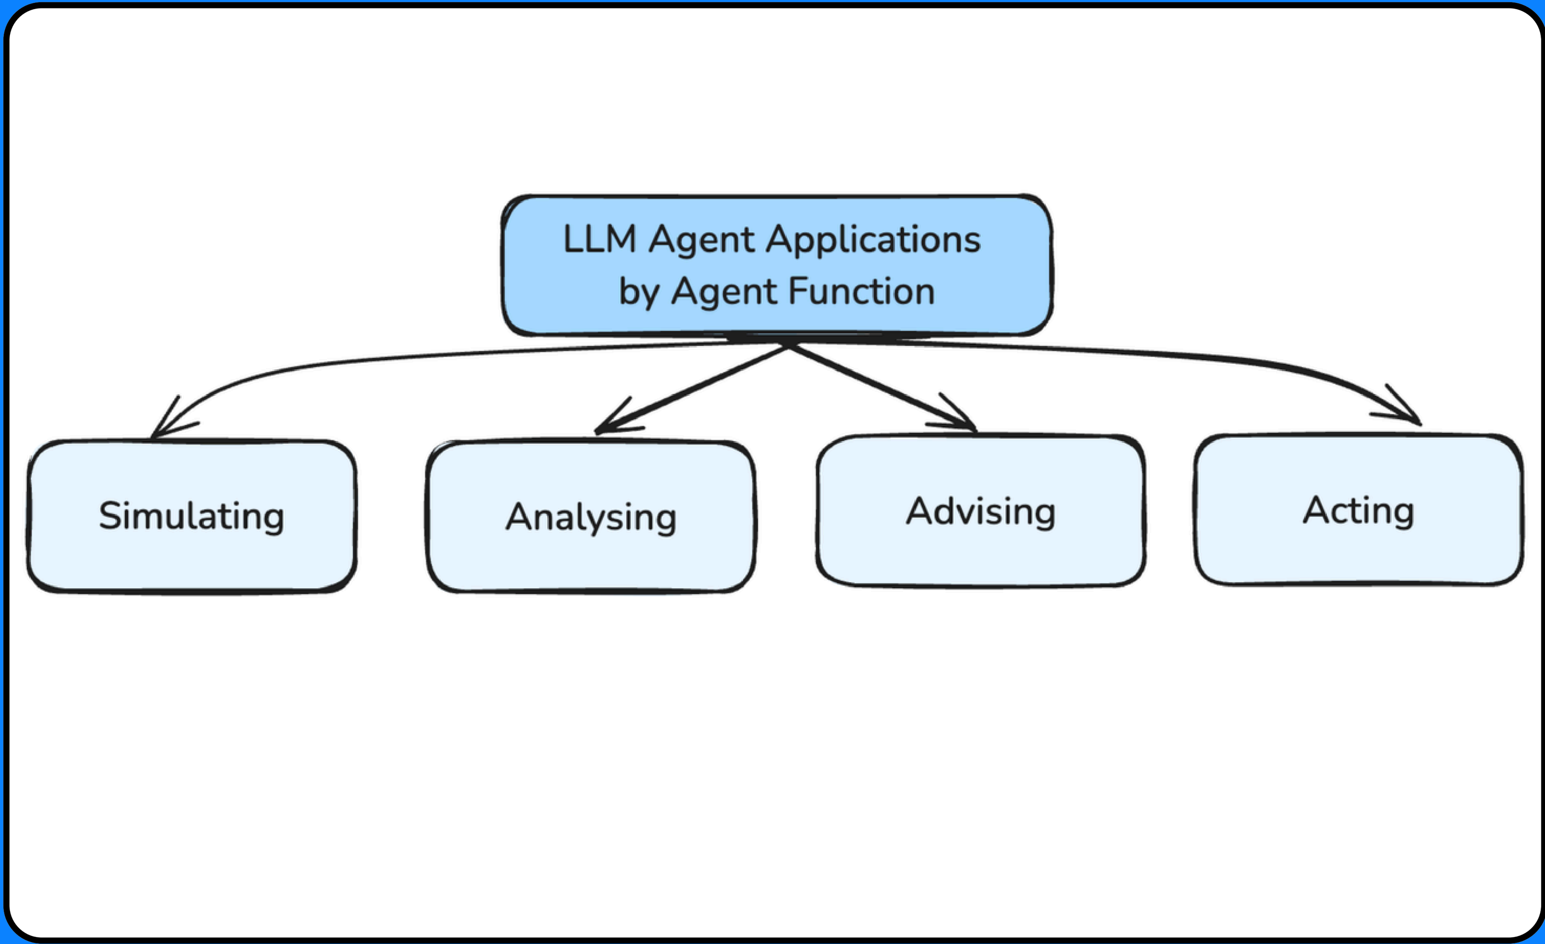

Survey charts LLM agents in finance across four roles: simulation, analysis, advising, acting. Evidence spans EconAgent, TradingGPT, StockAgent, trading, sentiment analysis, compliance, and DeFi. Surprises include tacit price collusion, agents out-rationalizing humans, and synthetic data rivaling real. Contributions: a unifying taxonomy, a risk and regulation lens citing EU AI Act, and a blueprint for hybrid oversight, bias control, evaluation, explainability, security, and integration. Gaps are scarce head to head benchmarks, deployments, and model herding risks.

AI, LLMs, or Machine Learning for Investment Management or Trading

Natural Language Processing in Financial Markets

Large Language Model Disagreement

Yang Liu|Jiatao Liu|Tianyu Wang

Hunan University Xi'an Jiaotong-Liverpool University Tsinghua University

Researchers show top AI models often disagree and that it moves markets. They build an LLM Disagreement index from six models scoring next quarter EPS direction on anonymized earnings calls from 2002 to 2024. Dispersion correlates with analyst and machine disagreement. High disagreement clusters in smaller, volatile losers. High minus low disagreement predicts lower returns of about 1.05 to 2.54 percent over 3 to 12 months, survives controls, and signals information asymmetry and low liquidity.

	LLMD	AFD	MFD	Size	Beta	BM	AG	OPE	$R_{-12,-1}$
Panel A: Descriptive statistics									
Mean	0.681	0.146	0.004	8574.548	1.127	0.545	0.164	0.098	0.160
Std	0.291	0.371	0.001	20132.510	0.449	0.405	0.378	0.110	0.418
Skew	0.063	4.838	0.981	4.202	0.315	1.430	3.035	0.390	1.233
Min	0.006	0.002	0.002	83.414	0.188	0.038	-0.302	-0.240	-0.528
5^{th}	0.220	0.007	0.003	167.712	0.438	0.094	-0.144	-0.056	-0.355
50^{th}	0.657	0.041	0.004	1783.560	1.096	0.449	0.069	0.091	0.095
95^{th}	1.177	0.563	0.005	41132.587	1.921	1.304	0.807	0.278	0.916
Max	1.433	2.602	0.007	126545.292	2.303	2.100	2.239	0.480	1.820
Panel B: Correlation matrix									
LLMD	1.000	0.136	0.070	-0.099	0.028	0.077	-0.032	-0.010	-0.085
AFD	0.136	1.000	0.259	-0.303	0.250	0.217	-0.100	-0.115	-0.138
MFD	0.070	0.259	1.000	-0.189	0.163	0.044	0.025	-0.125	-0.081
Size	-0.099	-0.303	-0.189	1.000	-0.007	-0.299	0.058	0.027	0.127
Beta	0.028	0.250	0.163	-0.007	1.000	-0.007	0.028	-0.093	-0.064
BM	0.077	0.217	0.044	-0.299	-0.007	1.000	-0.198	0.458	-0.300
AG	-0.032	-0.100	0.025	0.058	0.028	-0.198	1.000	-0.089	0.044
OPE	-0.010	-0.115	-0.125	0.027	-0.093	0.458	-0.089	1.000	-0.171
$R_{12,1}$	-0.085	-0.138	-0.081	0.127	-0.064	-0.300	0.044	-0.171	1.000

AI, LLMs, or Machine Learning for Investment Management or Trading

LLMs for Stock Analysis and Prediction

Technical Report: Full-Stack Fine-Tuning for the Q Programming Language

Brendan R. Hogan|Will Brown|Adel Boyarsky|Anderson Schneider|Yuriy Nevmyvaka

Morgan Stanley Prime Intellect

	3 month	6 month	9 month	12 month
Panel A: Large language model disagreement (LLMD)				
Low	2.90	5.98	9.59	12.94
2	3.27	6.16	9.67	12.76
3	2.62	5.22	8.74	12.11
4	2.09	4.69	8.01	11.35
High	1.82	4.34	7.18	10.36
High-Low	-1.05***	-1.56***	-2.34***	-2.54***
	(-4.54)	(-4.11)	(-4.91)	(-4.09)
Panel B: Analyst forecast disagreement (AFD)				
Low	2.73	5.67	8.78	11.77
2	2.65	5.51	8.72	11.66
3	2.91	5.59	8.66	11.43
4	2.81	5.69	8.90	11.91
High	2.26	5.16	8.50	11.96
High-Low	-0.47	-0.52	-0.28	0.19
	(-1.00)	(-0.69)	(-0.29)	(0.17)
Panel C: Machine learning forecast disagreement (MFD)				
Low	3.08	6.25	9.63	13.08
2	3.16	5.94	9.45	12.65
3	2.90	5.78	9.02	12.05
4	2.88	5.68	8.89	11.71
High	2.47	5.38	8.47	11.40
High-Low	-0.60	-0.88	-1.15	-1.68
	(-1.44)	(-1.37)	(-1.39)	(-1.76)

An open source recipe adapts LLMs to niche languages, proven on Q. A Qwen 2.5 32B reasoning model hits 59 percent pass@1 on a Q benchmark, beating Claude Opus 4 by 29.5 points. All sizes surpass GPT 4.1. Releases include a benchmark, a split solution and test pipeline, a corpus, and staged training with GRPO. Gains steady. Frontier models lag Q. SFT turns pythonic. Limits include dataset bias, noisy judging, and weak RL on 1.5B.

AI, LLMs, or Machine Learning for Investment Management or Trading

Quantum Computing in Financial Modeling

ZANISTA

Large Language Models as Simulated Human Investors: A Paradox of Model Sophistication

Ji Wang|Yongxin Xu|Leping Zhang

Tsinghua University Monash University Fudan University

Researchers find a paradox in asset-bubble labs. As LLMs get smarter, they act less human. In 10-trader, 10-round markets, DeepSeek-V3 formed human-sized bubbles (RAD 0.437 vs human 0.347) by misreading rules, while DeepSeek-R1 priced near fundamentals (RAD 0.023). OpenAI models echoed, with GPT-5 and o3 rational, GPT-4o mixed, 4o-mini off. A Goldilocks zone emerges. Trait-based role-play restores bubbles and isolates optimism as the key driver. Methods are transparent, but prompting and external validity limit claims.

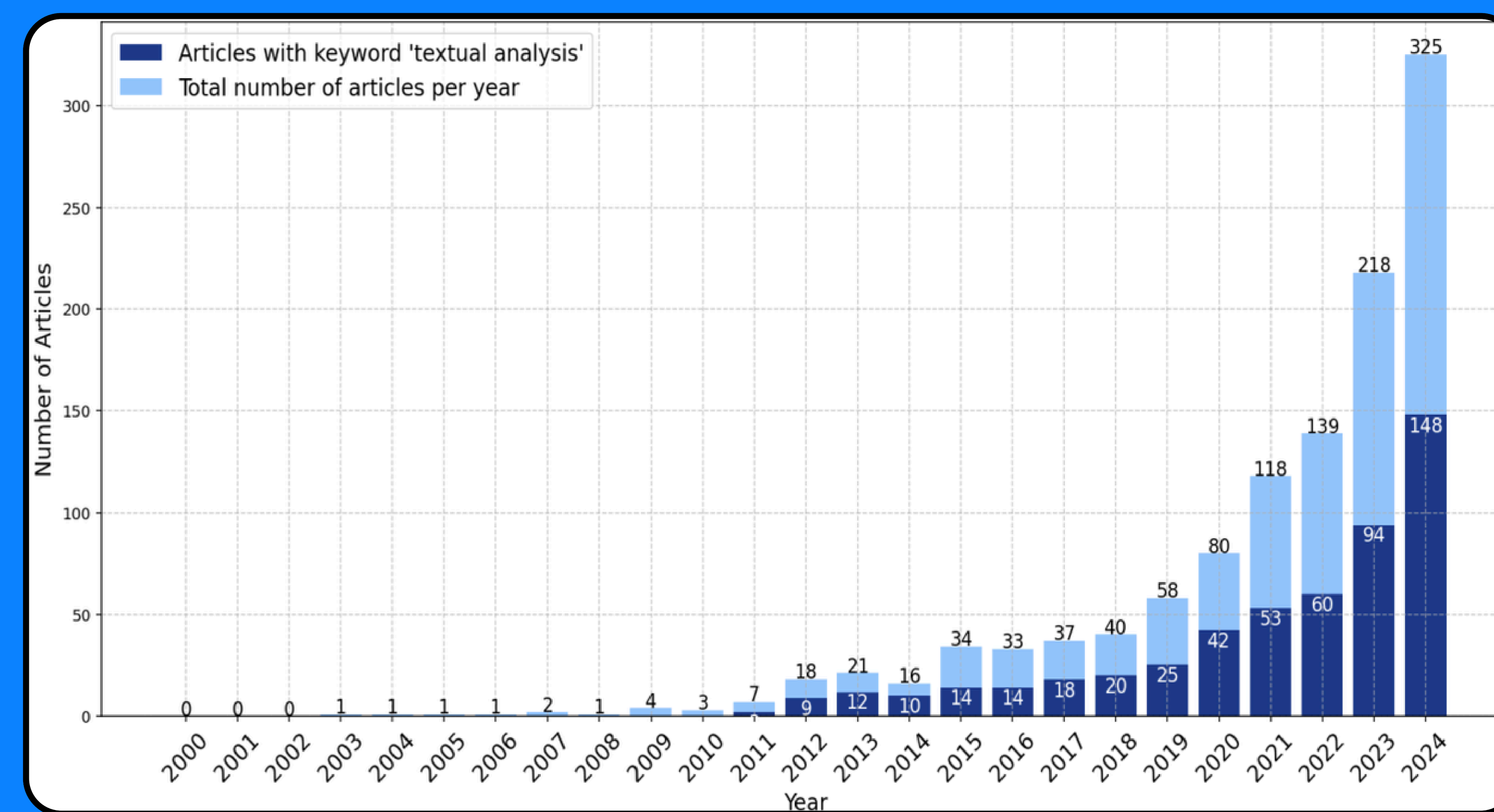

AI, LLMs, or Machine Learning for Investment Management or Trading

LLMs for Stock Analysis and Prediction

Global News Networks and Return Predictability

Gustavo Freire, Ali Moin, Alberto Quaini, Amar Soebhag

Erasmus University Rotterdam, Tinbergen Institute, Erasmus Research Institute of Management, Robeco Quantitative Investing

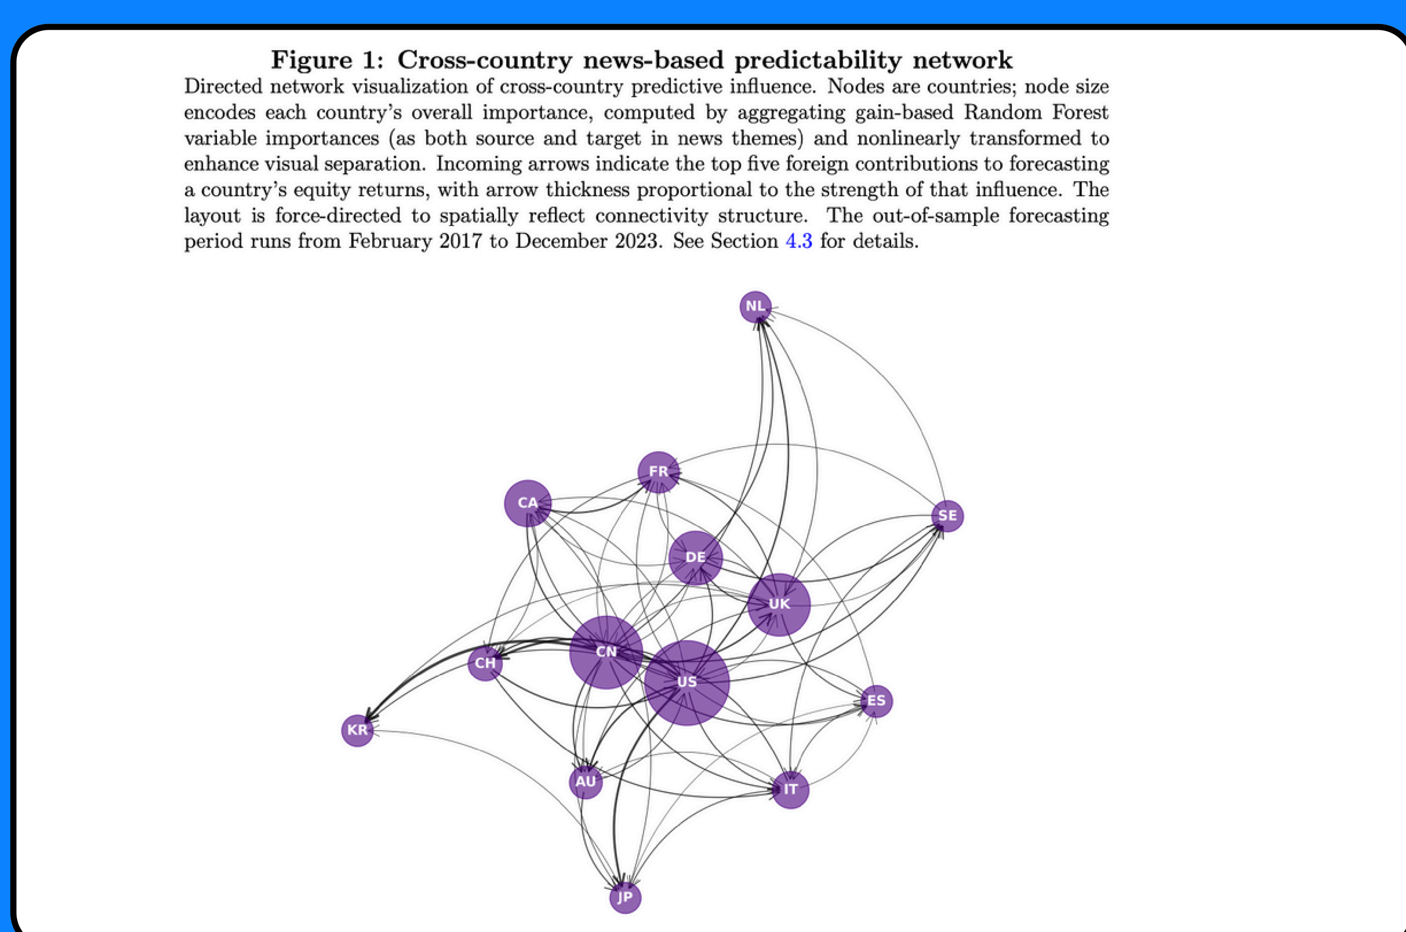

Figure 1: Cross-country news-based predictability network
Directed network visualization of cross-country predictive influence. Nodes are countries; node size encodes each country's overall importance, computed by aggregating gain-based Random Forest variable importances (as both source and target in news themes) and nonlinearly transformed to enhance visual separation. Incoming arrows indicate the top five foreign contributions to forecasting a country's equity returns, with arrow thickness proportional to the strength of that influence. The layout is force-directed to spatially reflect connectivity structure. The out-of-sample forecasting period runs from February 2017 to December 2023. See Section 4.3 for details.

The paper investigates how global news sentiment networks affect international equity return predictability using 520 million GDELT articles across 14 developed markets. By modeling local and cross-country sentiment linkages across 260 themes, the authors show that both local and global sentiment improve market-timing strategies. Random Forest forecasts yield significantly higher Sharpe ratios and alphas, highlighting market integration via shared sentiment

AI, LLMs, or Machine Learning for Investment Management or Trading

Language Model Guided Reinforcement Learning in Quantitative Trading

Adam Darmanin, Vince Vella

University of Malta

The paper proposes a hybrid LLM+RL system for quantitative trading, combining LLMs' strategic reasoning with RL's execution strengths. LLMs generate high-level trading strategies, refined through prompt engineering, memory, and news integration, while RL agents handle short-term execution. Experiments show improved Sharpe Ratios, reduced uncertainty, and smoother drawdowns compared to RL-only baselines, demonstrating modular, interpretable, and more trustworthy algorithmic trading architectures

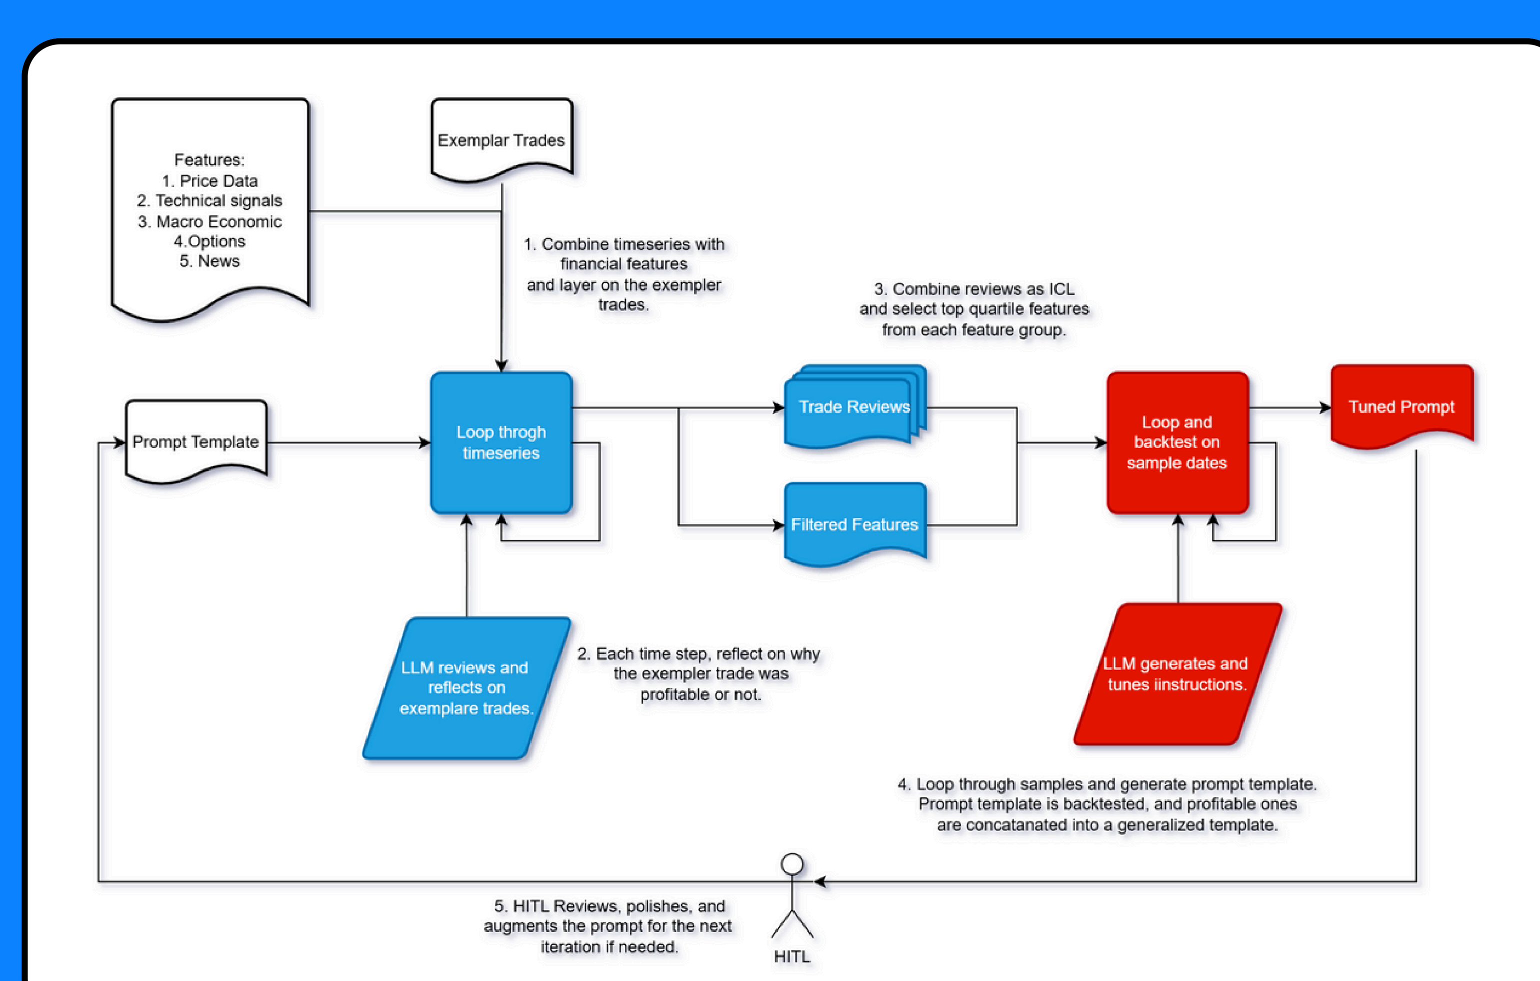

AI, LLMs, or Machine Learning for Investment Management or Trading

ZANISTA

(Deep) Learning to Trade: An Experimental Analysis of AI Trading and Market Outcomes

Andreas J. Bütler, Scott A. MacLellan
Robeco

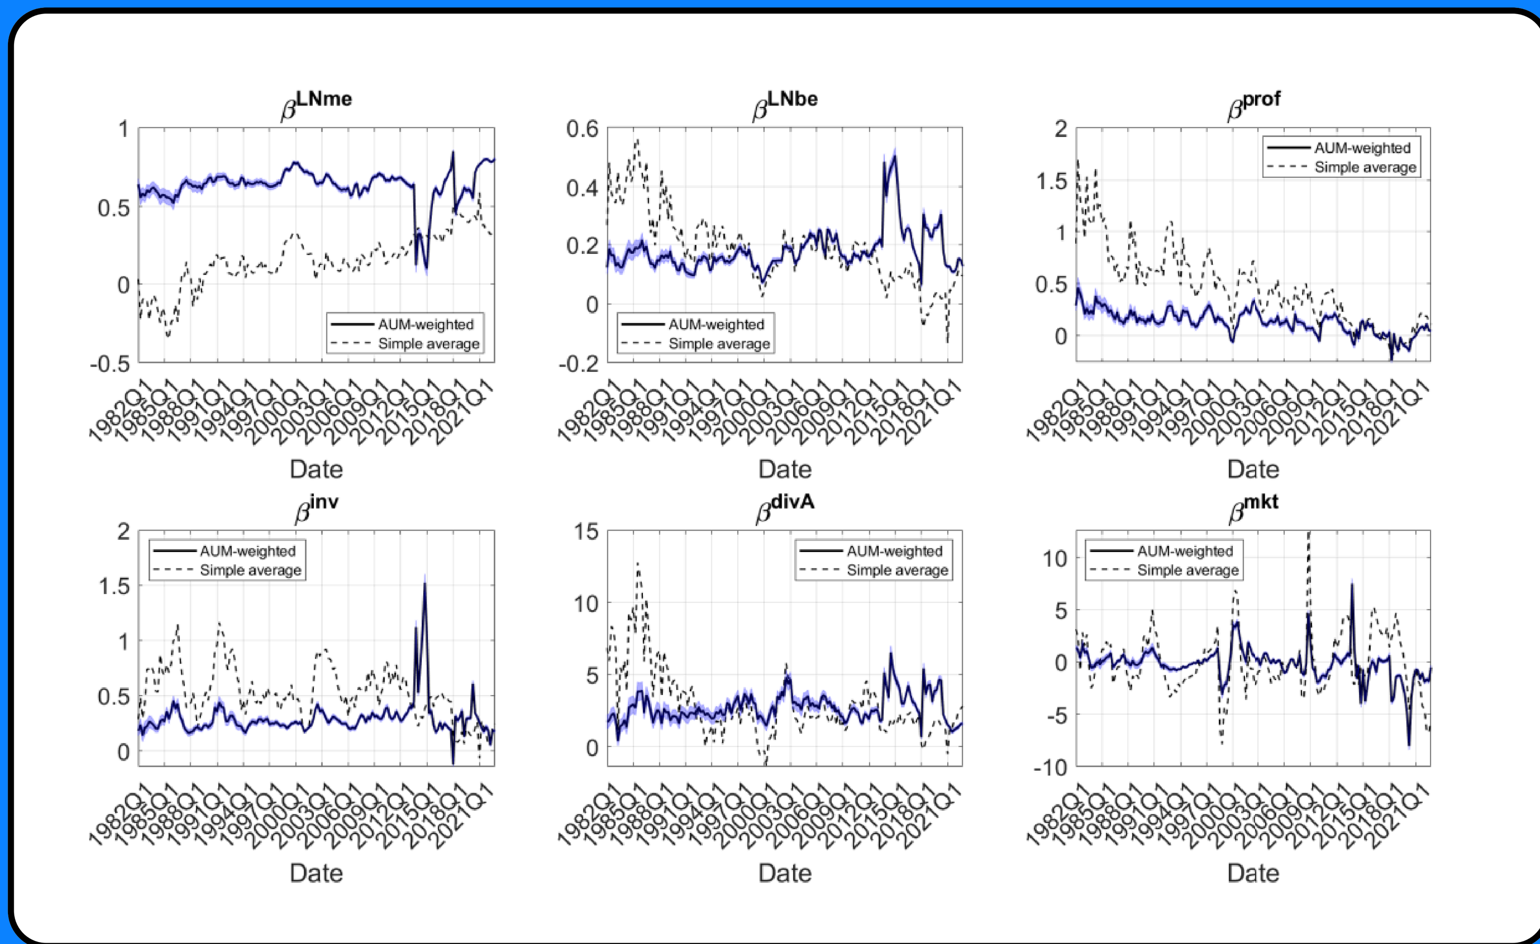

The paper explores the role of Generative AI in asset management, focusing on portfolio optimization, sentiment analysis, and forecasting. It reviews current applications, limitations, and regulatory concerns, highlighting both efficiency gains and risks like bias, hallucination, and compliance challenges. The study emphasizes the importance of explainability, human oversight, and robust governance to integrate AI responsibly into institutional investment decision-making .

AI, LLMs, or Machine Learning for Investment Management or Trading

FinReflectKG: Agentic Construction and Evaluation of Financial Knowledge Graphs

Abhinav Arun, Fabrizio Dimino, Tejas Prakash Agarwal, Bhaskarjit Sarmah, Stefano Pasquali
Domyn

The paper introduces FinReflectKG, an open-source framework and dataset for constructing large-scale financial knowledge graphs (KGs) from SEC 10-K filings of all S&P 100 companies. It addresses challenges in financial text processing through intelligent parsing, table-aware chunking, schema-guided extraction, and a novel reflection-agent workflow. The framework supports three extraction modes (single-pass, multi-pass, reflection-based) and incorporates holistic evaluation combining rules, statistical checks, and LLM judgments. Results show reflection-based extraction achieves the best trade-off between precision, coverage, and compliance

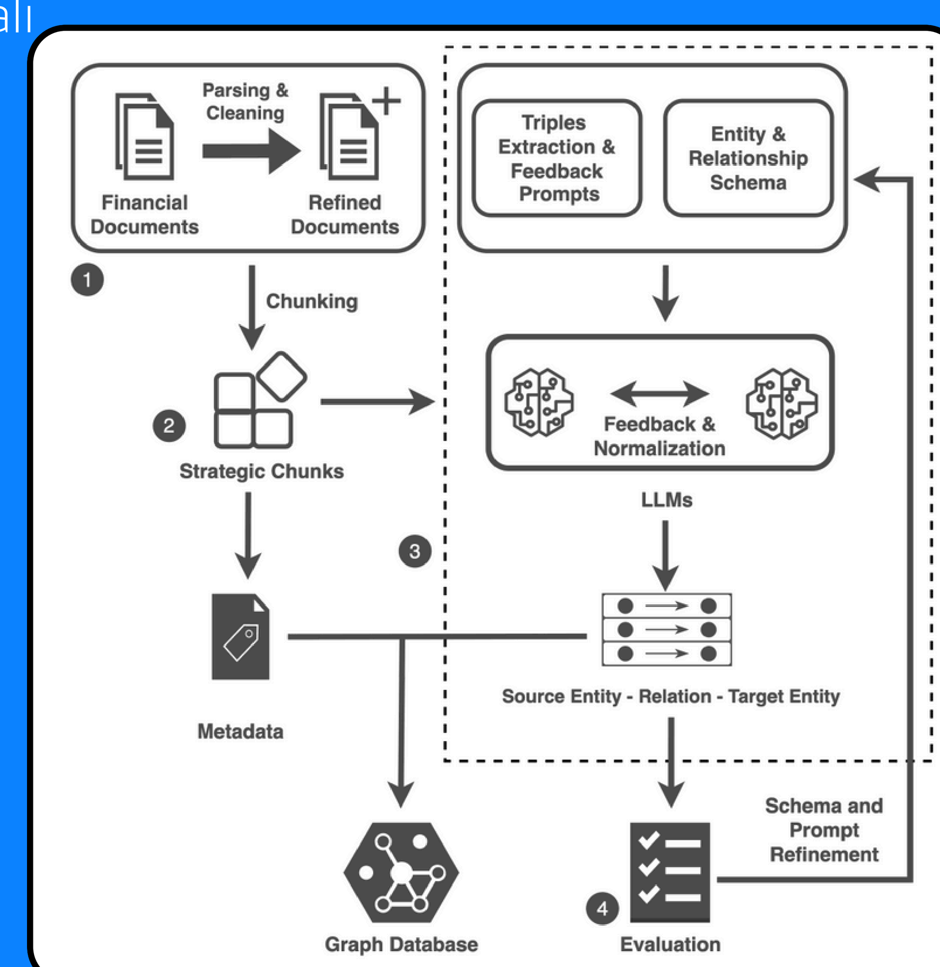

AI, LLMs, or Machine Learning for Investment Management or Trading

MountainLion: A Multi-Modal LLM-Based Agent System for Interpretable and Adaptive Financial Trading

Siyi Wu, Junqiao Wang, Zhaoyang Guan et al
Northwestern University, Xi'an University of Electronic Science and Technology, Kyoto University, Tsinghua University, University of Toronto

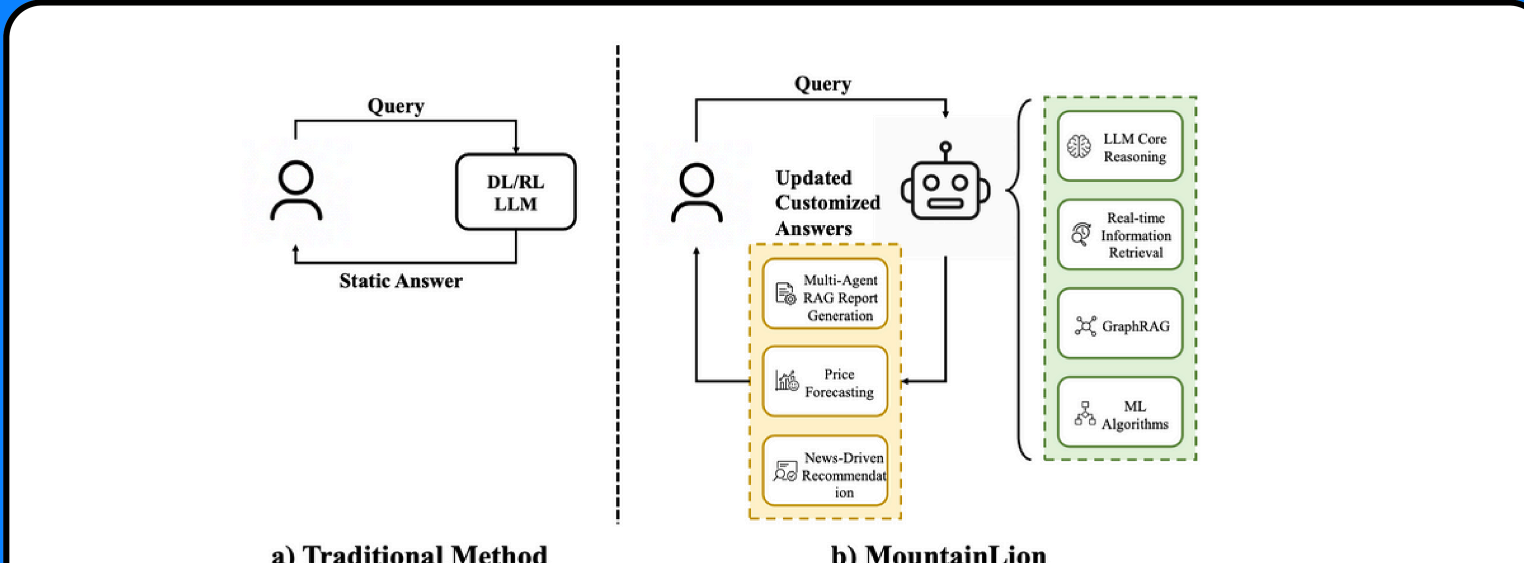

Figure 1: Comparison of (a) traditional DL/RL-based static QA pipelines and (b) our MountainLion framework. MountainLion produces updated, personalized answers through multi-agent collaboration, RAG, and reflective reasoning.

The paper introduces MountainLion, a multi-agent, multi-modal LLM-based system for cryptocurrency trading. It integrates news, candlestick charts, trading signals, and on-chain data to produce interpretable investment strategies. Specialized agents collaborate through a reflection module, while RAG ensures real-time adaptability and reduced hallucinations. Experiments demonstrate improved forecasting accuracy, robustness, and investor confidence compared to traditional DL/RL methods

AI, LLMs, or Machine Learning for Investment Management or Trading

ZANISTA

AlphaAgents: Large Language Model based Multi-Agents for Equity Portfolio Constructions

Tianjiao Zhao, Jingrao Lyu, Stokes Jones, Harrison Garber, Stefano Pasquali, Dhagash Mehta

BlackRock, Inc.

AlphaAgents uses a role based LLM team for equity research. Fundamental, Sentiment, and Valuation agents analyze filings, news, and prices, discuss and debate until they reach consensus. The setup reduces human bias and hallucination, produces buy sell calls, and encodes risk tolerance. In backtests, the multi agent portfolio beat single agents under risk neutral settings. Under risk averse settings it lagged benchmark but limited drawdowns. Signals can feed portfolio optimizers.

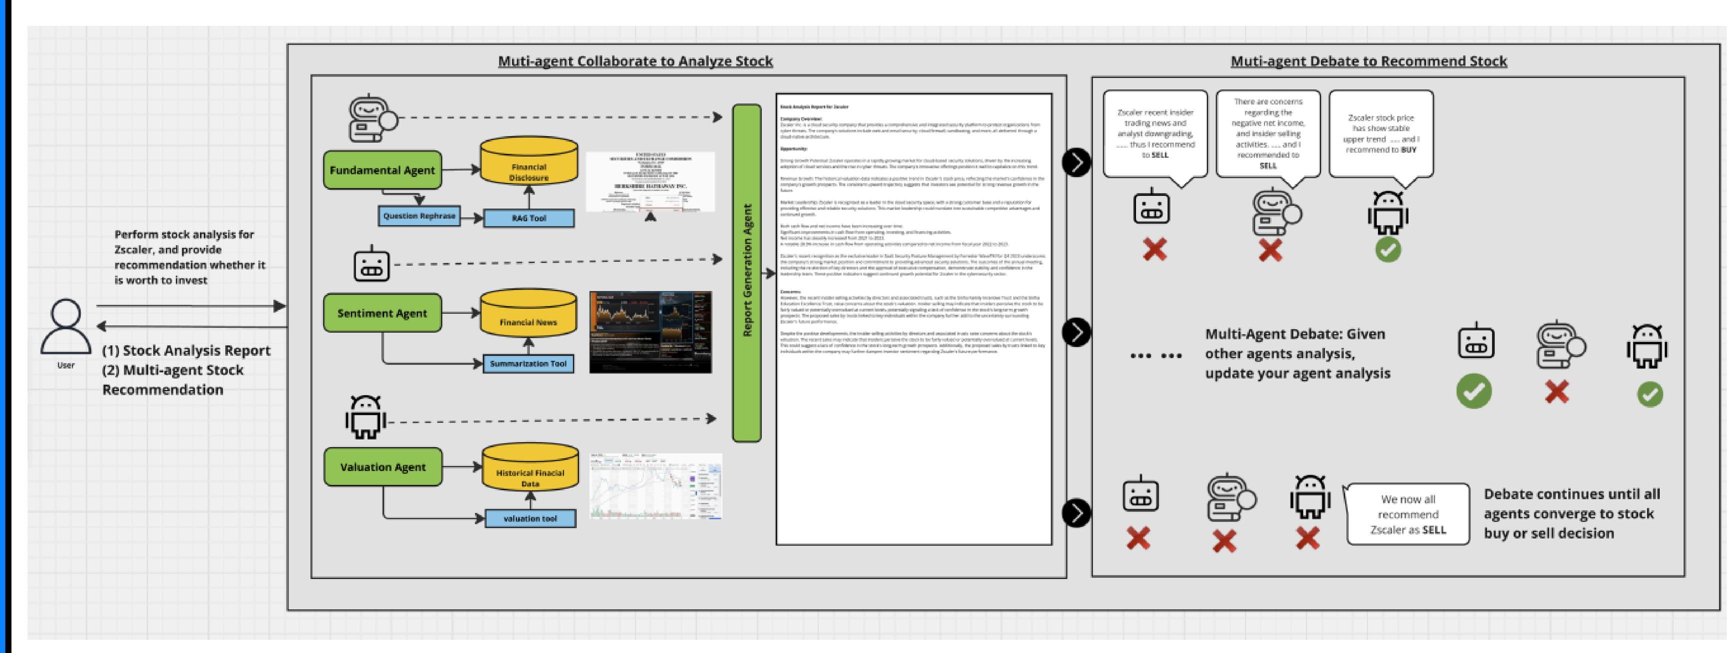

AI, LLMs, or Machine Learning for Investment Management or Trading

TradingGroup: A Multi-Agent Trading System with Self-Reflection and Data-Synthesis

Feng Tian, Flora D. Salim, Hao Xue

University of New South Wales

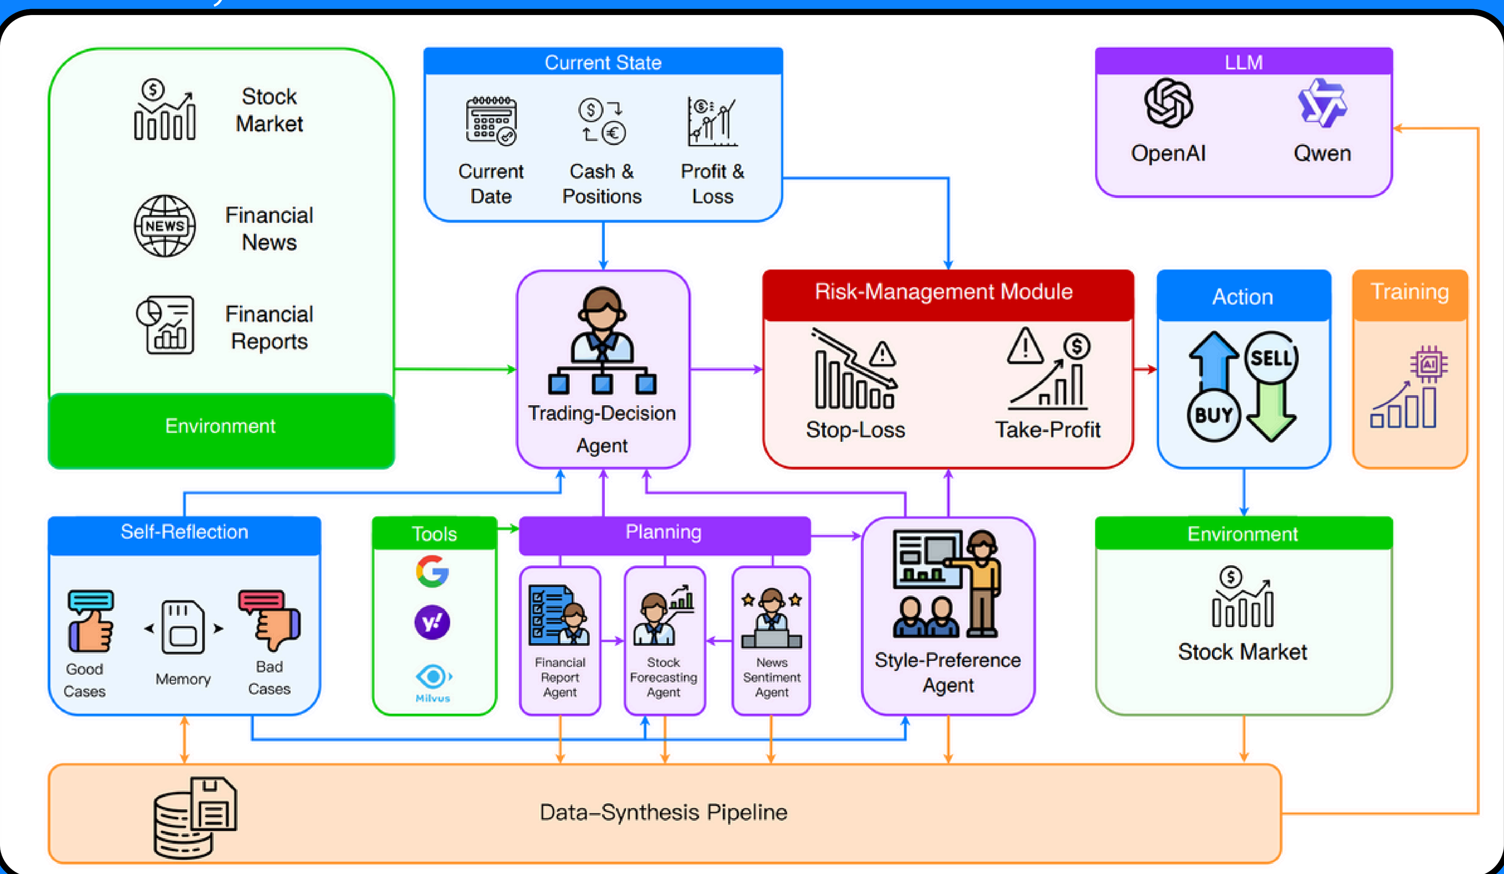

TradingGroup is a multi agent LLM trading system. News Sentiment, Financial Report, Stock Forecasting, Style Preference, and Trading Decision agents analyze signals, debate, and reach actions, supported by a dynamic risk management module. Built in self reflection mines past wins and mistakes to refine prompts and style, while a data synthesis pipeline auto labels trajectories for post training. In backtests on five stocks, TradingGroup beat rule based, ML, RL, and prior LLM agents, improving returns and drawdowns without sacrificing interpretability.

AI, LLMs, or Machine Learning for Investment Management or Trading

THEME: Enhancing Thematic Investing with Semantic Stock Representations and Temporal Dynamics

Hoyoung Lee, Wonbin Ahn, Suhwan Park, Jaehoon Lee, Minjae Kim, Sungdong Yoo, Taeyoon Lim, Woohyung Lim, Yongjae Lee

Ulsan National Institute of Science and Technology, LG AI Research

The paper introduces THEME, a hierarchical contrastive learning framework for thematic investing. It fine-tunes stock embeddings by aligning themes with constituent stocks and refining them via recent returns, enabling dynamic theme-aware retrieval. Using the Thematic Representation Set (TRS), THEME outperforms LLM baselines in retrieval precision and portfolio construction. Results show higher Sharpe ratios, improved cumulative returns, and reduced drawdowns compared to ETFs and general embeddings

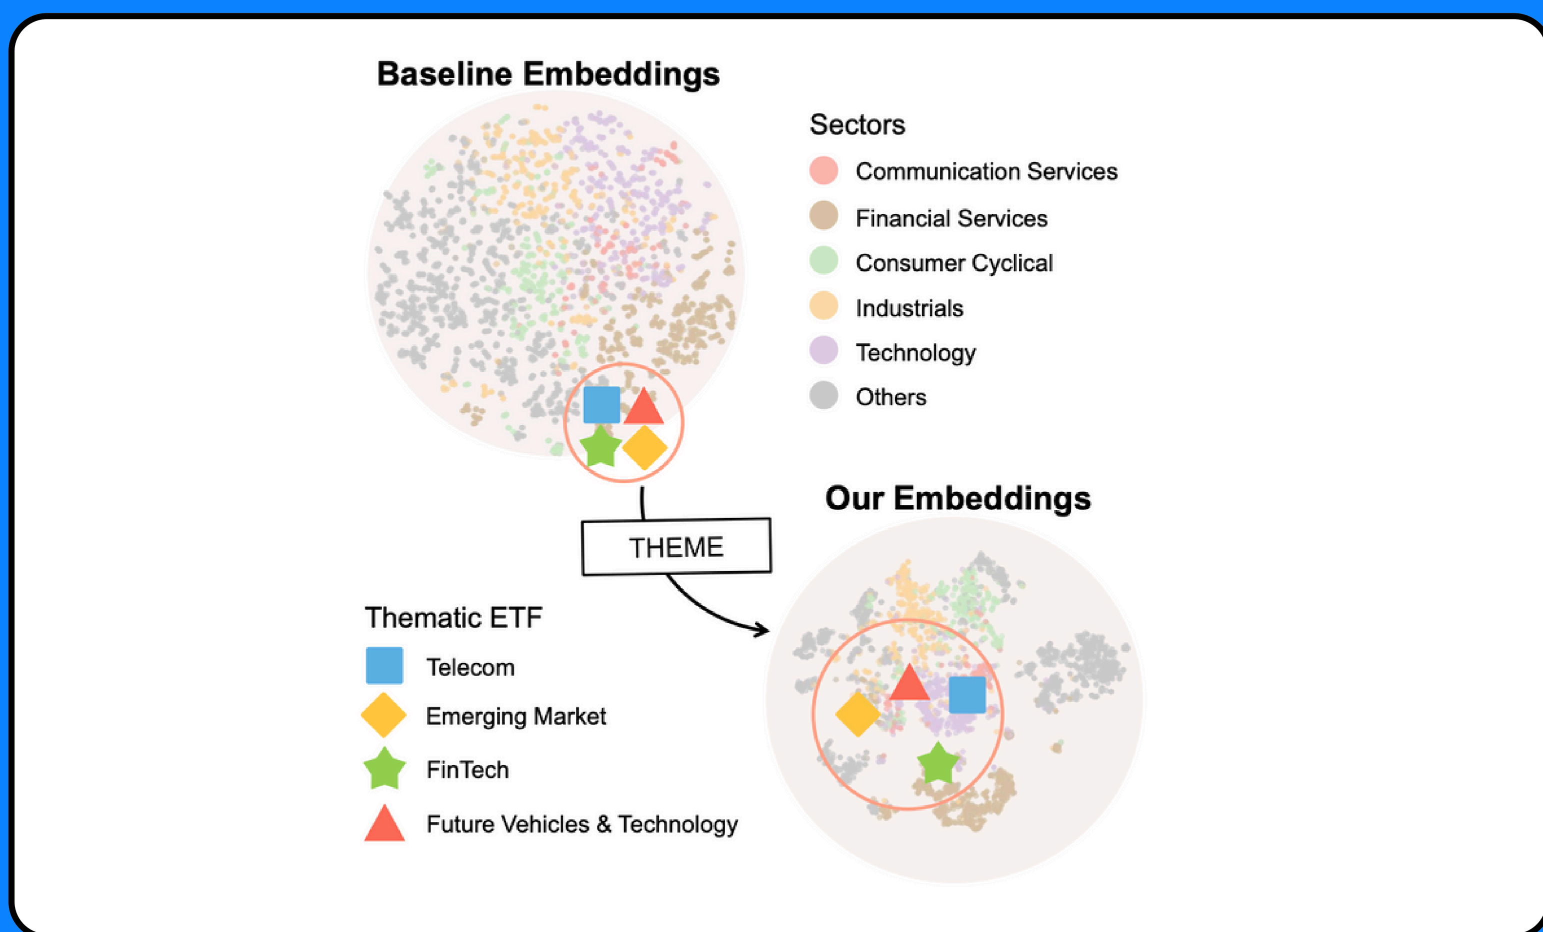

AI, LLMs, or Machine Learning for Investment Management or Trading

ZANISTA

Can large language models autonomously generate unique and profound insights in fundamental analysis?

Tao Xu, Zhe Piao, Tadashi Mukai, Yuri Murayama, Kiyoshi Izumi

The University of Tokyo, Nomura Holdings, Inc.

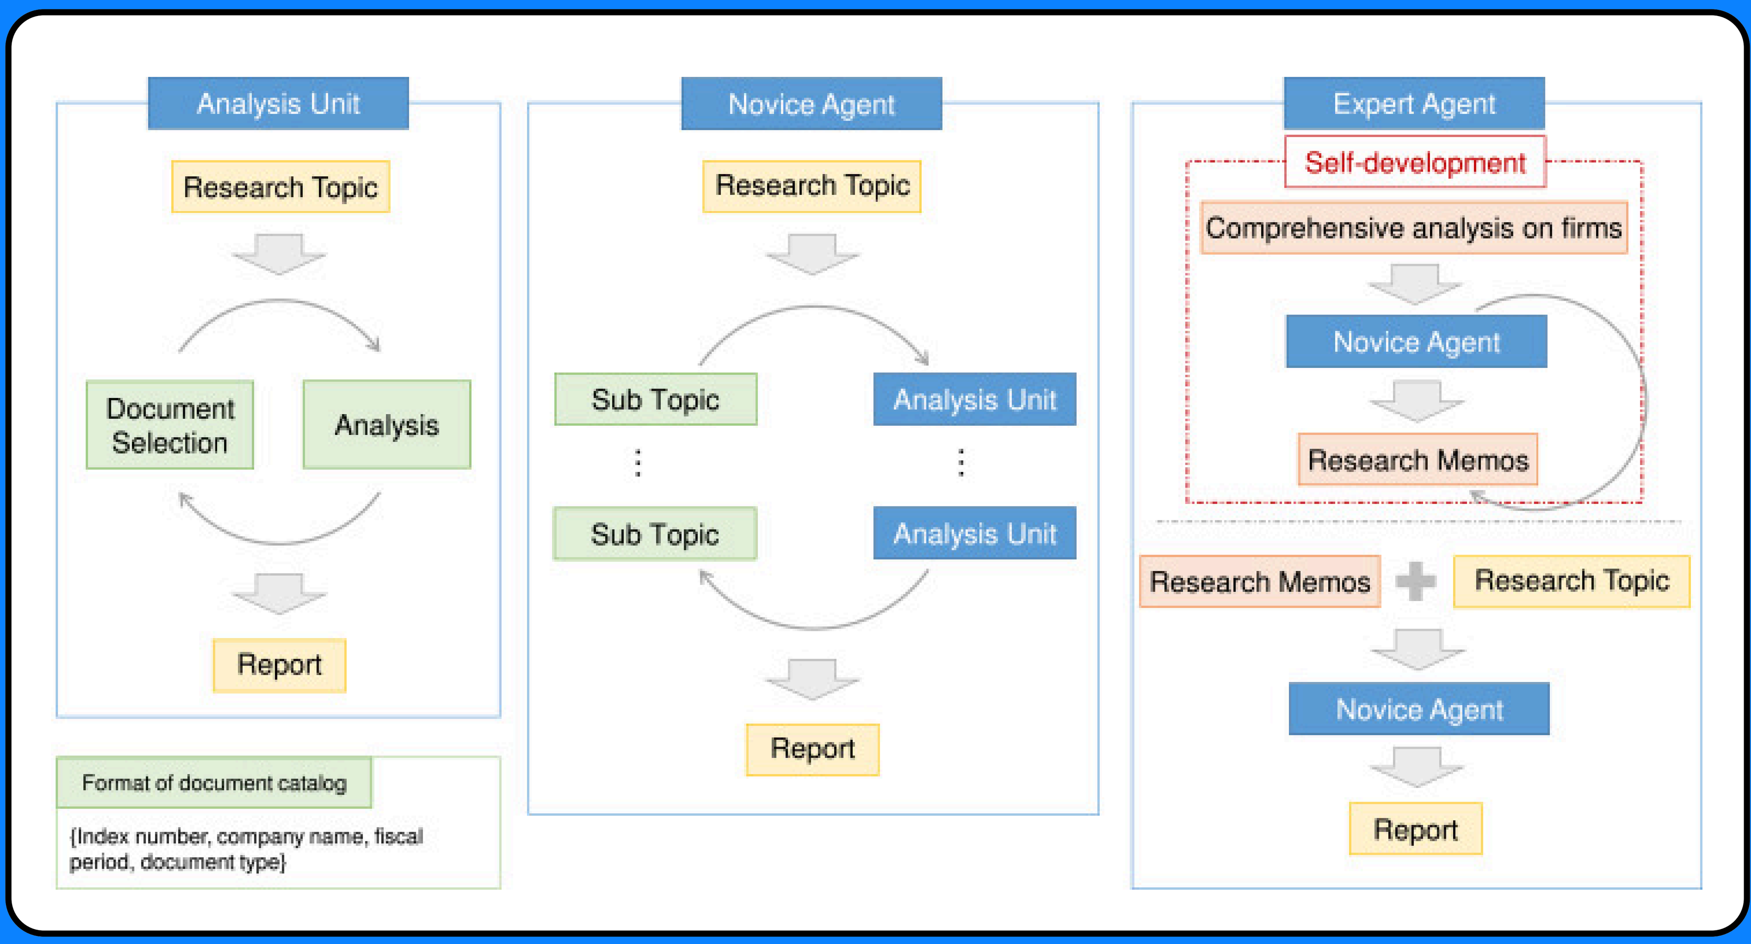

The paper proposes AutoFAS (Autonomous Fundamental Analysis System), enabling LLM agents to autonomously perform equity fundamental analysis. By mimicking human analysts' accumulation of prior research, AutoFAS develops "Expert Agents" that generate deeper, long-term, and more insightful analyses compared to "Novice Agents." Experiments across five Japanese sectors show Expert Agents consistently provide unique perspectives and profound insights, narrowing the gap between human and AI-driven fundamental analysis

AI, LLMs, or Machine Learning for Investment Management or Trading

Fin-PRM: A Domain-Specialized Process Reward Model for Financial Reasoning in Large Language Models

Yuanchen Zhou, Shuo Jiang, Jie Zhu, Junhui Li, Lifan Guo, Feng Chen, Chi Zhang

Qwen DianJin Team (Alibaba Cloud Computing), Osaka University, Soochow University

Fin-PRM is a finance specific process reward model that scores reasoning steps and full trajectories. It adds knowledge verification and coverage to reduce hallucinations. Trained on a curated 3k dataset, it guides data selection, reinforcement learning, and Best of N inference. Across CFLUE and FinQA, it beats general PRMs clearly. Reported gains: plus 12.9 percent for supervised tuning, plus 5.2 percent for RL, and plus 5.1 percent at test time.

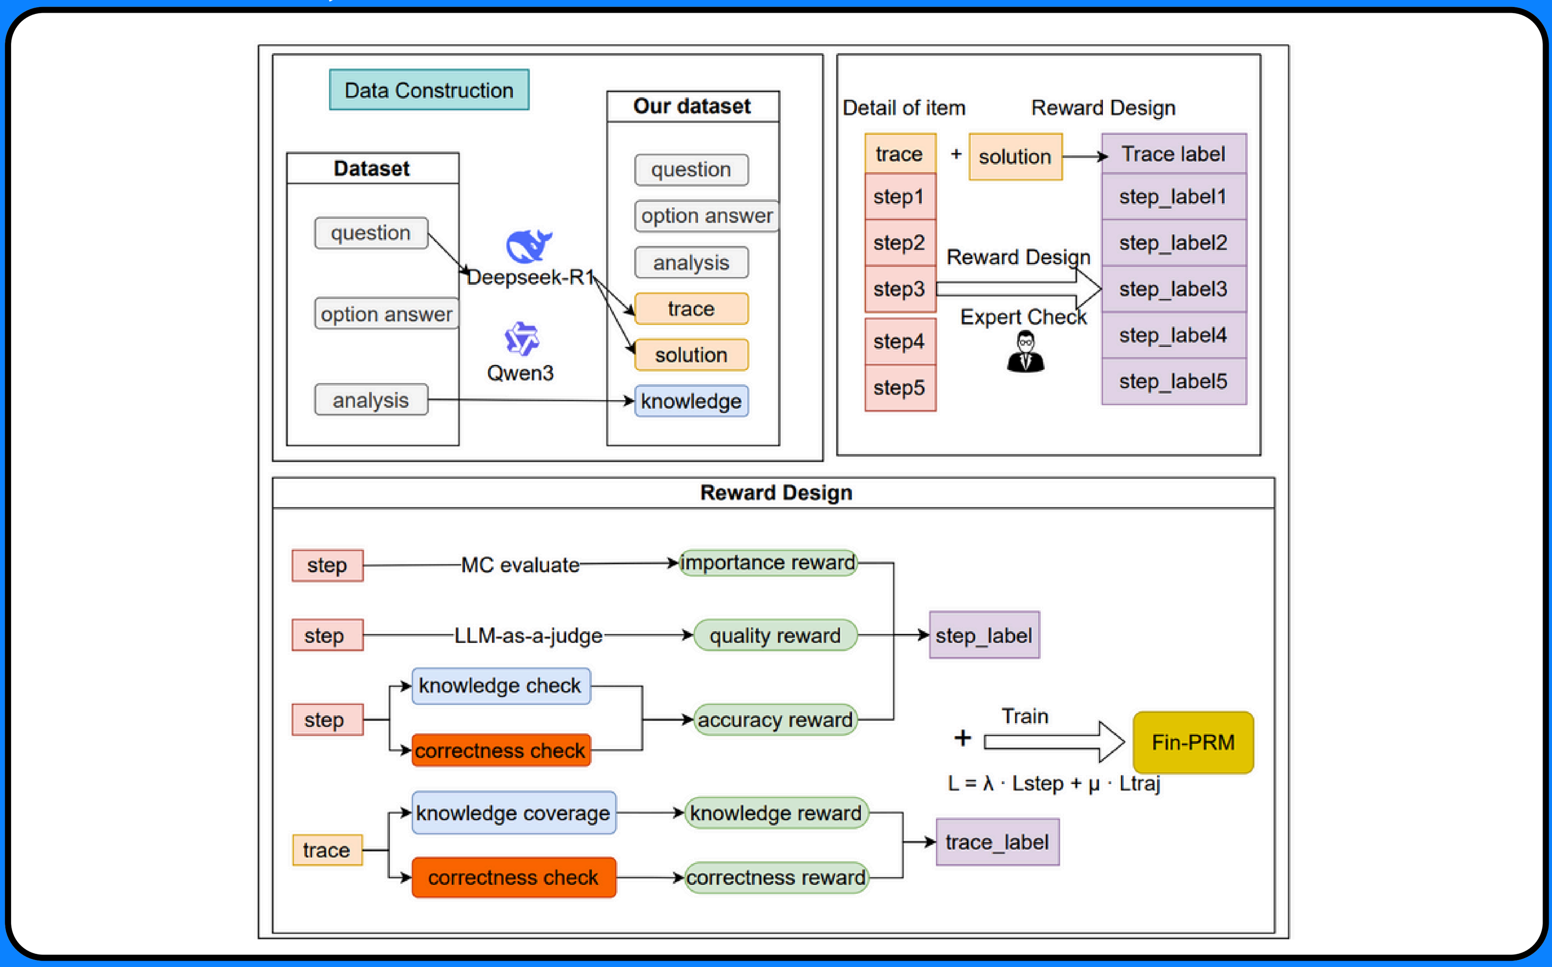

AI, LLMs, or Machine Learning for Investment Management or Trading

The Market's Mirror: Revealing Investor Disagreement with LLMs

Vineet Bhagwat|J. Anthony Cookson|Chukwuma Dim|Marina Niessner

George Washington University University of Colorado Boulder Indiana University

- **LLM Persona Disagreement**: Researchers use local Llama 3.1 8B to simulate 216 demographically varied personas, scoring 207,372 S&P 500 headlines (2010--2025) for 44.4 million belief updates; a FINRA-weighted disagreement metric validated against divides.
- **Disagreement Dynamics**: Disagreement spikes around elections, rises income near TCJA and Brexit, jumps racially after George Floyd; lowest on hard fundamentals (earnings, credit), highest on social issues; extreme bad news compresses more.
- **Disagreement Alpha**: A 1 SD rise predicts 1.2% higher volume and ~0.1% next-day gains after controls and StockTwits; long--short ~20 bps monthly with Fama--French+momentum alphas; weaknesses: LLM biases, headline-level prompts, limited ground-truth.

Researchers used Llama 3.1 8B to simulate 216 personas and score sentiment on 207,372 S&P 500 headlines from 2010 to 2025, generating 44.4 million belief updates. Their FINRA weighted disagreement measure matches splits at elections, TCJA, Brexit, and George Floyd. Disagreement is low on fundamentals and high on social issues. A one standard deviation rise lifts same day abnormal volume 1.2 percent and predicts 0.1 percent next day. A long short earns 20 bps monthly.

AI, LLMs, or Machine Learning for Investment Management or Trading

Earnings Estimates and Analyst Ratings as AI Features

ZANISTA

What the London Stock Exchange can Teach Us About Private Equity

Richard Ennis|Daniel Rasmussen

Ennis and Rasmussen use prices of European listed private equity funds to test the asset class. They find PE is riskier, more tied to public markets, and not better after risk than NAV studies claim. Volatility is far higher. Correlations near 0.94. Shares trade at 10 to 35 percent NAV discounts, averaging 20, bigger than secondaries. Beta is 1.6. CAPM alpha is minus 4.2 percent. Since 2022, LPE lagged MSCI World by 6.9 percent annually.

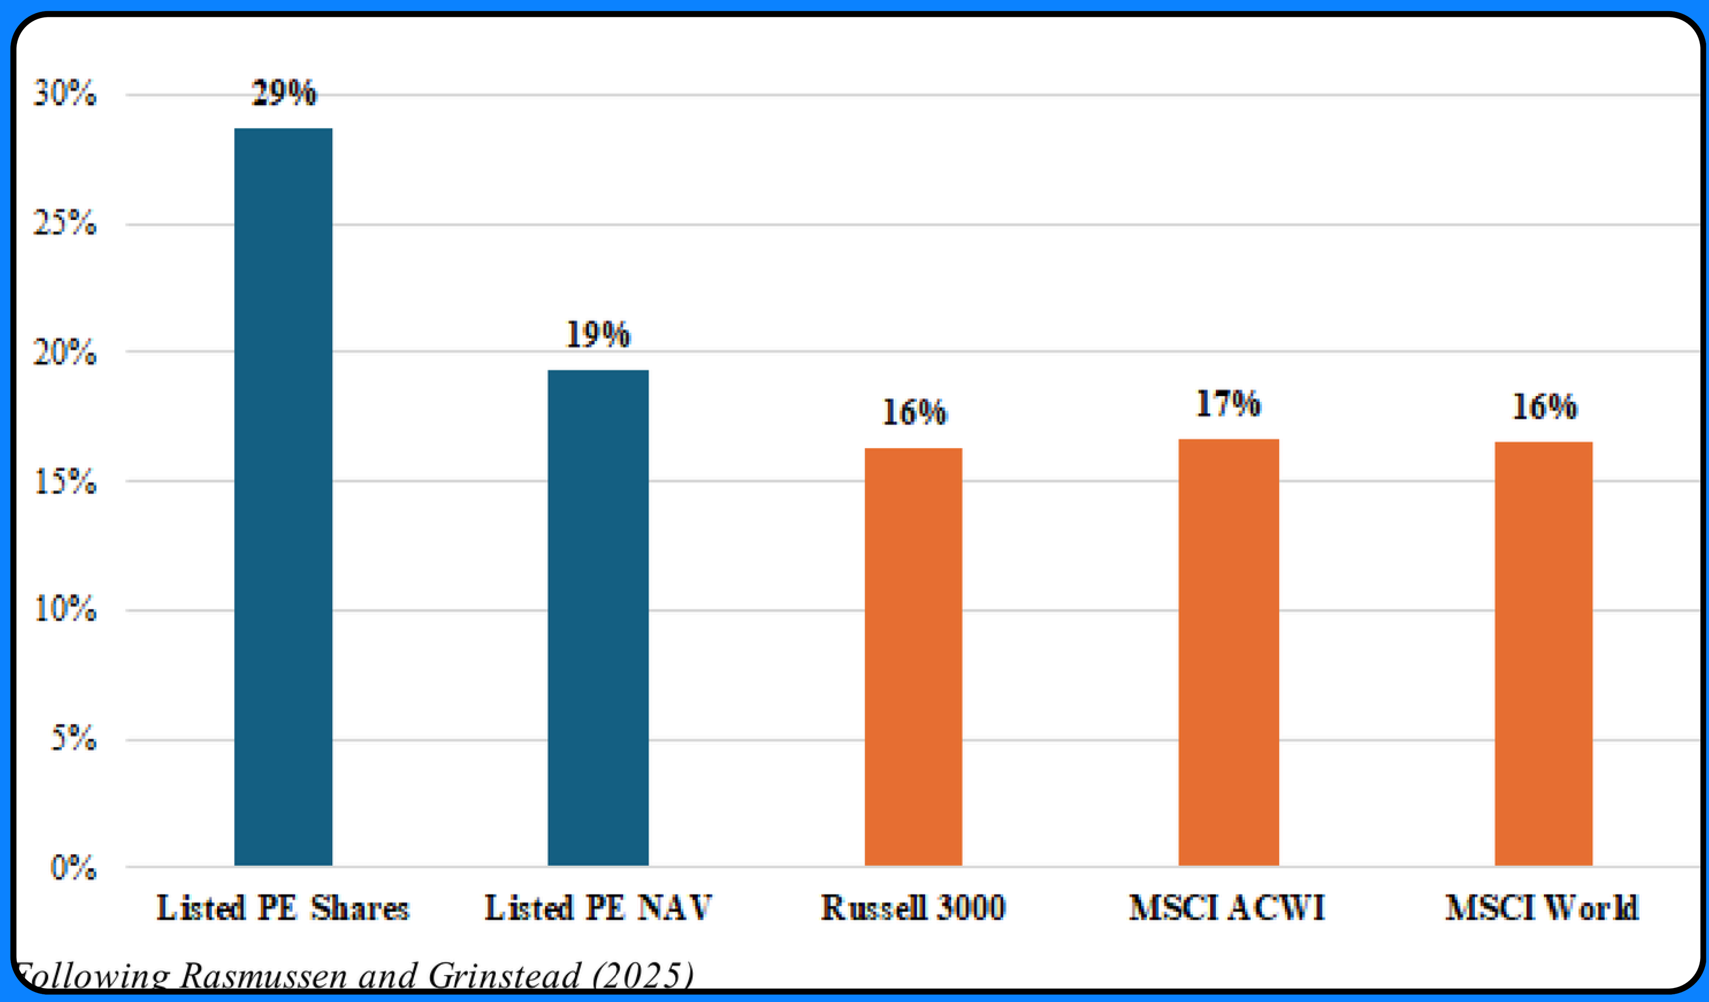

Stock Market, Equity Market

Pricing of Securities

Accounting vs technical information: what matters more for stock return predictability?

Arnoud W.A. Boot, Thomas Philippon
University of Amsterdam, New York University

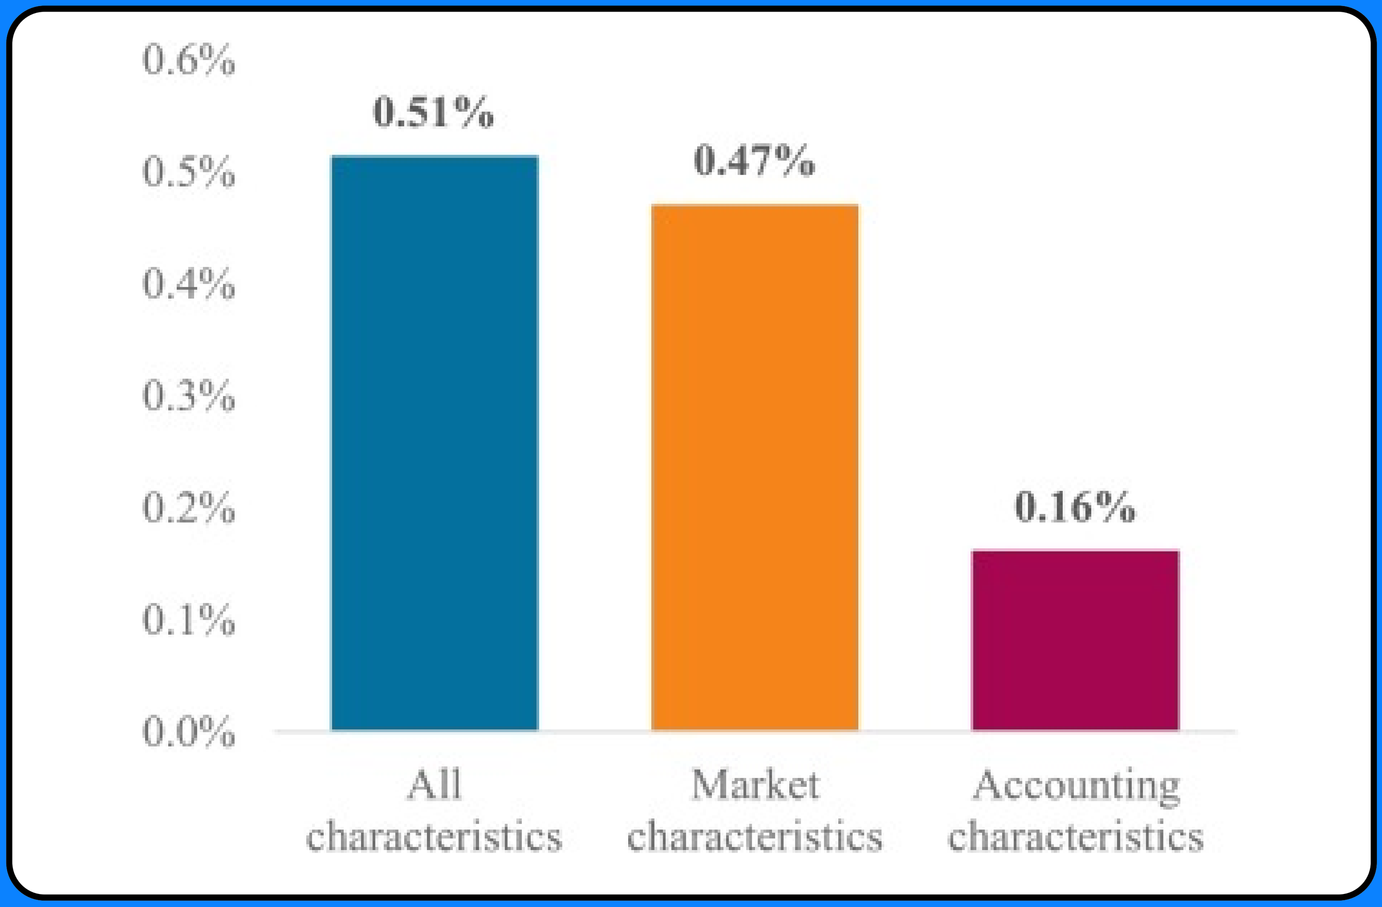

The paper explores the application of generative AI in financial intermediation, focusing on its role in enhancing productivity, client services, and market efficiency. It highlights opportunities such as personalized financial advice and risk modeling, while also examining challenges including data privacy, bias, regulatory compliance, and systemic risks. The authors argue for balanced governance frameworks to ensure responsible, transparent, and beneficial deployment of generative AI .

Stock Market, Equity Market

Interpretable Factors of Firm Characteristics

Yuxiao Jiao|Guofu Zhou|Wu Zhu|Yingzi Zhu
Tsinghua University Washington University in St. Louis Central University of Finance and Economics

C-IPCA clusters 94 stock characteristics from 1985 to 2021 and extracts one interpretable factor per cluster. Thirteen factors emerge. Top performers are Operating Illiquidity, Return Volatility, Operating Efficiency, and Size and Illiquidity. With ordered or Bayesian selection, C-IPCA often matches or beats IPCA on Sharpe ratios. FF3, FF5, and q explain only weak factors, while strong ones earn alphas. A zero correlation factor recovers the market. Limits include mis clustering, hyperparameters, strength, and US scope.

- **C-IPCA Breakthrough**: Researchers unveil C-IPCA, clustering 94 firm characteristics (1985--2021) to extract one factor per cluster, producing 13 interpretable factors. StandoutsOperating Illiquidity, Return Volatility, Operating Efficiency, Size-Illiquidityoften match or beat IPCA out-of-sample.
- **Classics Challenged**: Classic models (FF3, FF5, q) mainly explain weak C-IPCA factors; strongest show alphas. Combines priors, Chameleon clustering using value-weighted correlations, loadings restrictions, and a zero-correlation factor recovering market (corr >0.99).
- **Robustness and Limits**: Placebo tests show gains arise from economic priors and data-driven clustering, not fewer parameters, and persist under equal-weighting

Stock Market, Equity Market

Earnings Estimates and Analyst Ratings as AI Features

ZANISTA

How much money is really left on the table? Reassessing the measurement of IPO underpricing

Joseph J. Henry|Terrence M. O?Brien

Northeastern University University of Maryland

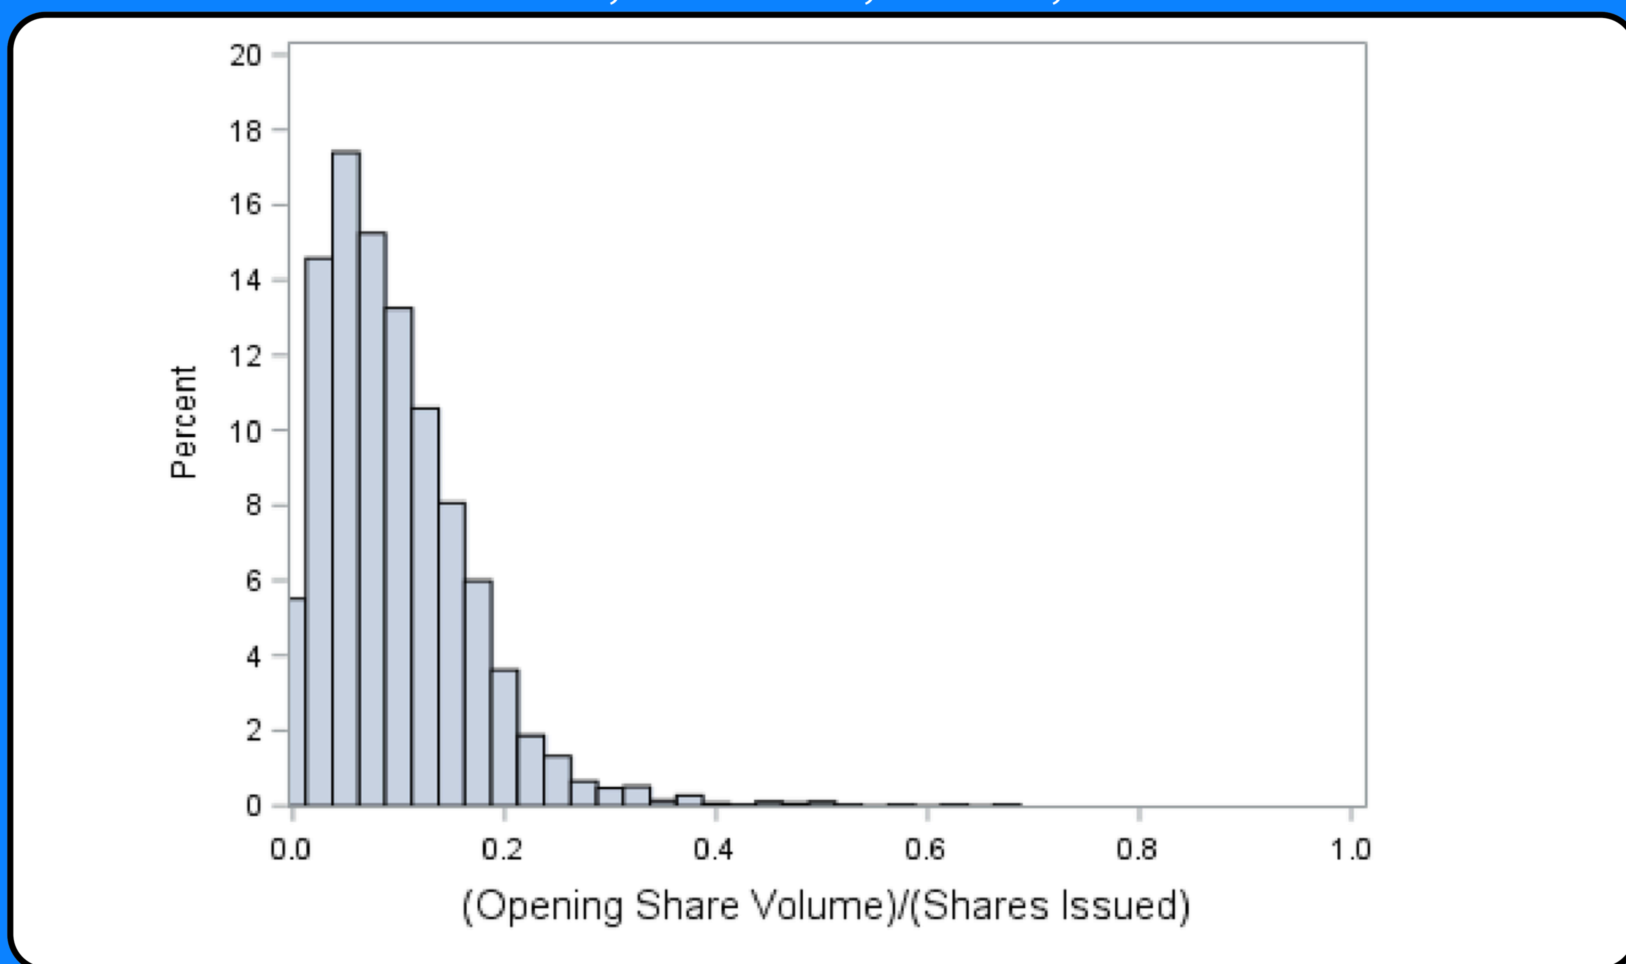

New research says the IPO pop exaggerates underpricing by comparing prices at mismatched quantities. Using a UVD measure on 2,937 IPOs from 1993 to 2023, underpricing is 40 percent lower than initial returns suggest. Only 10 percent of shares trade at open, reflecting buyers. The authors infer demand from IPO and opening data, simulate pay bid money left, urge disclosure of clearing prices, and flag limits static demand, unobserved clearing prices, and understatement convex demand.

Stock Market, Equity Market | Pricing of Securities

Virtual Trading in Multi-Settlement Electricity Markets

Agostino Capponi|Garud Iyengar|Bo Yang|Daniel Bienstock

Columbia University The Hong Kong University of Science and Technology

The paper shows that in two-settlement electricity markets, load-serving entities (LSEs) underbid in the day-ahead (DA) market to depress prices below expected real-time (RT) levels. Virtual trading reduces and, with many traders, eliminates this DA–RT price gap. However, it does not align quantities: DA demand remains below true expectations as LSEs avoid overcommitting. Renewable suppliers cannot offset this distortion. Empirical data from CAISO and NYISO confirm reduced price gaps and lowered DA demand after virtual trading.

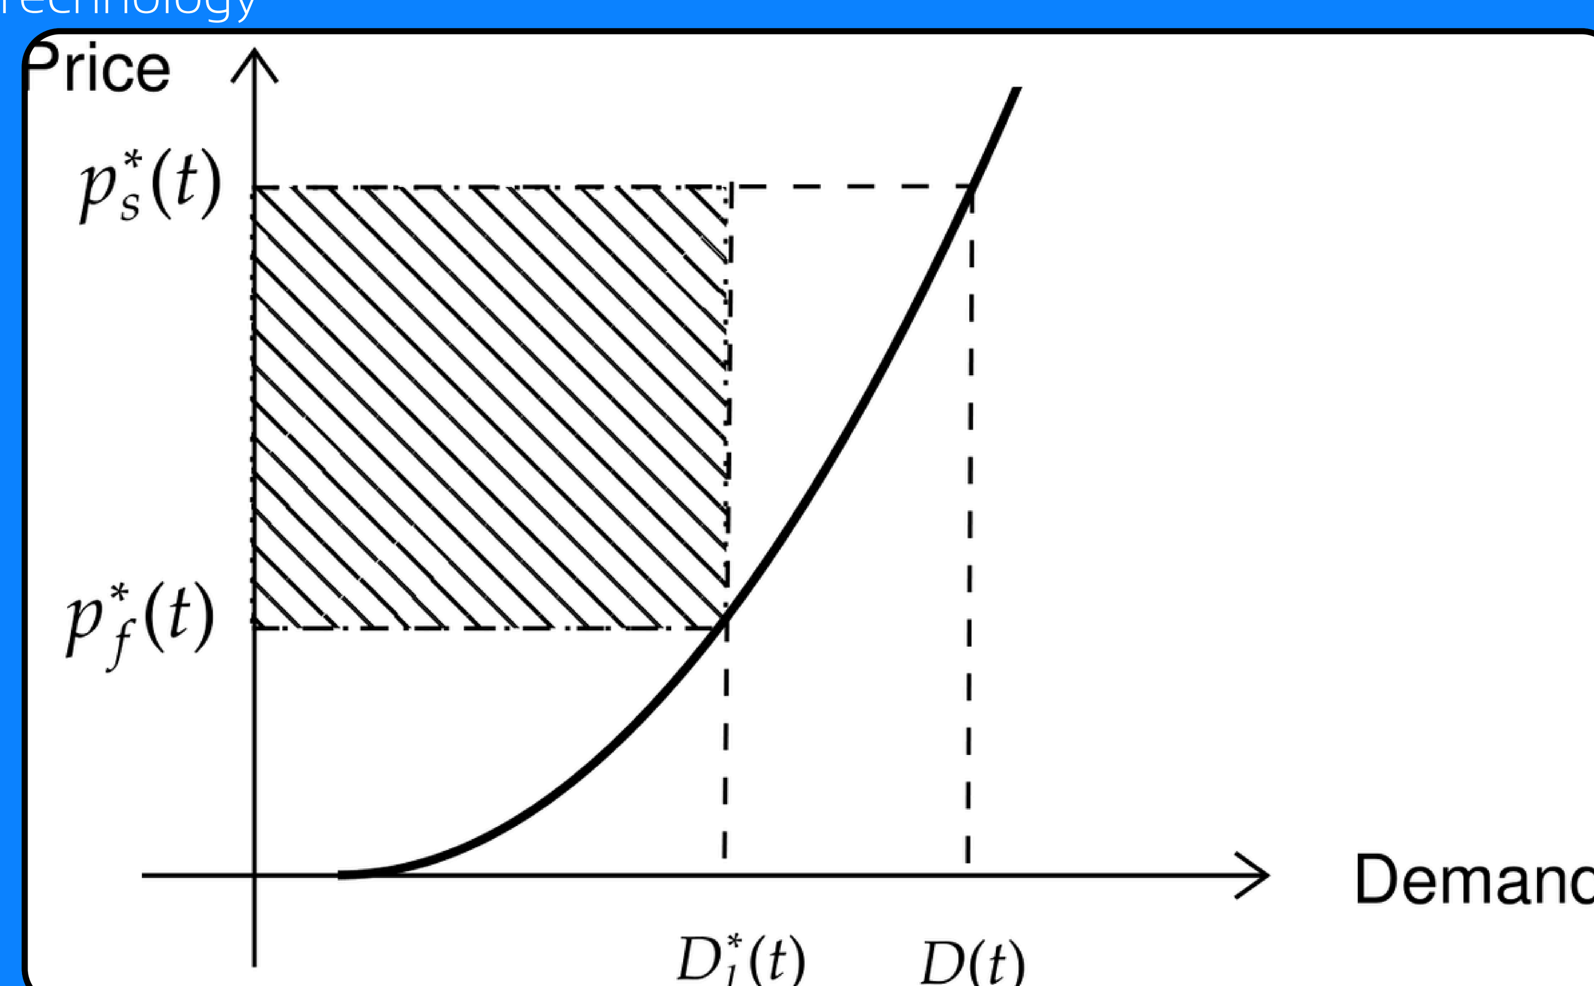

Commodity Market | Trading and Market Microstructure

The Global "Carbon-Energy-Intelligence" Framework: Decoding Cross-Market Interlinkages

M Tao|S Poletti|D Roubaud|AK Tiwari

University of Auckland Indian Institute of Management Bodh Gaya Montpellier Business School

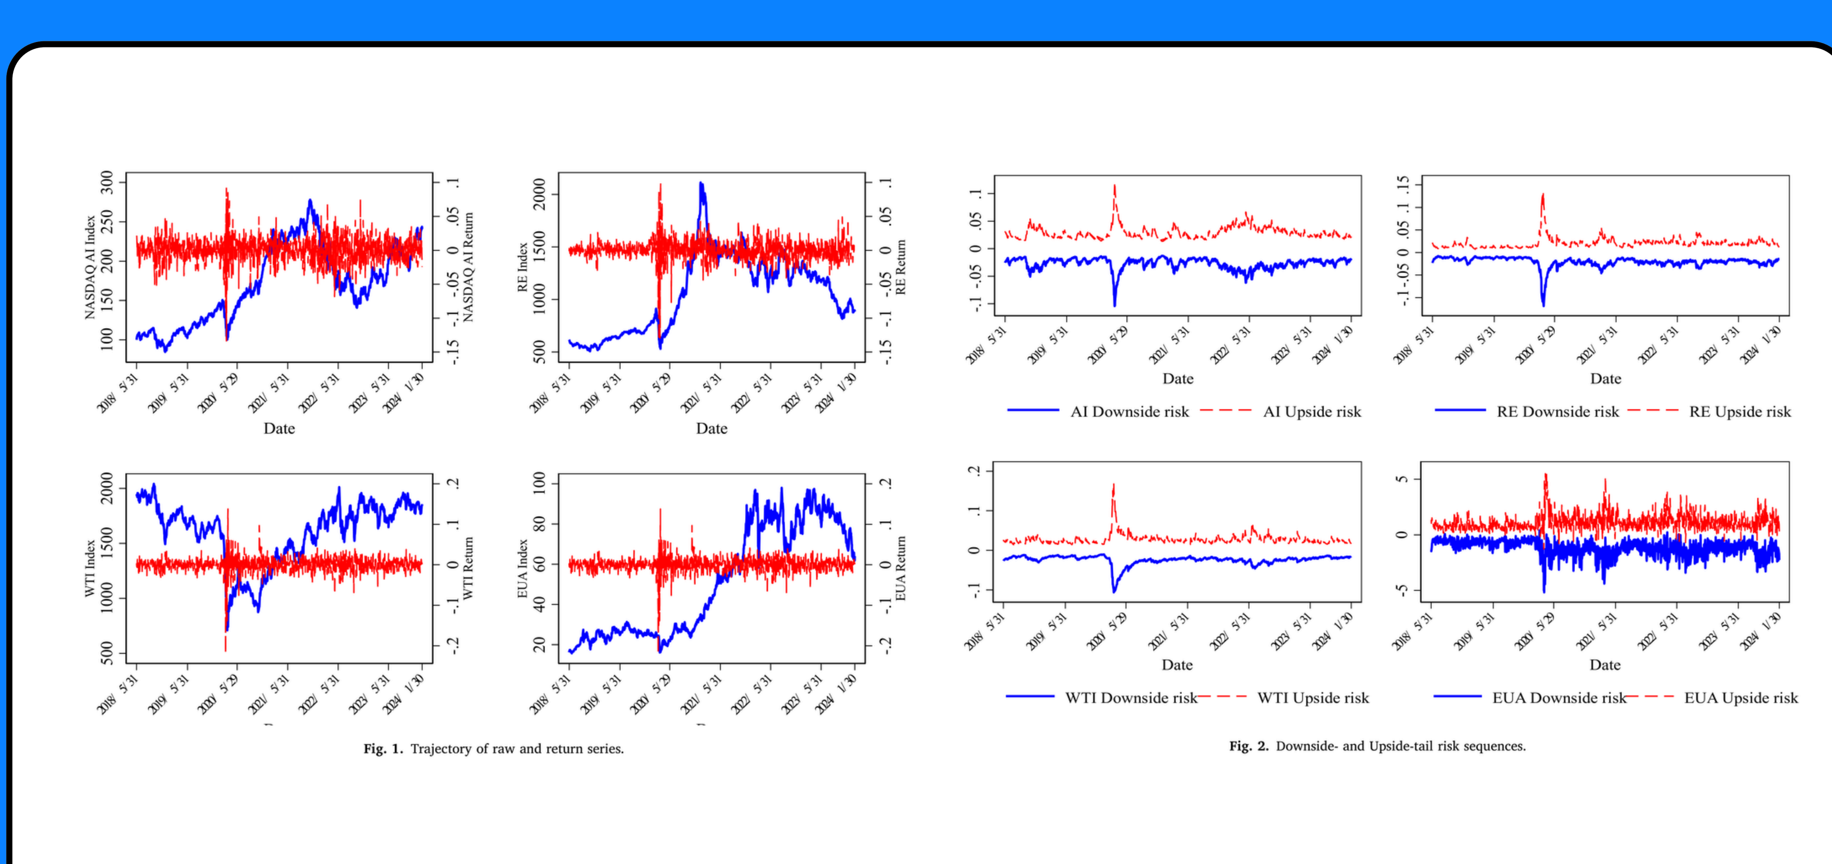

This study links AI stocks, oil, renewables, and EU carbon to track how shocks move across markets. Using daily data since 2018 and time varying Granger tests, TVP SV VAR, and CAViaR plus QVAR, it finds two way causality that intensifies in crises. AI lifts oil persistently but has mixed, regime dependent effects on renewables and carbon. Tail risk spikes systemwide. AI transmits most tail risk. EUA and oil receive, with EUA absorbing extremes better.

Commodity Market | AI-Enhanced Trading and Algorithmic Techniques

ZANISTA

Commodity trading: Evidence from traders' profits and an executive survey

Google Scholar SSRN

Antti Belt|Antti Kaskela|Matti Suominen
Boston Consulting Group Aalto University

The study shows hedge funds and merchant traders profit differently, fueled by volatility and imbalance. Using hedge fund returns and nine merchant firms, it finds funds earn from carrying commodity risk and 12 to 1 momentum, with factors explaining 43 percent of monthly variation and gains peaking when industry profit margins diverge. Merchant profits rise with short term volatility and a Revenue over SG&A squared proxy, but fall with steeper curves. A 400 survey aligns.

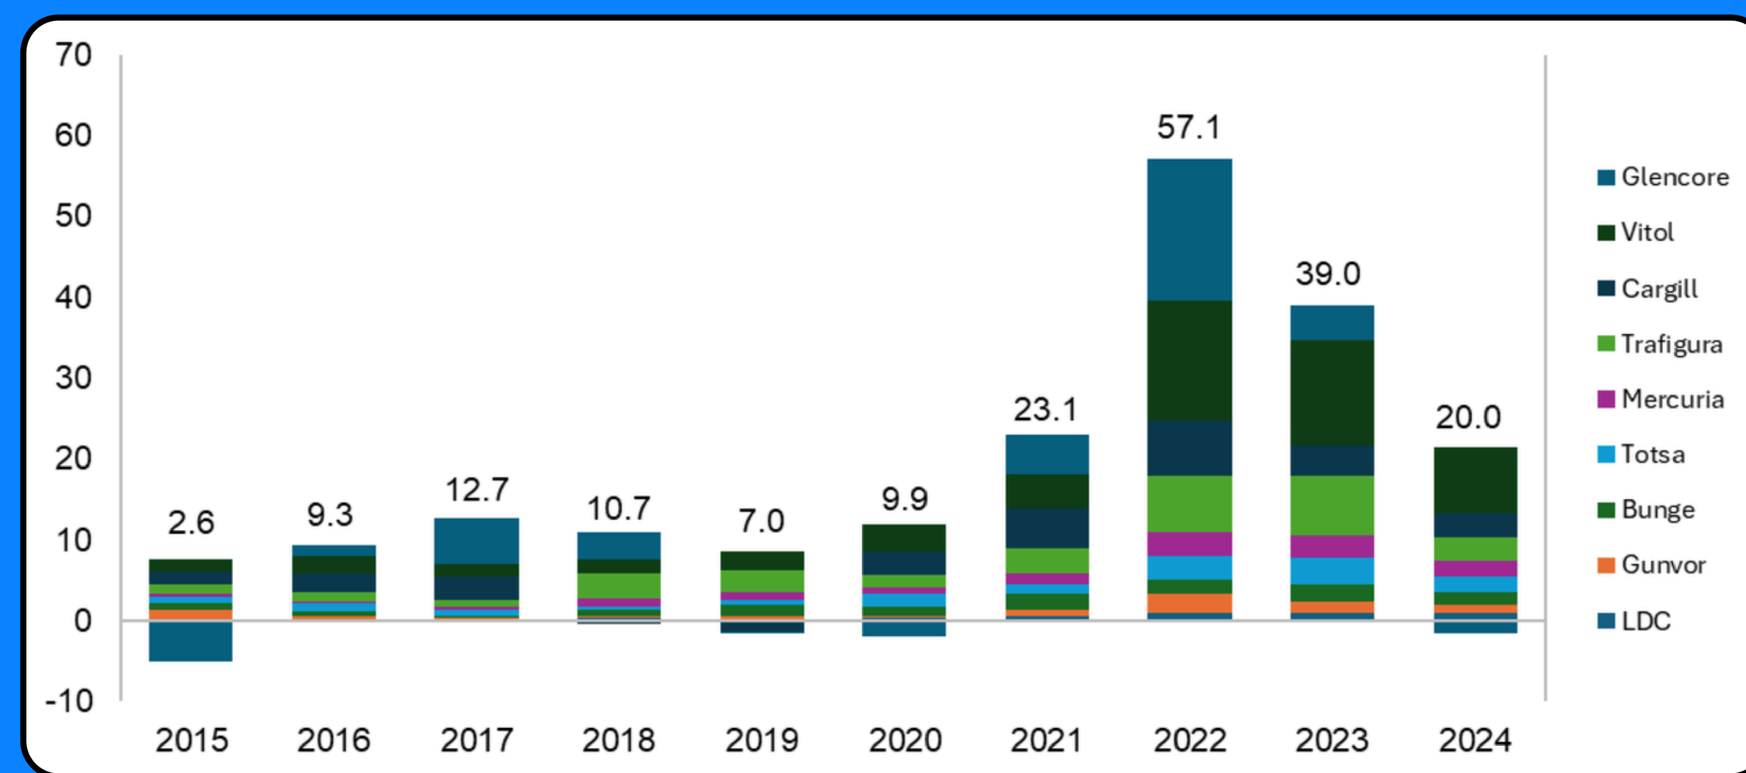

Commodity Market | Trading and Market Microstructure

FCM Concentration, Mandatory Clearing, and Commodity Futures Market

Google Scholar SSRN

Shengwu Du, Ketan Patel, Sean Wijesekera
Federal Reserve Board, Federal Reserve Bank of Chicago

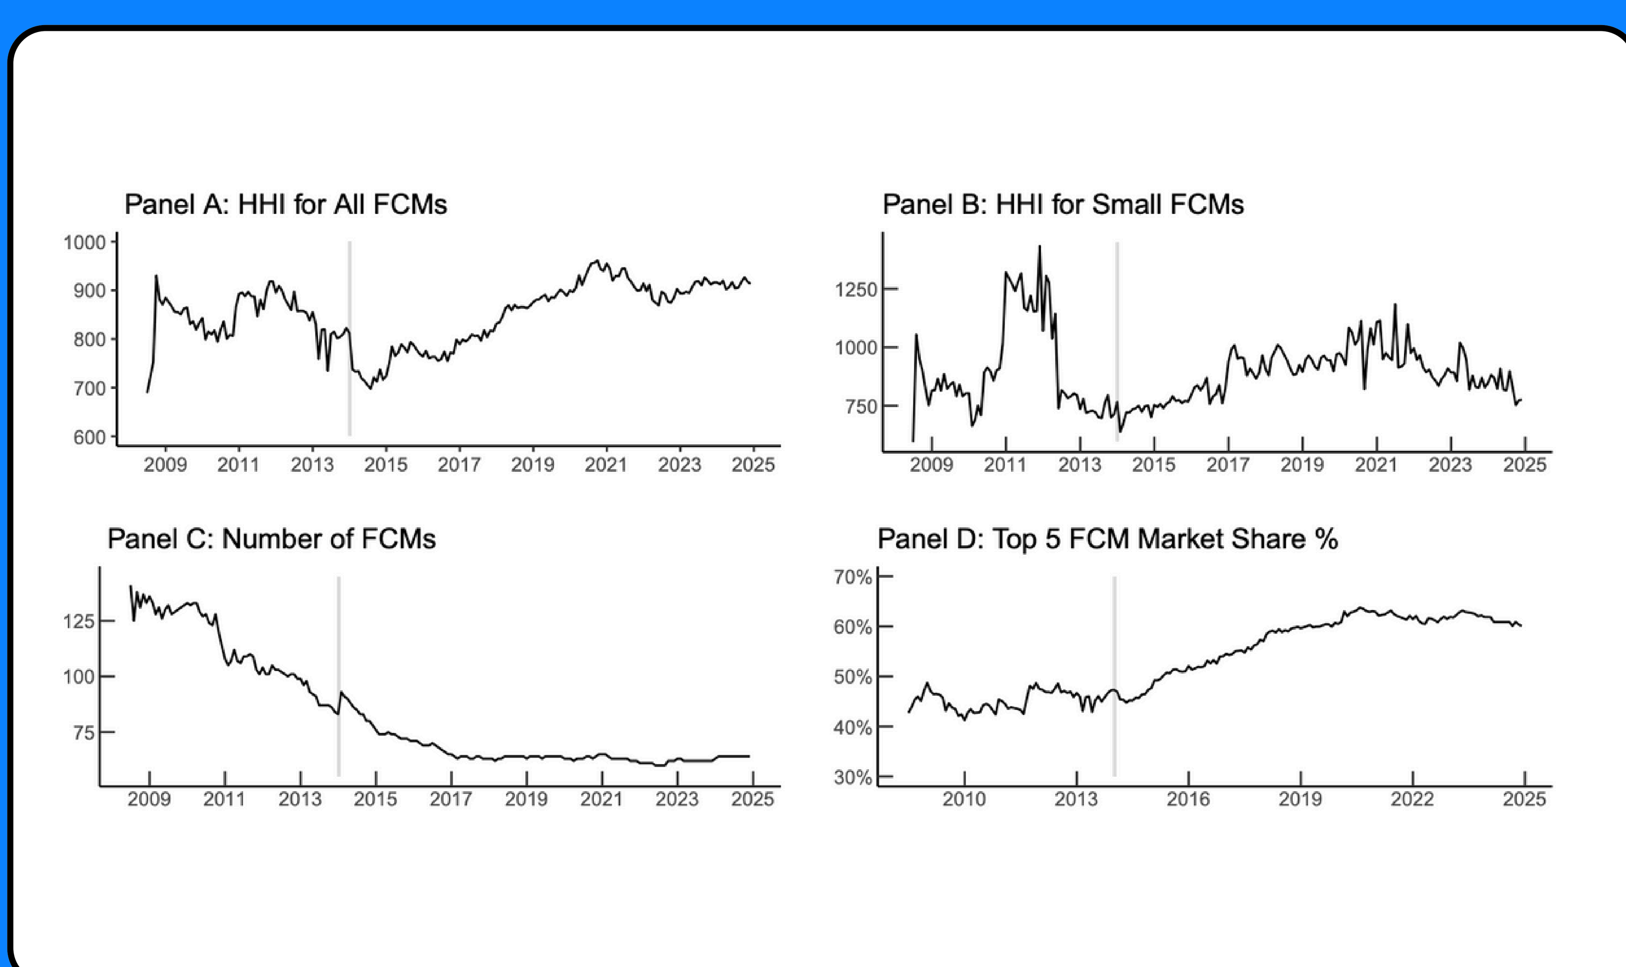

This study examines the effects of increasing concentration among Futures Commission Merchants (FCMs) in commodity futures markets, driven by post-crisis mandatory central clearing reforms. It finds no broad market-wide impact on volatility, liquidity, or price discovery, but identifies a significant reduction in small trader participation in several markets, highlighting distributional effects despite overall market resilience.

Commodity Market

Is Social Media Information Noise or Fundamentals? Evidence from the Crude Oil Market

Google Scholar SSRN

Gaoping Ma, Alireza Tourani-Rad, Yahua Xu, Z. Ivy Zhou
Central University of Finance and Economics, Auckland University of Technology, University of Wollongong

This study investigates whether social media sentiment predicts crude oil returns through fundamental information or noise. Positive sentiment reflects fundamentals and has lasting effects, while negative sentiment is noise-driven and leads to temporary price changes. Social media also provides incremental information beyond traditional news. Trading strategies based on sentiment yield significant economic gains.

	i=1	i=3	i=5	i=10	i=22	i=44	i=66
Intercept	-0.001	0.007	0.017**	0.045***	0.074***	0.146***	0.150***
	(-0.45)	(1.31)	(1.98)	(2.52)	(3.04)	(3.92)	(2.55)
$\Delta PosSent_{t-1}$	0.000	0.014**	0.022*	0.022**	0.023**	0.038**	0.035**
	(0.02)	(2.21)	(1.76)	(2.25)	(2.19)	(2.08)	(2.04)
$\Delta NegSent_{t-1}$	-0.024*	0.003	0.015	0.007	0.029	0.061**	0.067**
	(-1.74)	(0.33)	(0.72)	(0.46)	(1.35)	(2.24)	(1.98)
$R_{t-i,t-1}$	0.070	0.067	0.090	-0.060	0.091	-0.216*	-0.320
	(1.45)	(-0.03)	(0.90)	(-0.58)	(1.07)	(-1.79)	(-4.67)
$Volatility_{t-1}$	0.767***	0.885***	0.754***	1.412***	2.589***	2.580***	2.277***
	(18.26)	(9.59)	(4.73)	(3.27)	(7.32)	(4.09)	(4.43)
$Volume_{t-1}$	-1.009**	-2.011	-3.714	-9.766*	-14.783**	-26.742***	-24.110***
	(-1.85)	(-1.60)	(-1.61)	(-1.80)	(-2.56)	(-3.54)	(-3.25)
$Weekday_t$	Control	Control	Control	Control	Control	Control	Control
R^2 (%)	10.95	5.02	2.88	7.09	9.28	10.77	13.29

Commodity Market

ZANISTA

Does Home Bias Amplify the Disposition Effect? Evidence from Retail Forex Trading

Guiming Han, Alex Preda
King's College London, Lingnan University

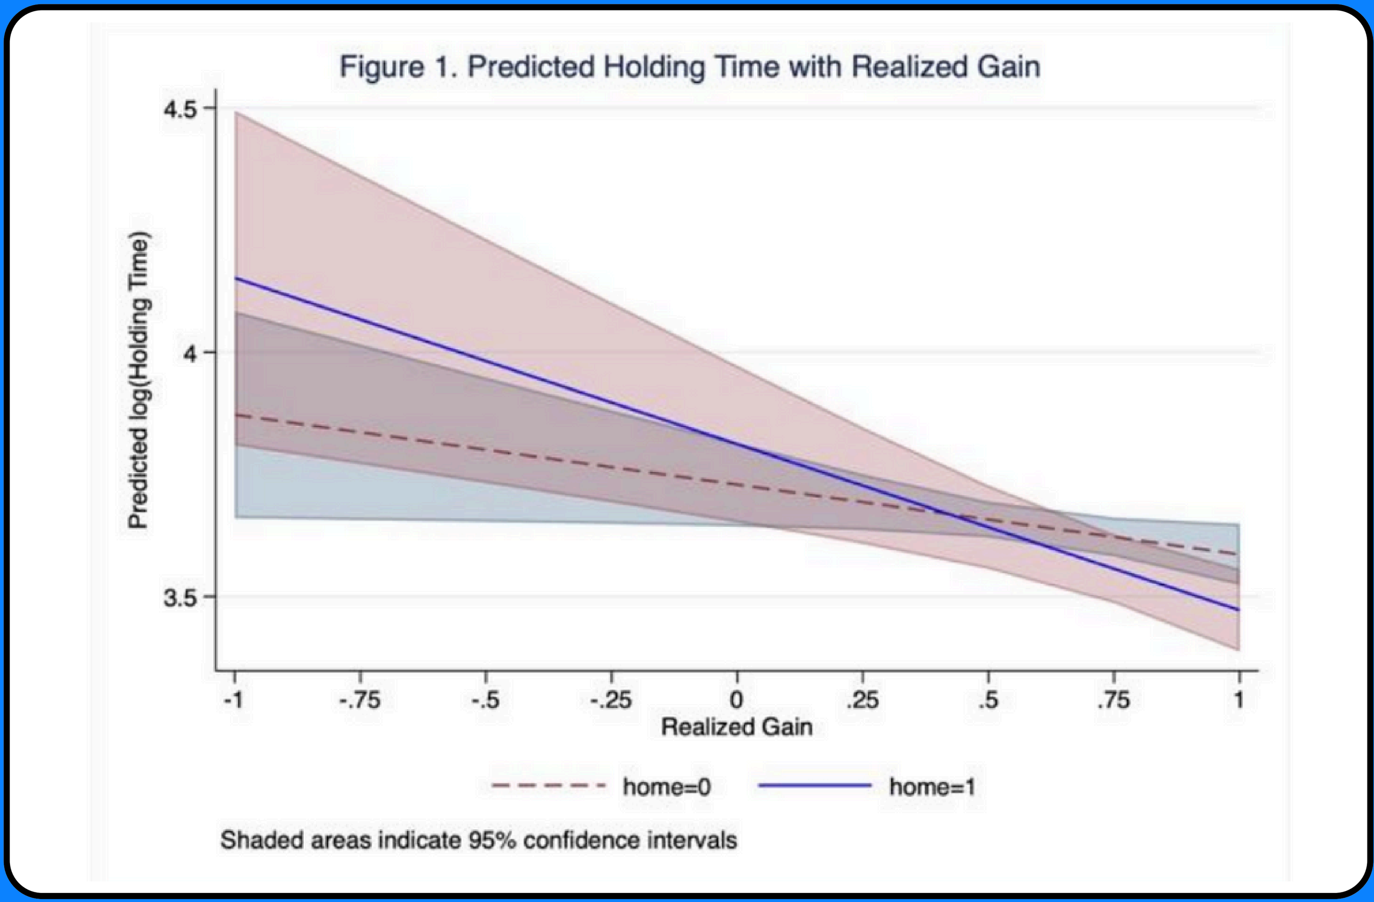

The paper examines whether home bias amplifies the disposition effect in retail forex trading. Using over one million trades by 4,226 investors across 126 countries, results show investors sell gains faster and hold losses longer when trading home currencies. Effects are strongest among older investors, high-balance accounts, and U.S. traders, demonstrating familiarity-driven behavioral distortions with significant implications for financial decision-making

FX (Foreign Exchange)

Hedge funds and the Treasury cash-futures basis trade

D Barth|RJ Kahn
Federal Reserve System

The paper finds hedge funds' Treasury cash futures basis trade is central, topping $1 trillion in gross exposure. Basis traders make up over 60 percent of hedge fund Treasury positions and about 70 percent of repo use. Strikingly, wider basis spreads coincide with larger positions when margins, haircuts, and repo limits bind, as in March 2020. Using data, tests, and a simple model, the study maps financing frictions to systemic risk. Limits include identification issues.

- **Treasury Basis Dominance**: Hedge funds Treasury cashfutures basis trade is vast, at times topping $1 trillion in gross exposures, dominating activityover 60% of Treasury positions and 70% of repo usemaking it marketplumbing central.
- **Arbitrage Frictions**: Strikingly quantities and basis spreads move together when frictions bitedefying textbook arbitrageas margins, haircuts, and repo constraints slow convergence and can amplify stress, a pattern visible during March 2020 turmoil.
- **Systemic Risk Link:** Study quantifies basis trading dominance and links financing frictions to systemic risk using regulatory data, empirical tests, simple model; strong measurement, policy relevance; limited causality, leverageterm granularity, Treasuriesonly generalizability noted

Fixed Income and Bonds Markets | Fixed Income and Bonds Markets

Analytical fixed income pricing in discrete time: A new family of models

L Stentoft|X Ye
University of Western Ontario Western University

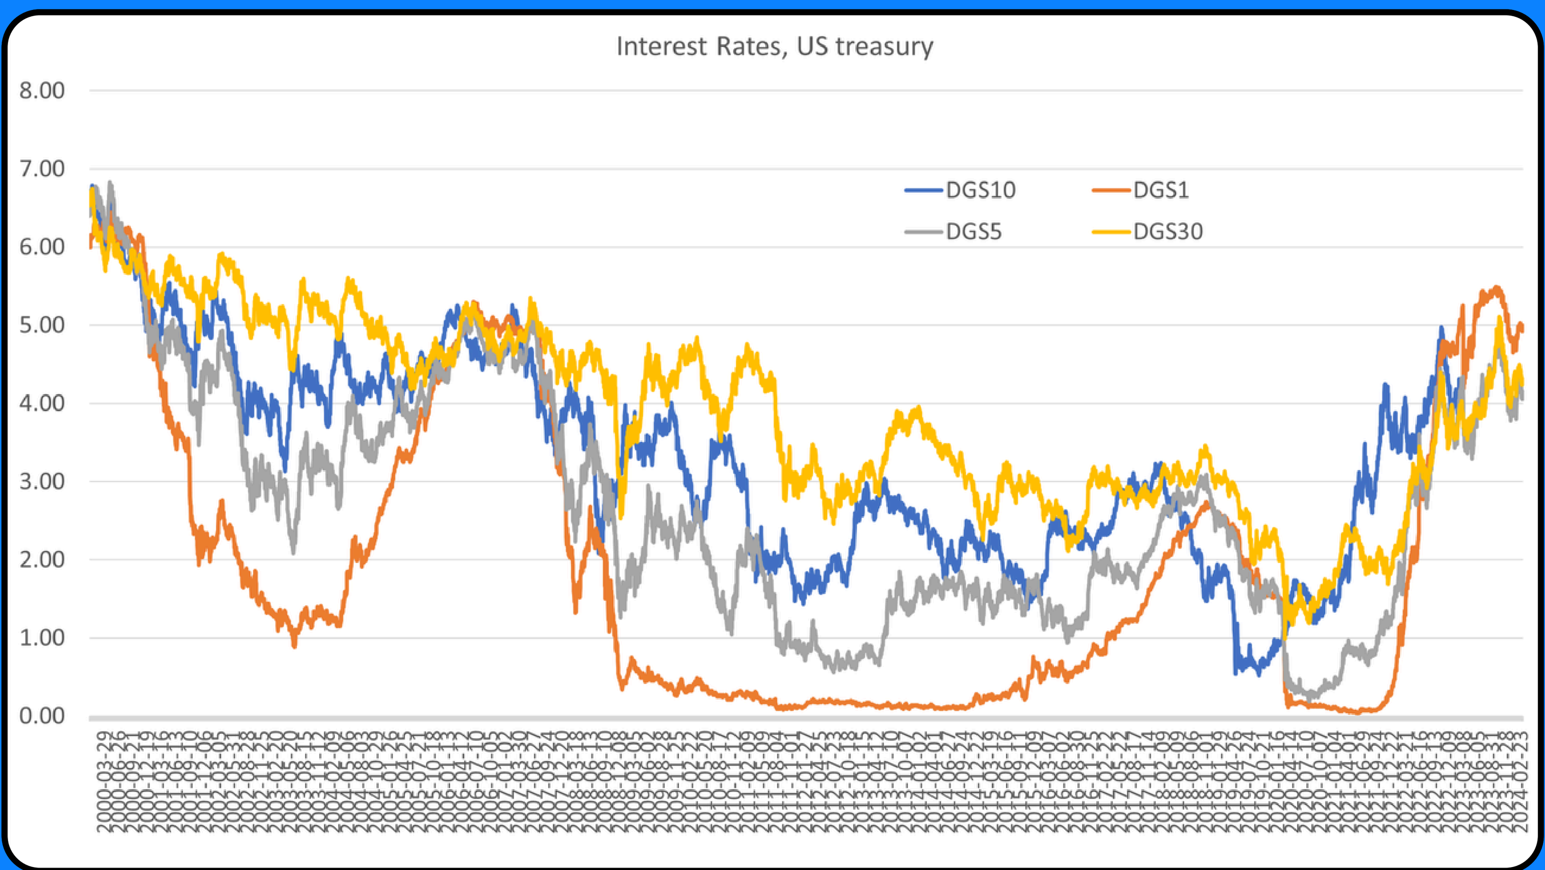

Researchers propose discrete time affine term structure models that keep closed form bond pricing while allowing a floor for interest rates, including negatives. The leading factor design mixes Gaussian dynamics with a noncentral chi square component. Estimated by maximum likelihood on daily U.S. Treasuries and monthly G7 rates, it beats Vasicek autoregressions and CIR on likelihood and AIC or BIC. It stays tractable. Gaps include out of sample tests, identification, robustness checks, and scaling costs.

Fixed Income and Bonds Markets | Fixed Income and Bonds Markets

ZANISTA

Multivariate Network Modeling: A Novel Approach To Corporate Bond Pricing

Tomaso Aste, Davide Controzzi

University College London, Splenetra Ltd.

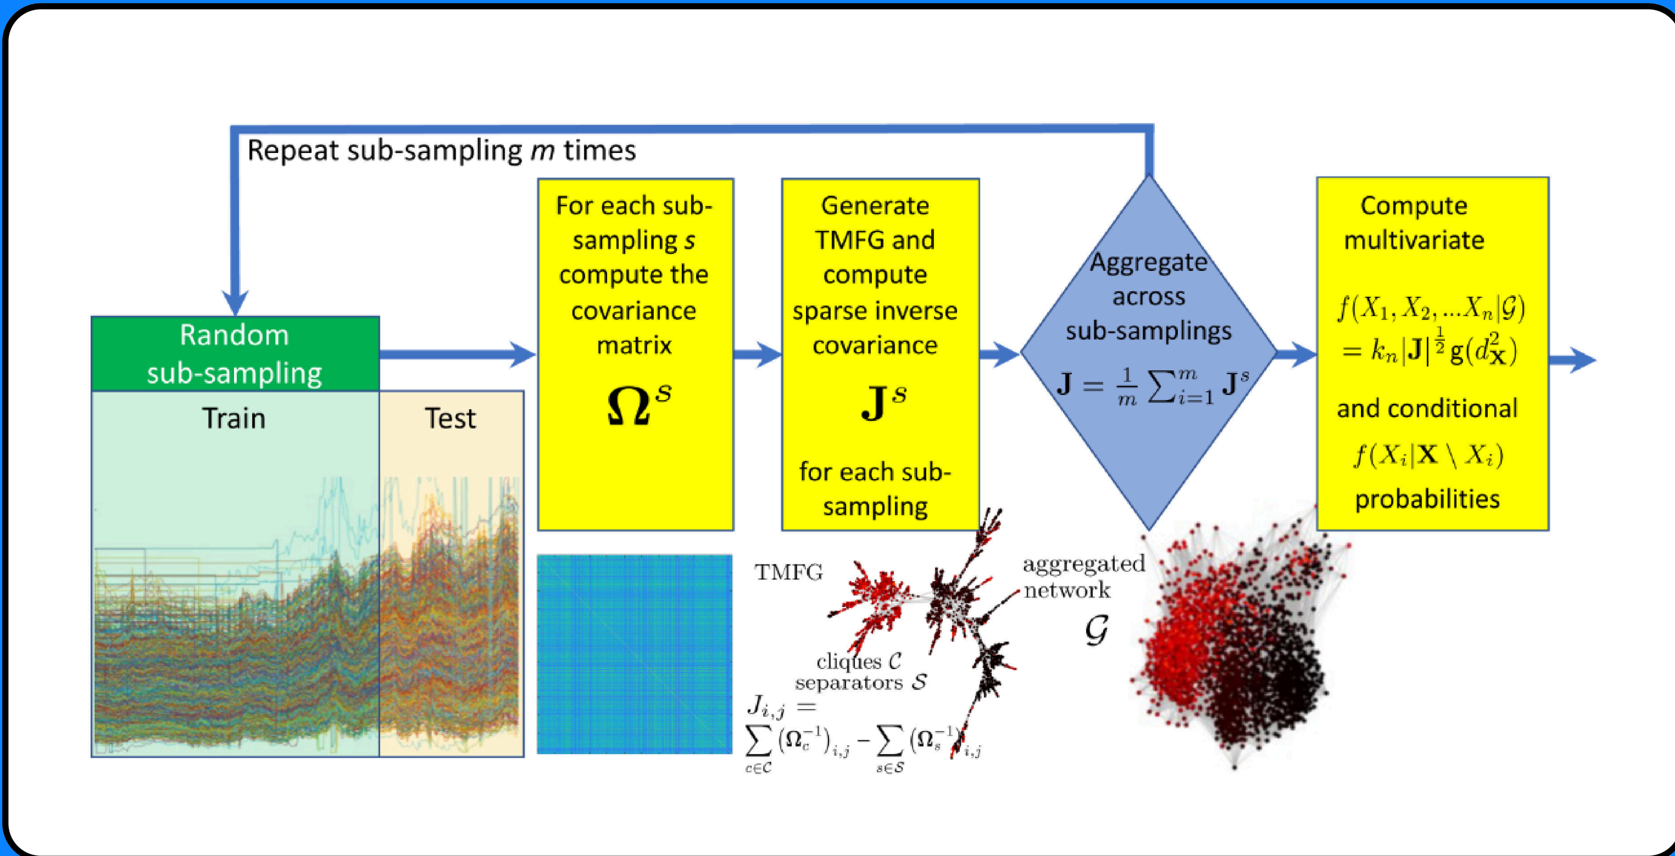

This paper introduces a novel multivariate network model combining sparse information filtering networks with Student-t distributions to price corporate bonds. It captures complex interdependencies between bonds, outperforming traditional factor models by providing accurate price predictions and confidence intervals, especially for illiquid bonds, using a purely data-driven approach without predefined factors.

Fixed Income and Bonds Markets

Twitter-based economic uncertainty and corporate bond credit spreads

Ali K. Malik, Gonul Colak

University of Leicester, University of Sussex, Hanken School of Economics

This study examines the impact of Twitter-based economic uncertainty (TEU) on corporate bond credit spreads. It finds a significant positive relationship, with TEU increasing spreads, especially for non-investment grade bonds and post-Volcker Rule. TEU also affects liquidity, reducing inter-dealer activity but increasing customer trading during uncertainty.

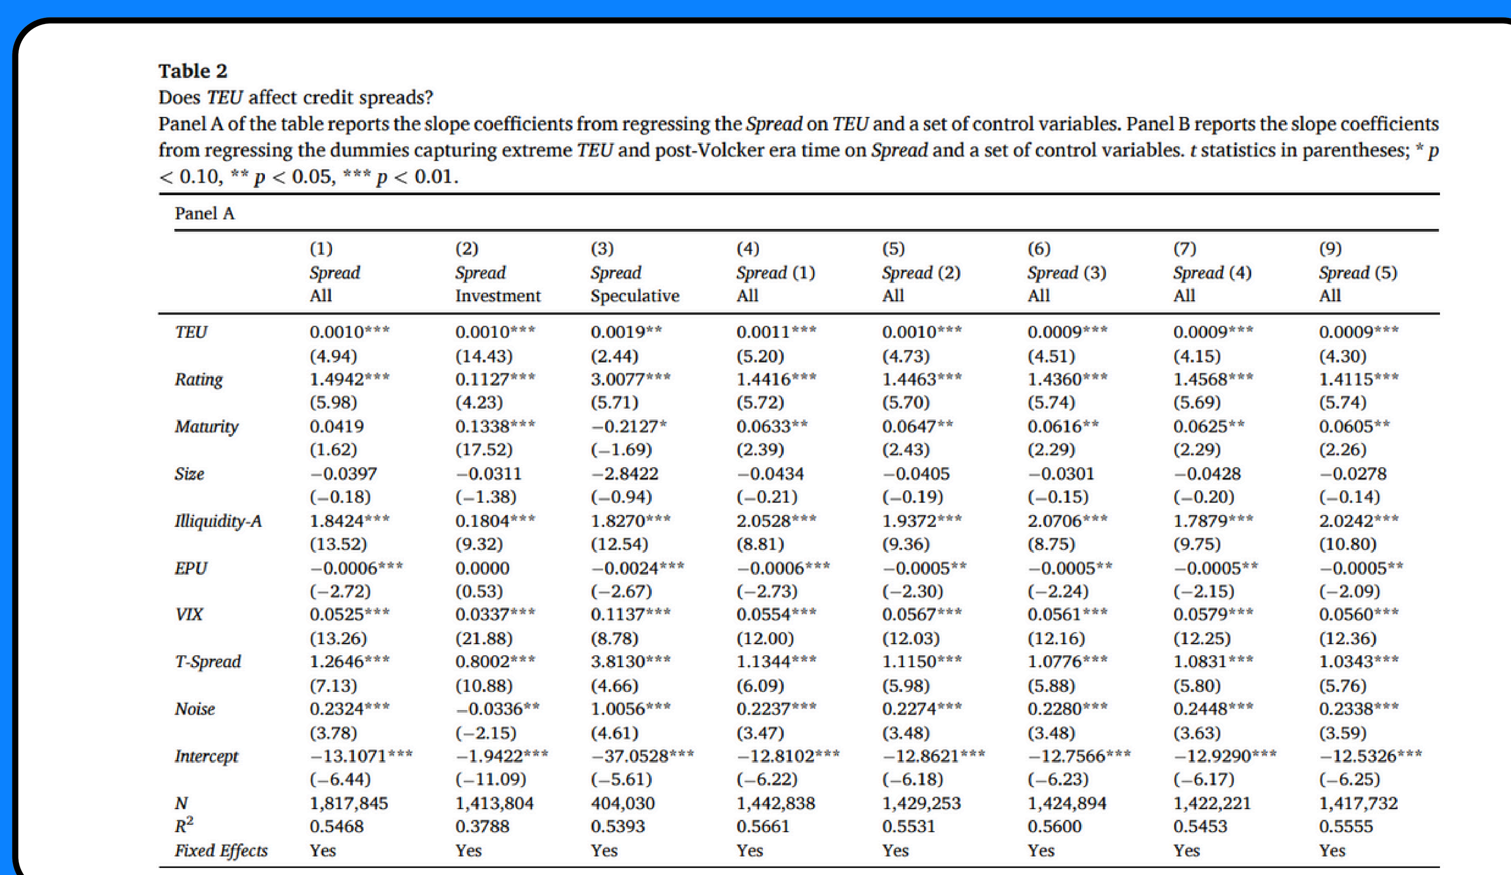

Table 2

Does *TEU* affect credit spreads?

Panel A of the table reports the slope coefficients from regressing the *Spread* on *TEU* and a set of control variables. Panel B reports the slope coefficients from regressing the dummies capturing extreme *TEU* and post-Volcker era time on *Spread* and a set of control variables. t statistics in parentheses; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A

	(1) *Spread* All	(2) *Spread* Investment	(3) *Spread* Speculative	(4) *Spread* (1) All	(5) *Spread* (2) All	(6) *Spread* (3) All	(7) *Spread* (4) All	(9) *Spread* (5) All
TEU	0.0010***	0.0010***	0.0019**	0.0011***	0.0010***	0.0009***	0.0009***	0.0009***
	(4.94)	(14.43)	(2.44)	(5.20)	(4.73)	(4.51)	(4.15)	(4.30)
Rating	1.4942***	0.1127***	3.0077***	1.4416***	1.4463***	1.4360***	1.4568***	1.4115***
	(5.98)	(4.23)	(5.71)	(5.72)	(5.70)	(5.74)	(5.69)	(5.74)
Maturity	0.0419	0.1338***	−0.2127*	0.0633**	0.0647**	0.0616**	0.0625**	0.0605**
	(1.62)	(17.52)	(−1.69)	(2.39)	(2.43)	(2.29)	(2.29)	(2.26)
Size	−0.0397	−0.0311	−2.8422	−0.0434	−0.0405	−0.0301	−0.0428	−0.0278
	(−0.18)	(−1.38)	(−0.94)	(−0.21)	(−0.19)	(−0.15)	(−0.20)	(−0.14)
Illiquidity-A	1.8424***	0.1804***	1.8270***	2.0528***	1.9372***	2.0706***	1.7879***	2.0242***
	(13.52)	(9.32)	(12.54)	(8.81)	(9.36)	(8.75)	(9.75)	(10.80)
EPU	−0.0006***	0.0000	−0.0024***	−0.0006***	−0.0005**	−0.0005**	−0.0005**	−0.0005**
	(−2.72)	(0.53)	(−2.67)	(−2.73)	(−2.30)	(−2.24)	(−2.15)	(−2.09)
VIX	0.0525***	0.0337***	0.1137***	0.0554***	0.0567***	0.0561***	0.0579***	0.0560***
	(13.26)	(21.88)	(8.78)	(12.00)	(12.03)	(12.16)	(12.25)	(12.36)
T-Spread	1.2646***	0.8002***	3.8130***	1.1344***	1.1150***	1.0776***	1.0831***	1.0343***
	(7.13)	(10.88)	(4.66)	(6.09)	(5.98)	(5.88)	(5.80)	(5.76)
Noise	0.2324***	−0.0336**	1.0056***	0.2237***	0.2274***	0.2280***	0.2448***	0.2338***
	(3.78)	(−2.15)	(4.61)	(3.47)	(3.48)	(3.48)	(3.63)	(3.59)
Intercept	−13.1071***	−1.9422***	−37.0528***	−12.8102***	−12.8621***	−12.7566***	−12.9290***	−12.5326***
	(−6.44)	(−11.09)	(−5.61)	(−6.22)	(−6.18)	(−6.23)	(−6.17)	(−6.25)
N	1,817,845	1,413,804	404,030	1,442,838	1,429,253	1,424,894	1,422,221	1,417,732
R^2	0.5468	0.3788	0.5393	0.5661	0.5531	0.5600	0.5453	0.5555
Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Fixed Income and Bonds Markets

Predicting Rating Transitions using Machine Learning

V Kumar|K Giesecke|P Kamenski

Stanford University

- **ML Beats Matrices**: Machine learning beats traditional matrices at predicting rating migrations, capturing non-linear default-risk patterns, improving forecasts and aiding portfolio allocation, capital planning, and risk management across banks, asset managers, hedge funds
- **Predictive Migration Framework**: Authors frame rating migration as a predictive learning task, benchmark algorithms against statistical baselines, propose: issuer/macro/market features; out-of-sample validation; horizon-specific transition probabilities, yielding asymmetric gains in downgrades and fallen-angel risk.
- **Implications and Risks:** risks include overfitting, interpretability, procyclicality. Despite applied focus and links, paper needs ablation-studies and regulatory checks

Machine learning models predict credit rating transitions better than legacy matrices by capturing nonlinear patterns in migration and default risk. The paper casts rating changes as a predictive task and reports gains for banks and investors in forecasting and stress tests. Benefits look strongest for downgrades and fallen angels. Design choices include issuer, macro, and market features, out of sample validation, and horizon probabilities. Limits include disclosure, overfitting, interpretability, procyclicality, and need for regulatory tests.

Credit and Debt Markets

Machine Learning in Finance

ZANISTA

Attention Deficits and Asset Prices: A Theory of Information Backlogs

Nicole Beevers, Hasan Fallahgoul
Monash University

This study introduces "information backlogs" as a new source of return predictability, occurring when financial news flow exceeds investors' attention capacity. Using YouTube data from major news providers, the authors show that attention deficits cause immediate positive returns but create persistent backlogs that depress future returns and reduce volatility less effectively, revealing a market-wide cognitive constraint that cannot be arbitraged away.

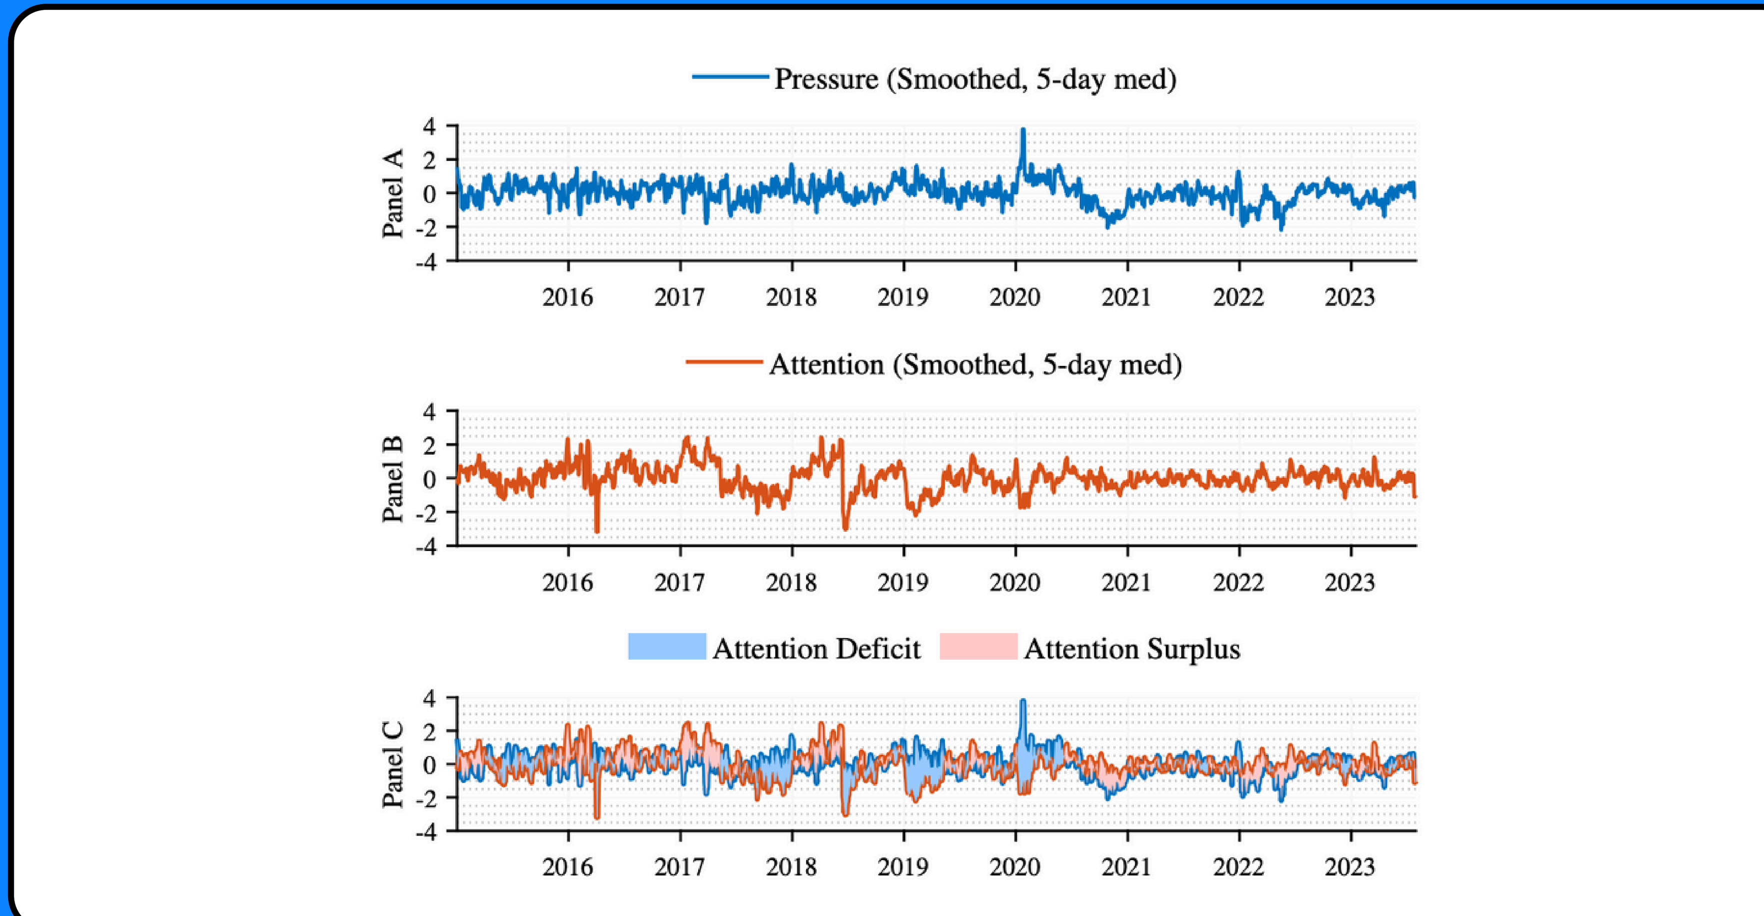

Portfolio Optimization and Market Prediction

Deep portfolio selection with contrastively aligned cross-modal attention

H Huang|R Liu|Q Wu
King's College London City University of Hong Kong

- **Caracal Outperforms**: Caracal debuts as a contrastively aligned cross-modal attention model fusing price time series with textual signals and optimizing joint return distributions, outperforming across three datasets versus leading baselines in benchmarks.
- **Architecture Highlights**: Key modules: Inter-Modal Contrastive Fusion aligning text--numeric pairs via attention; Inner-Modal Contrastive Learning capturing coarse and fine temporal structure from similarity graphs; and a distribution-aware objective, with ablations validating results.
- **Insights and Caveats:** Insights include cross- and intra-modal alignment reduces noise in scarce data, sharpening signals when assets co-move. Caveats include dataset/regime detail, text sensitivity, overhead, unclear handling of transaction costs and turnover

Contrastively Aligned Cross-modal (Cross)-Attention Model (Caracal) is a contrastively aligned cross modal attention model for portfolio selection that fuses price time series with text and optimizes the joint distribution of returns. It outperforms state of the art on three datasets. Components include inter modal contrastive fusion, inner modal contrastive learning from similarity graphs, and a distributional portfolio loss. Alignment cuts noise under data scarcity and similar moves. Ablations confirm impact. Limits include dataset detail, alignment latency, transaction costs, and compute.

Portfolio Optimization and Market Prediction | AI-Driven Potfolio Optimization

Tackling estimation risk in Kelly investing using options

Fabrizio Lillo|Piero Mazzarisi|Ioanna-Yvonni Tsaknaki
University of Bologna University of Siena Scuola Normale Superiore

Adding a European put to a binomial stock bond market does not lift long run Kelly growth when parameters are known. Kelly with options matches classical Kelly. Under estimation risk, misjudging the up factor can let hedged Kelly outperform, but only by misspecification direction. No single hedge dominates. The authors propose Kelly with Options (KOc), a convex blend that converges to the better hedge. They use a binomial tree, option pricing, Karush–Kuhn–Tucker conditions KKT, and simulations, and note asymptotic limits.

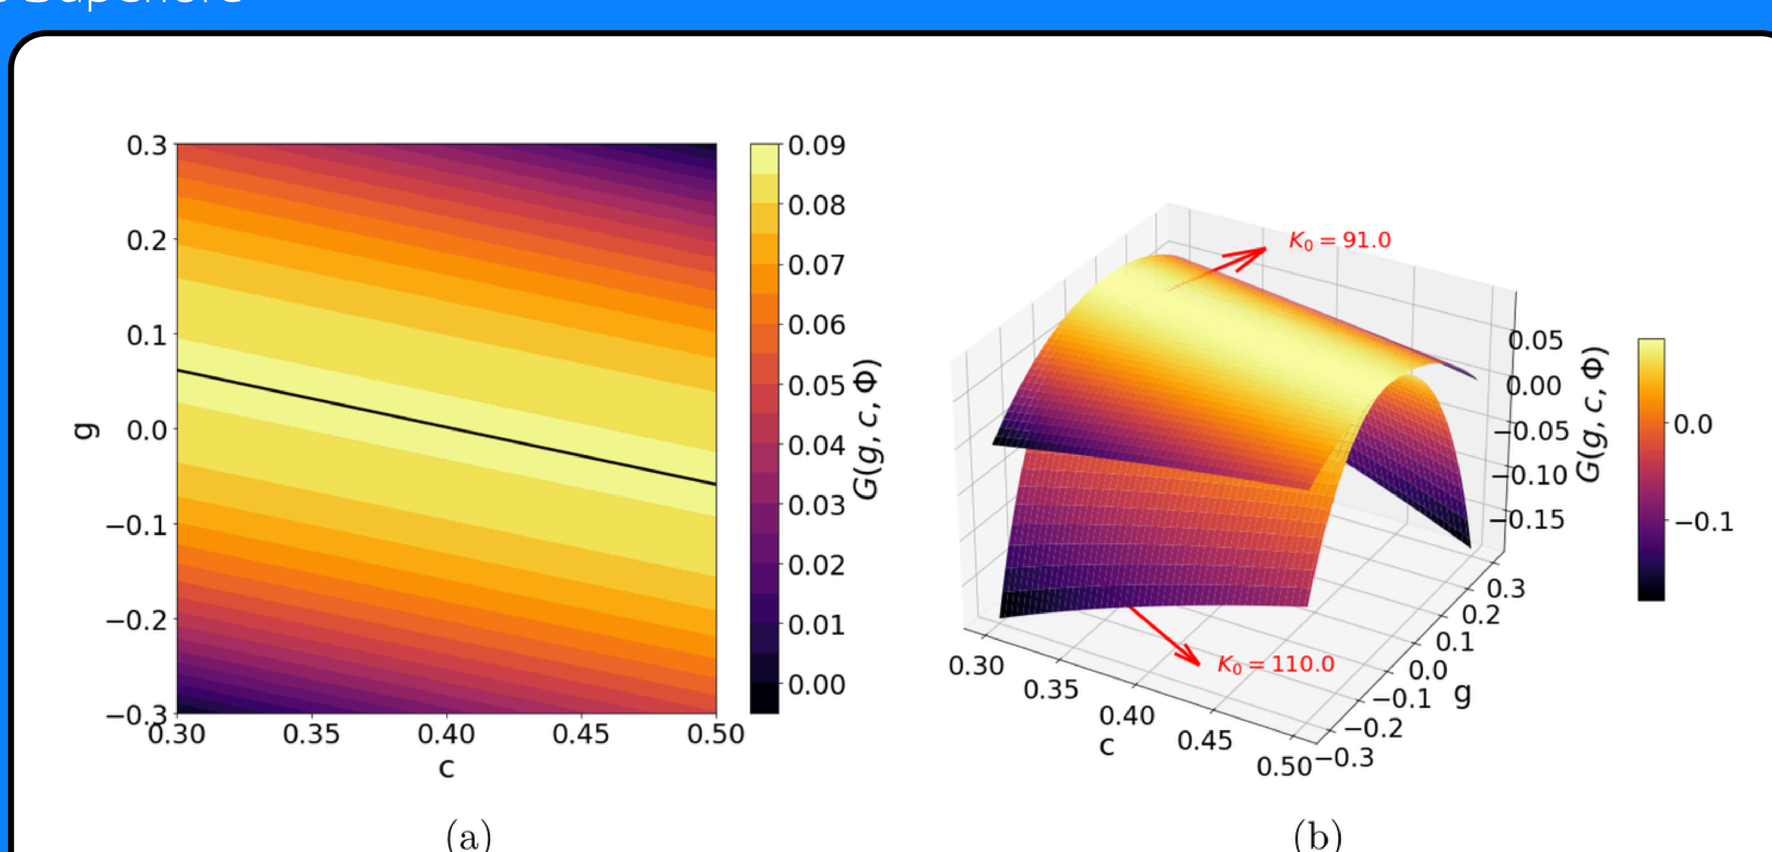

(a) (b)

Portfolio Optimization and Market Prediction | Mathematical Finance

ZANISTA

False Confidence in Systematic Trading: Illusion of Speed and Mirage of Performance

Daniel Alexandre Bloch
Quant Finance Ltd

#	ϵ_μ	n^T_μ (IID / Mod. / Strong)	Years	ϵ_{σ^2}	$n^T_{\sigma^2}$ (IID / Mod. / Strong)	Years
	0.40	24 / 34 / 48	0.10 / 0.13 / 0.19	0.70	17 / 26 / 51	0.07 / 0.10 / 0.20
	0.20	96 / 134 / 192	0.38 / 0.53 / 0.76	0.50	32 / 48 / 96	0.13 / 0.19 / 0.38
	0.10	384 / 538 / 768	1.5 / 2.1 / 3.1	0.20	193 / 290 / 579	0.8 / 1.1 / 2.3
	0.05	1,537 / 2,152 / 3,074	6.1 / 8.5 / 12.2	0.10	769 / 1,154 / 2,307	3.0 / 4.6 / 9.2
	0.01	38,416 / 53,782 / 76,832	152 / 214 / 306	0.05	3,074 / 4,611 / 9,222	12.2 / 18.3 / 36.6
				0.04	4,803 / 7,205 / 14,409	19.1 / 28.6 / 57.2
				0.02	19,208 / 28,812 / 57,624	76.3 / 114 / 229

Table 2: Required sample sizes for mean and variance estimation under IID, moderate, and strong dependence ($\alpha = 0.05$, $\sigma = 1$). Years computed assuming 252 trading days per year.

Short lookback windows in systematic trading create an illusion of speed and a mirage of performance. Their apparent responsiveness is often statistical noise from insufficient data, not genuine predictive power. This small-sample paradox makes performance metrics like Sharpe ratios highly unreliable. The article argues that true robustness requires longer validation horizons, not faster signals.

Portfolio Optimization and Market Prediction

High risk aversion Merton's problem without transversality conditions

Enrico Biffis, Cristina Di Girolami, Salvatore Federico, Fausto Gozzi
Imperial Business School; Università Alma Mater Studiorum Bologna; Libera Università degli Studi Sociali "Guido Carli"

This paper solves Merton's infinite-horizon consumption-investment problem for an agent with high risk aversion (γ>1). The authors provide a direct method to establish the existence of an optimal strategy and the form of the value function, avoiding the technical difficulties of verifying transversality conditions. Their approach uses the problem's homogeneity and a "half-verification" of the solution the Hamilton-Jacobi-Bellman equation.

3.4 The "half-verification" and the optimal feedback strategy

In this subsection we finalize our approach by providing the true optimality of the candidate optimal feedback map that comes out from the optimization in $\mathcal{H}_{cv}$.

Using (3.2), the candidate optimal feedback map provided by the maximization in (3.4) (see (3.5)) is the map $G:(0,\infty)\to\mathbb{R}^+\times\mathbb{R}$ defined by

$$G(x) = (G^c(x), G^\pi(x)) = \left(c_{\max}(x, V'(x), V''(x)),\ \pi_{\max}(x, V'(x), V''(x))\right) = \left(a^{-1/\gamma}x,\ \frac{\lambda}{\sigma\gamma}x\right),$$

where a is given in (3.15). Plugging this map in the state equation (2.2), we get the following closed loop equation associated to G

$$\begin{cases} dX_t = \left(r + \frac{\lambda^2}{\gamma} - a^{-1/\gamma}\right) X_t\, dt + \frac{\lambda}{\gamma} X_t\, dW_t, \\ X_0 = x > 0, \end{cases} \tag{3.16}$$

whose explicit solution is

$$\widehat{X}_t = x \exp\left[\left(r + \frac{\lambda^2}{\gamma} - a^{-1/\gamma} - \frac{\lambda^2}{2\gamma^2}\right)t + \frac{\lambda}{\gamma}W_t\right] = x\exp\left[\left(\frac{r-\rho}{\gamma} + \frac{\lambda^2}{2\gamma}\right)t + \frac{\lambda}{\gamma}W_t\right]. \tag{3.17}$$

Portfolio Optimization and Market Prediction

AI-Driven Potfolio Optimization

ONLINE SUPPLEMENT--Deep Learning Alpha Signals from Limit Order Books: Practical Insights and Lessons Learned (Main Article Published at Risk. net)

Petter N. Kolm|Nicholas Westray
New York University

Ticker	Updates (000)	Trades (000)	Price Changes (000)	Price (USD)	Spread (bps)	Volume (USD MM)
AMAT	611.97	18.05	17.15	47.36	2.62	91.72
AMD	1213.04	42.06	16.26	31.33	3.64	343.63
AMGN	137.56	11.06	25.61	198.36	3.85	138.33
AMZN	304.76	31.14	64.23	1795.51	2.59	1531.40
ANSS	63.55	3.37	15.22	206.76	9.39	25.86
ASML	147.63	3.31	21.21	223.92	7.06	46.42
ATVI	423.01	17.97	20.25	49.87	2.81	91.82
AVGO	136.68	12.20	32.10	289.33	4.62	182.39
BIDU	155.90	13.37	26.50	132.58	4.79	129.79
BIIB	88.20	9.80	23.17	265.36	6.01	129.65
BKNG	51.44	4.20	15.00	1881.52	8.68	169.08
BMRN	61.69	5.08	10.81	82.87	10.20	25.70
CDNS	128.56	7.32	15.40	64.99	4.28	32.49
CDW	53.73	4.20	9.73	112.53	7.93	28.02
CERN	117.85	7.57	10.59	66.41	3.53	38.58

The supplement finds LSTMs with stationary order flow imbalance features match complex models for short-horizon limit order book prediction. Attention Seq2Seq helps with long sequences, but LSTMs remain competitive. Contributions include a 10TB LOBSTER pipeline for 115 Nasdaq stocks, preprocessing, and alpha structures using a characteristic time clock. Evaluation uses out-of-sample R^2 against a return benchmark with latency buffers. Predictive power is modest yet meaningful. Limits include per-symbol models and a 2019 to 2020 window.

Market Microstructure, High Frequency Trading, Execution, and Limit Order Books (LOB)

Machine Learning for Option Pricing

ZANISTA

DiffVolume: Diffusion Models for Volume Generation in Limit Order Books

Zhuohan Wang, Carmine Ventre
King's College London

This paper introduces DiffVolume, a conditional diffusion model for generating realistic future Limit Order Book (LOB) volume snapshots. Conditioned on past data, time of day, and target liquidity, it outperforms GANs in realism, enables controllable counterfactual generation for extreme scenarios, and improves downstream liquidity prediction tasks when used for data augmentation.

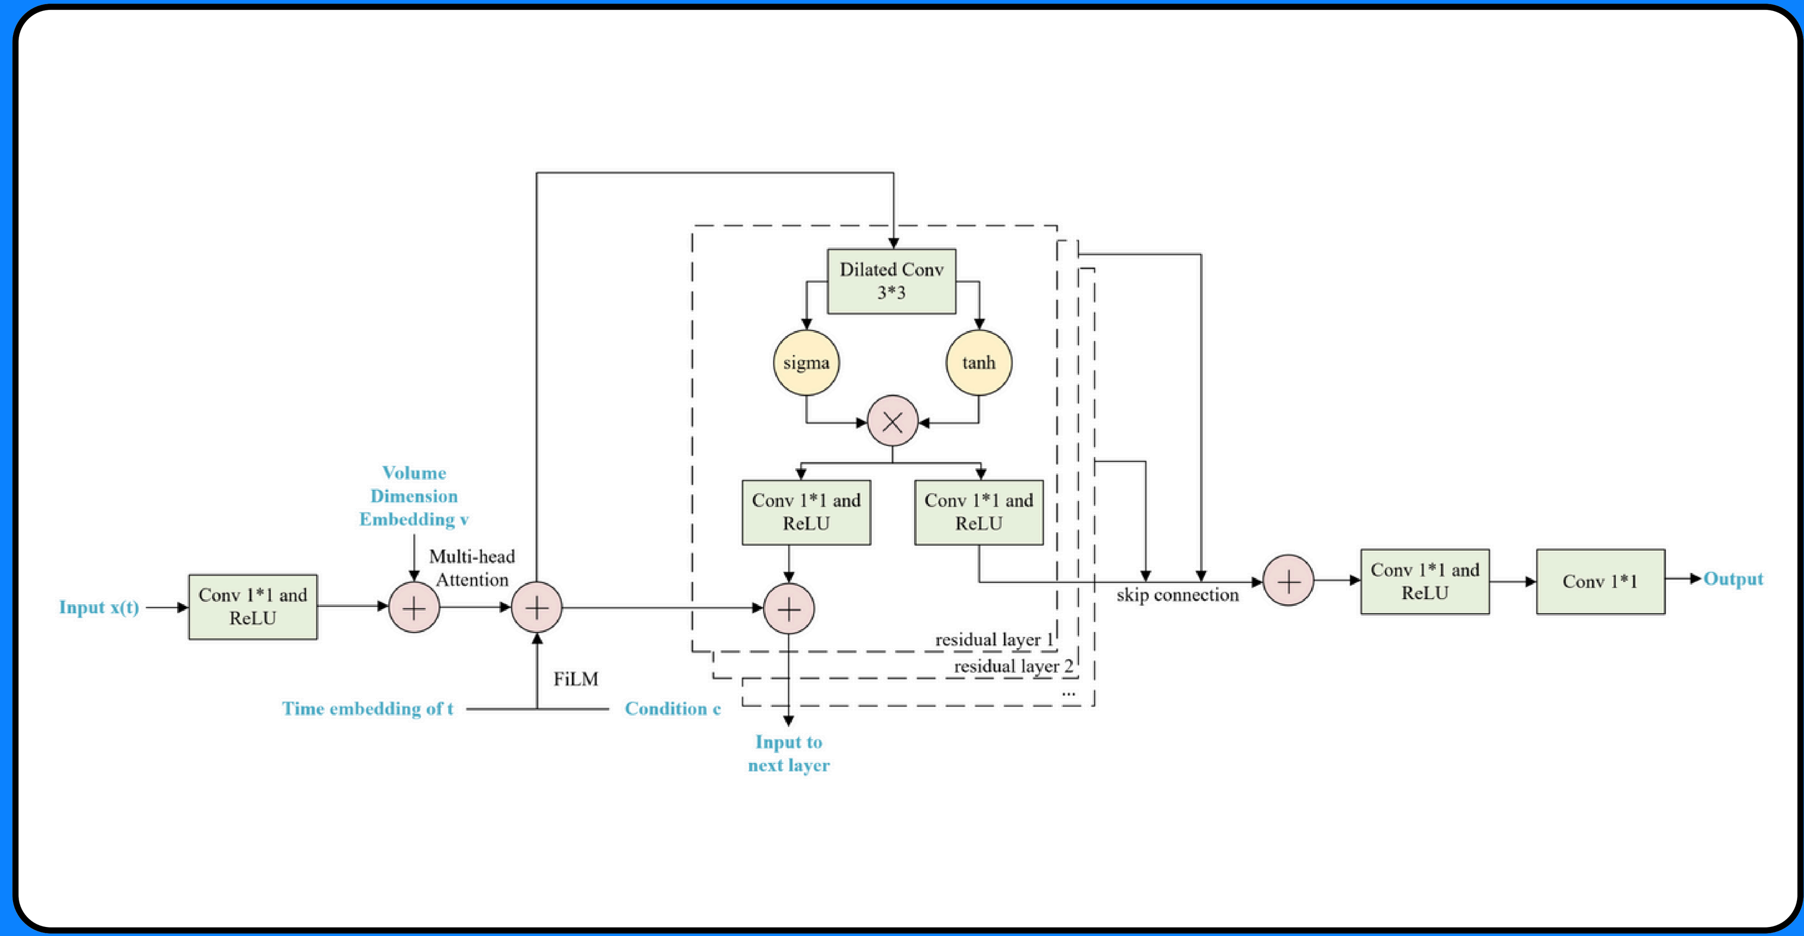

Market Microstructure, High Frequency Trading, Execution, and Limit Order Books (LOB)

Optimal Quoting under Adverse Selection and Price Reading

Alexander Barzykin|Philippe Bergault|Olivier Guéant|Malo Lemmel
CNRS Université Paris Dauphine-PSL Université Paris 1 Panthéon-Sorbonne

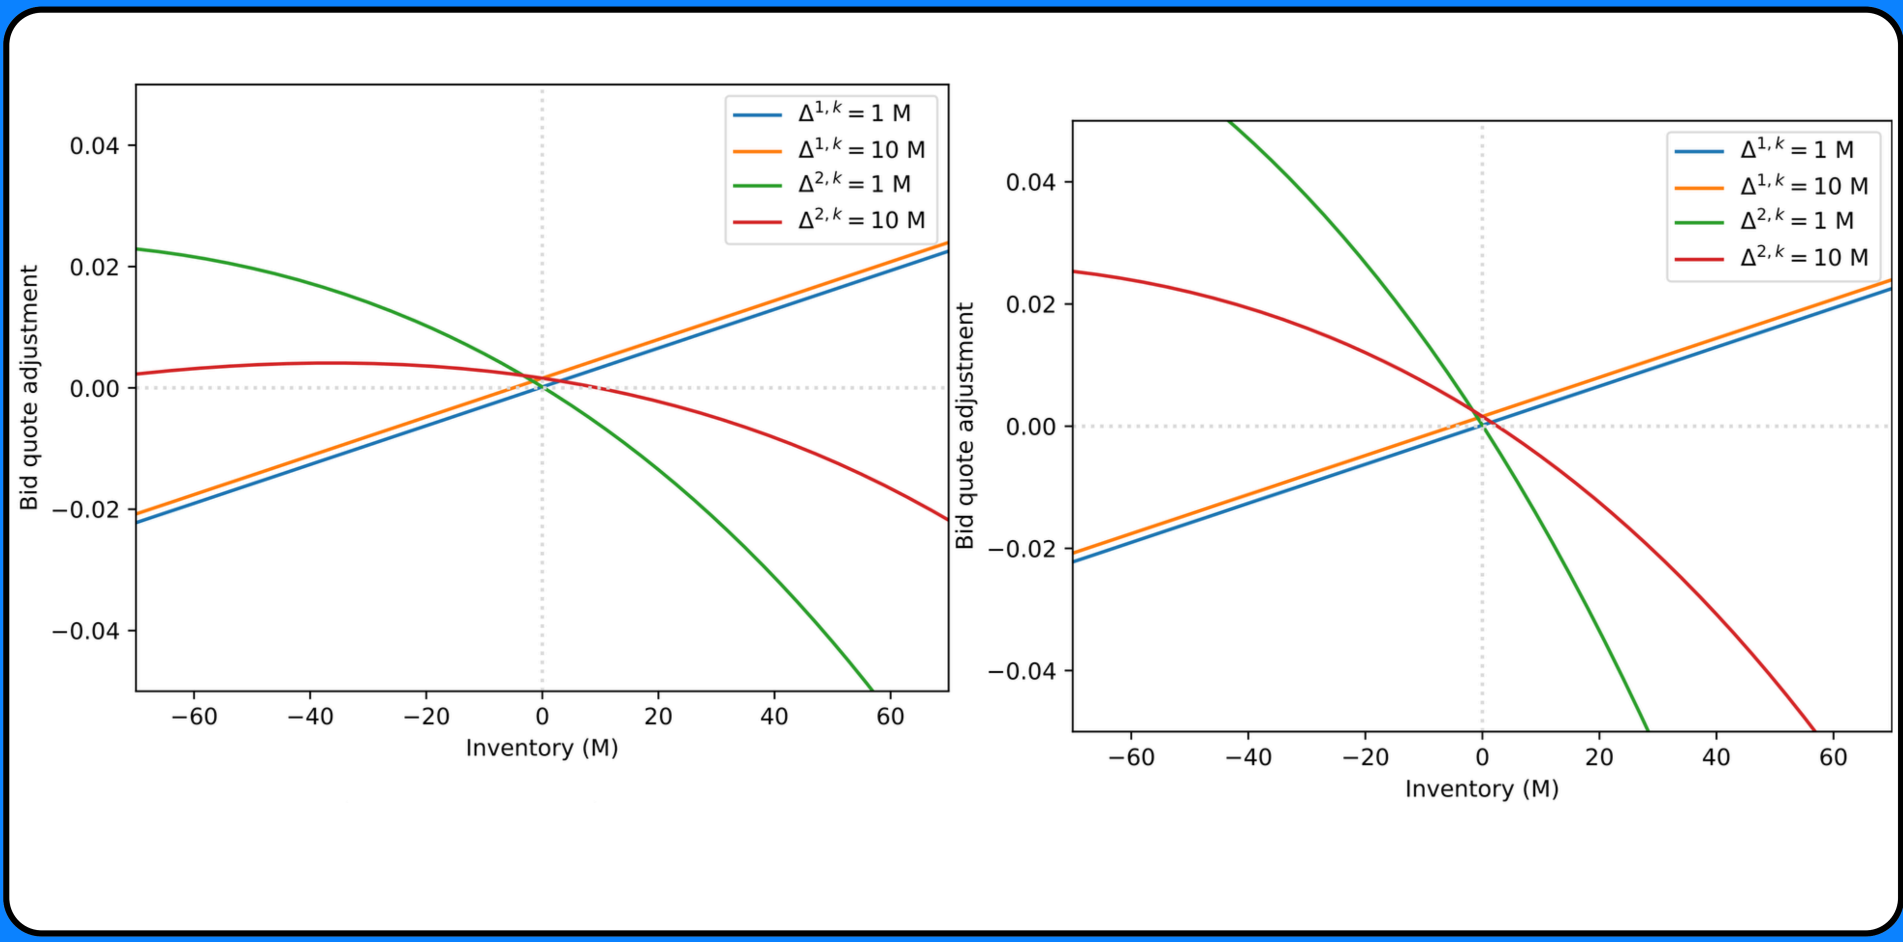

This paper proposes a practical quoting framework that embeds informational risk in market making. Optimal quotes get a two part adjustment. A global term widens spreads and moderates skew. A tier and size term tunes quotes to client toxicity. Price reading makes makers act more risk averse and show less skew. Adverse selection depends on signal speed. Slow signals invite tighter global risk and aggressive top of book quotes. Fast signals force defense. Uses HJB.

Market Microstructure, High Frequency Trading, Execution, and Limit Order Books (LOB)

Trading and Market Microstructure

Performative Market Making

Charalampos Kleitsikas|Stefanos Leonardos|Carmine Ventre
King's College London

Researchers formalize financial performativity with a performative SDE that steers mid price toward reservation price at a tunable rate. In Avellaneda Stoikov market making, closed forms show inventory driven drift that breaks martingale and speeds alpha decay. They build a performativity aware market maker that separates drift from noise and arbitrages A and S. Simulations show higher PnL, lower variance, and Sharpe gains from theta enhanced variant. Limits include linearization, HJB sub solution, and simulations.

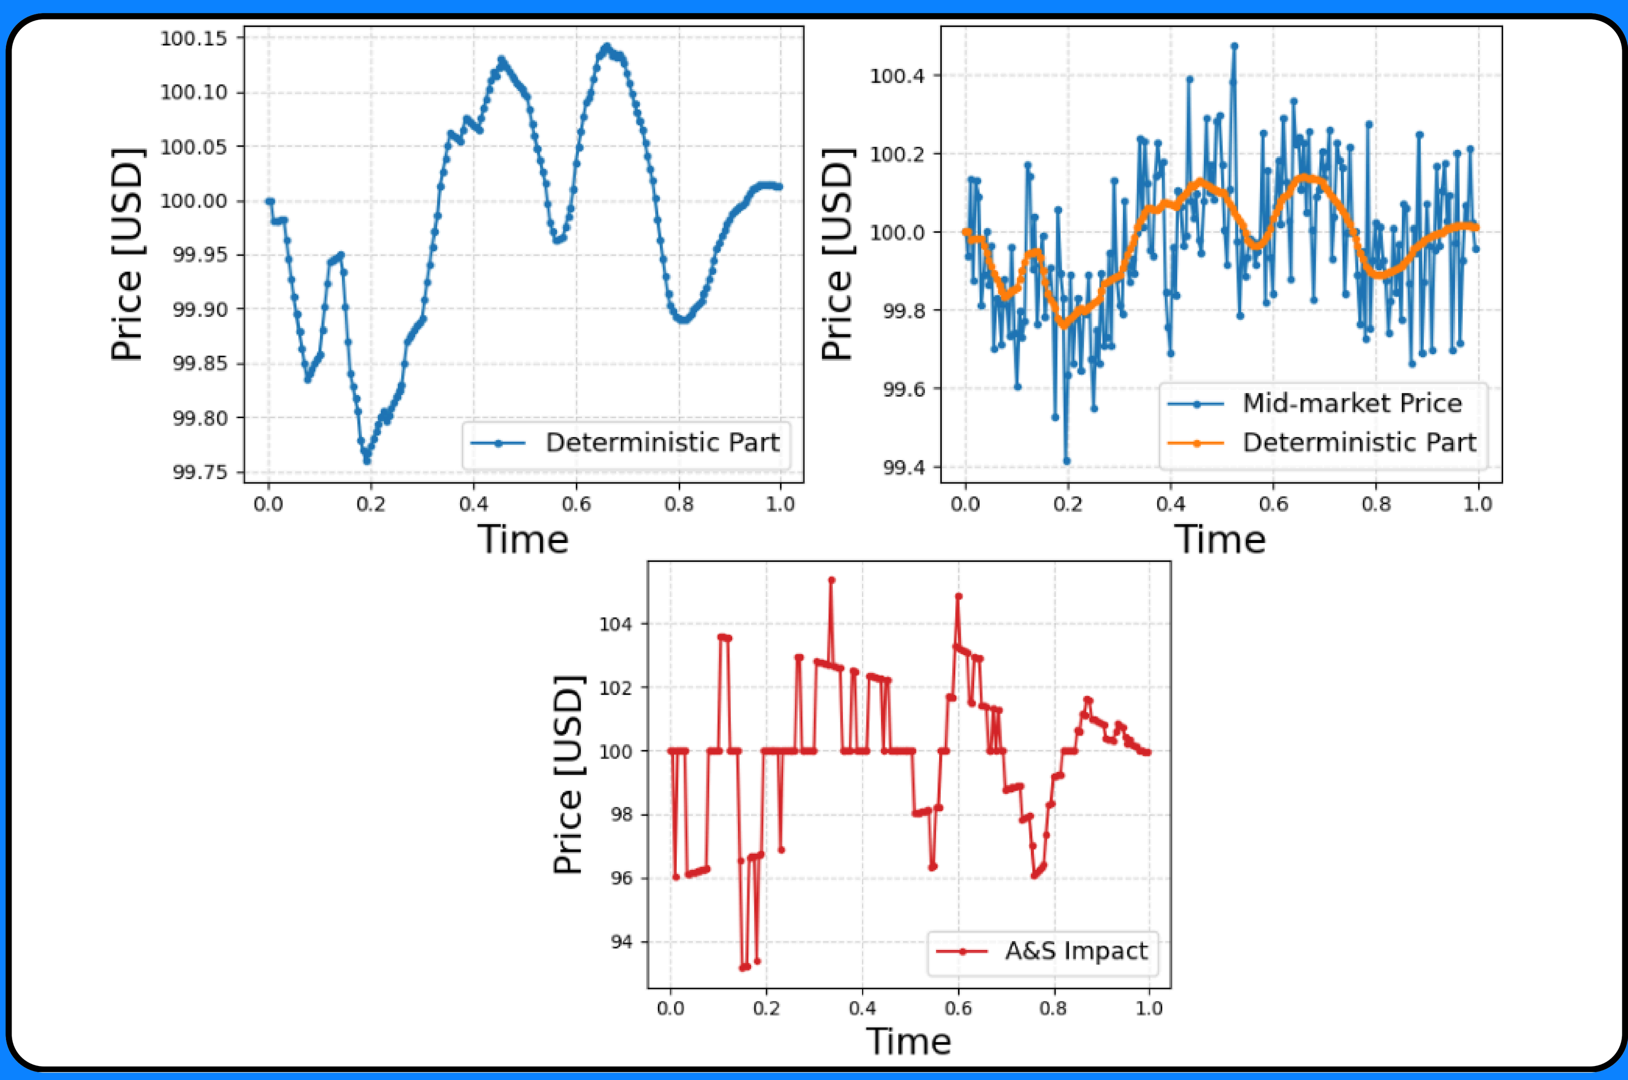

Market Microstructure, High Frequency Trading, Execution, and Limit Order Books (LOB)

Trading and Market Microstructure

ZANISTA

LiT: Limit Order Book Transformer

Yue Xiao, Carmine Ventre, Yuhan Wang, Haochen Li, Yuxi Huan, Buhong Liu
King's College London; Birkbeck, University of London; University College London; SOAS University of London

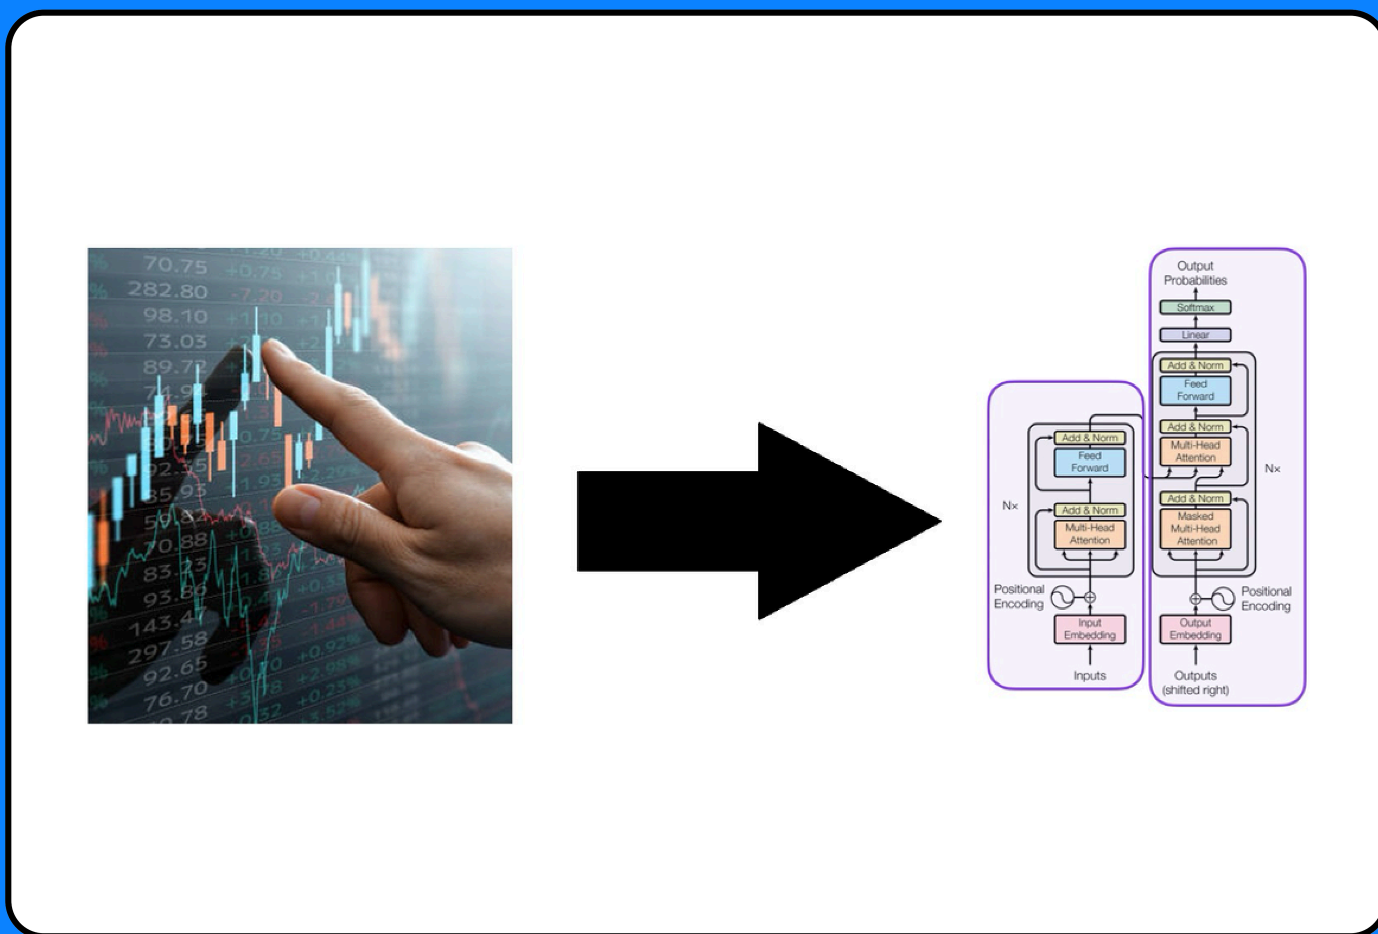

This paper presents Limit Order Book Transformer (LiT), a novel architecture for forecasting short-term market movements using high-frequency limit order book data. Unlike prior methods relying on convolutions, LiT applies structured patches and transformer self-attention to capture spatial and temporal dependencies in market microstructure. Evaluated across multiple datasets and prediction horizons, LiT consistently surpasses traditional machine learning and deep learning baselines. Moreover, it demonstrates robustness under distributional shifts through fine-tuning, highlighting its practicality in fast-paced financial environments.

Market Microstructure, High Frequency Trading, Execution, and Limit Order Books (LOB)

Refined Expansions of the Skew-Stickiness Ratio in Stochastic Volatility Models

F Bourgey|J Delemotte|S De Marco
Bloomberg Ecole Polytechnique 80 Technologies SAS

Authors deliver explicit second order formulas for the Skew Stickiness Ratio in forward variance models for Variance Swap (VS) and at-the-money-forward (ATMF). Using Bergomi Guyon expansion and Fréchet differentiation, they avoid nested Monte Carlo. VS and ATMF SSR coincide at first order but diverge at second. Results recover SSR ≈ 2 for Markovian and SSR ≈ H+3/2 for rough models. Accuracy is good on SPX days but weak in stress and low H. A hybrid fix improves accuracy.

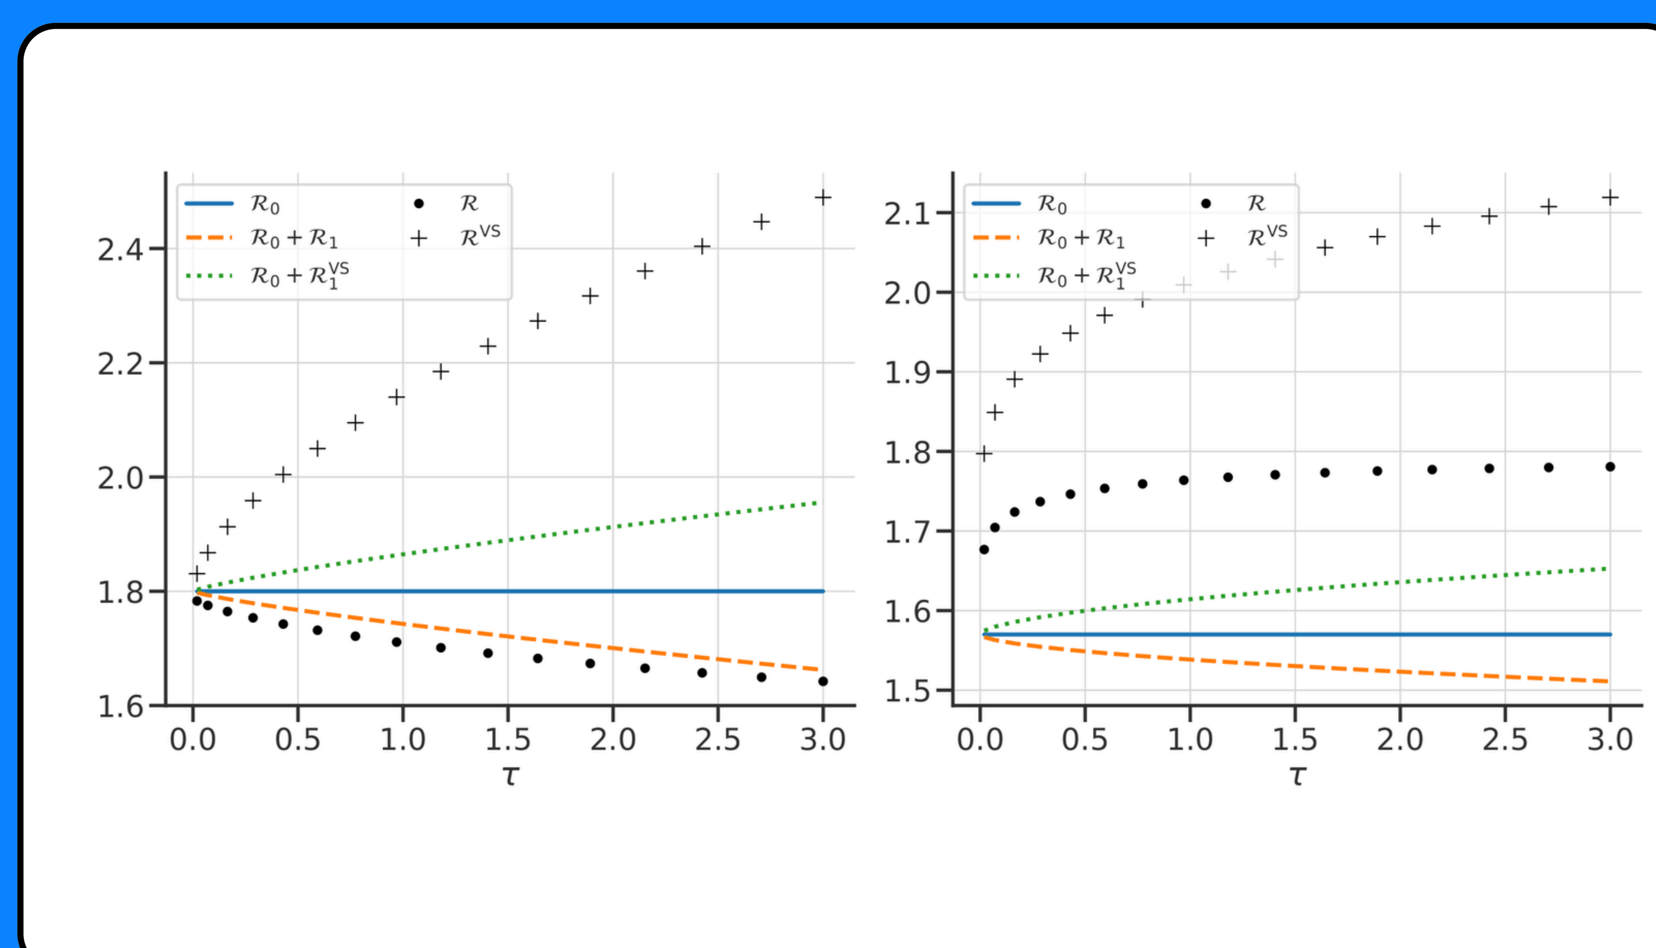

Derivative Modeling, Option trading and Volatility Modeling

FX and options

Option pricing under non-Markovian stochastic volatility models: A deep signature approach

Jingtang Ma|Xianglin Wu|Wenyuan Li
Southwestern University of Finance and Economics The University of Hong Kong

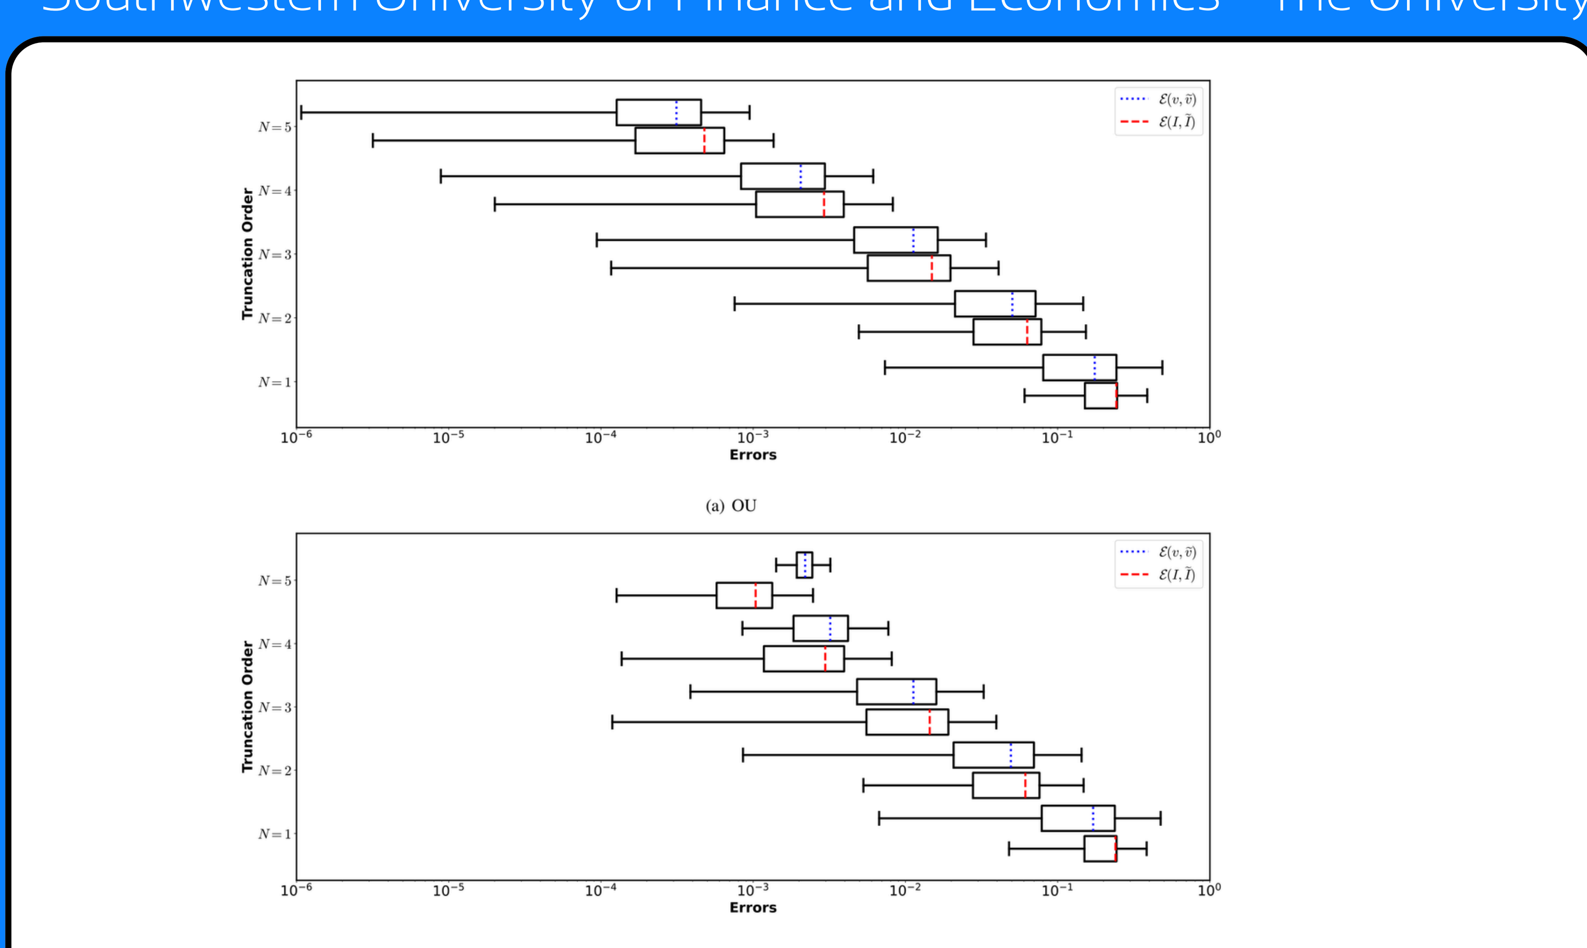

Researchers present an option pricing method for path dependent volatility by turning a non Markovian model into a Markovian one. They lift dynamics to a rough SDE and expand the rough path into Brownian signatures, yielding SDE and PDE via Feynman Kac. Linear and nonlinear signature models are tested; the nonlinear model is more accurate. Errors decay like

$$\frac{C^{N+1}}{(N+1)!}$$

and small N near 3 works. Tests on OU, mGBM, rough Heston, rough Bergomi confirm accuracy.

Derivative Modeling, Option trading and Volatility Modeling

Machine Learning for Option Pricing

ZANISTA

Fast reliable pricing and calibration of the rough Heston model

Svetlana Boyarchenko|Marco de Innocentis|Sergei Levendorskii
University of Texas at Austin Deutsche Bank Calico Science Consulting

Researchers unveil a pricer for rough Heston using a BL-modified fractional Adams solver, SINH Fourier inversion, and a Conformal Bootstrap check. Fourier pricers can produce ghost calibrations, like flat smiles with El Euch Rosenbaum parameters. The SINH CB scheme hits 1e-9 accuracy with 20 to 60 terms, millisecond grids at 1e-4 error, and stability near expiry and OTM. Benchmarks flag Laguerre sensitivity and slow BL2. A TSLA case shows fits and out of sample, limits.

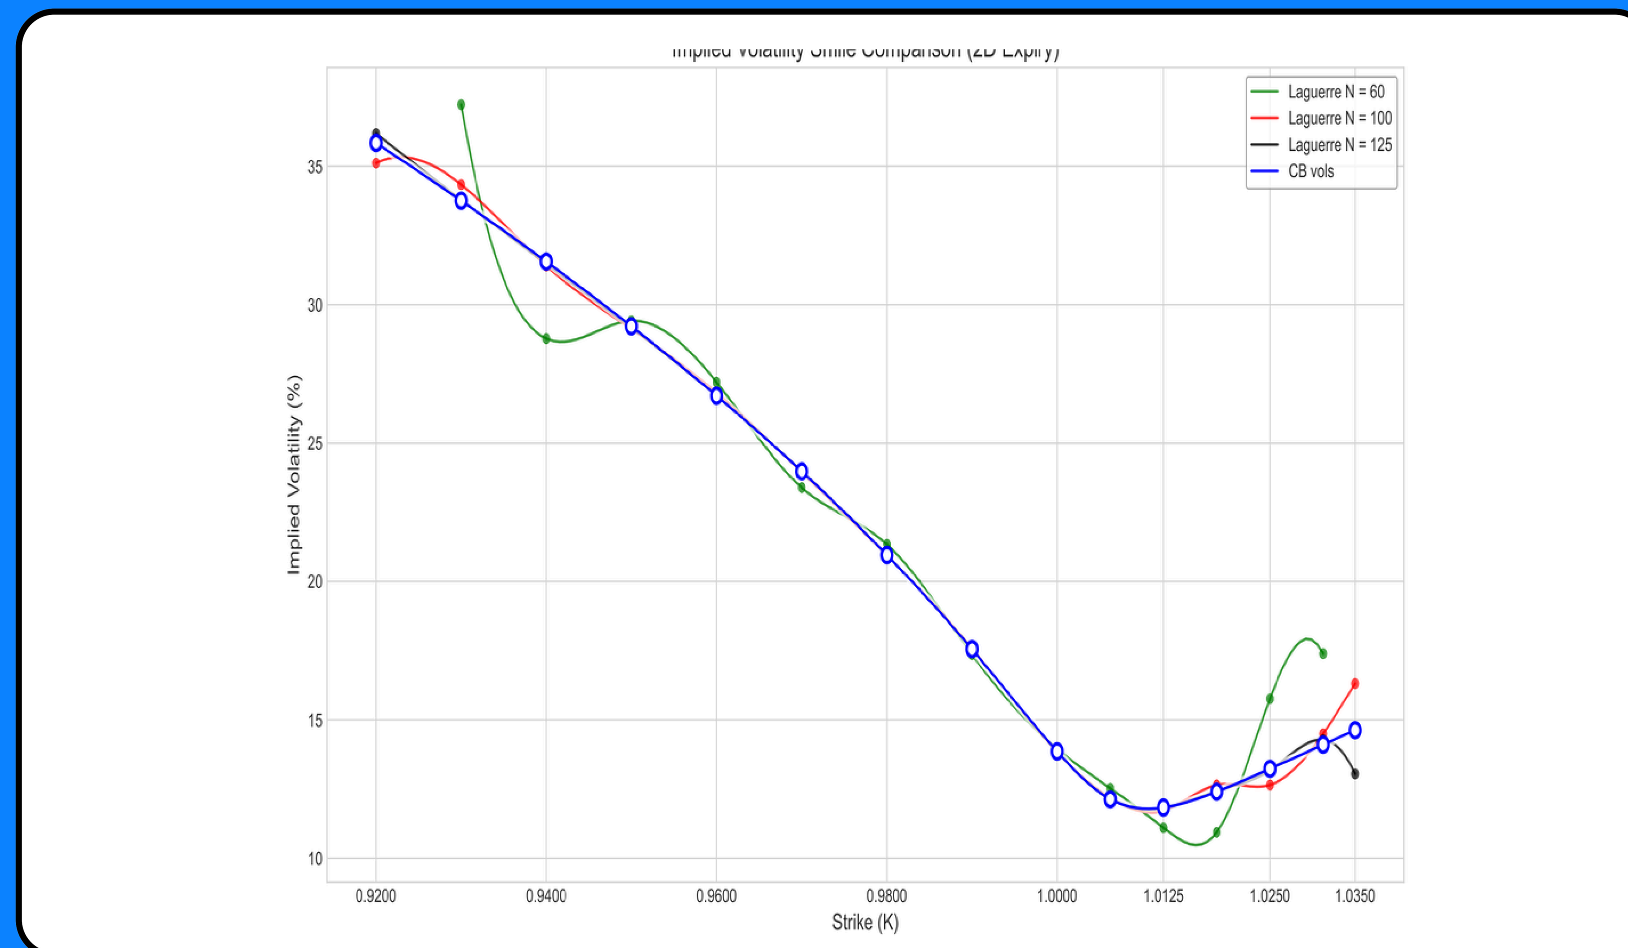

Derivative Modeling, Option trading and Volatility Modeling

FX and options

Structural Equity Option Pricing: Implications for Credit Risk

Nicola Fusari|Sujan Lamichhane
Johns Hopkins University International Monetary Fund

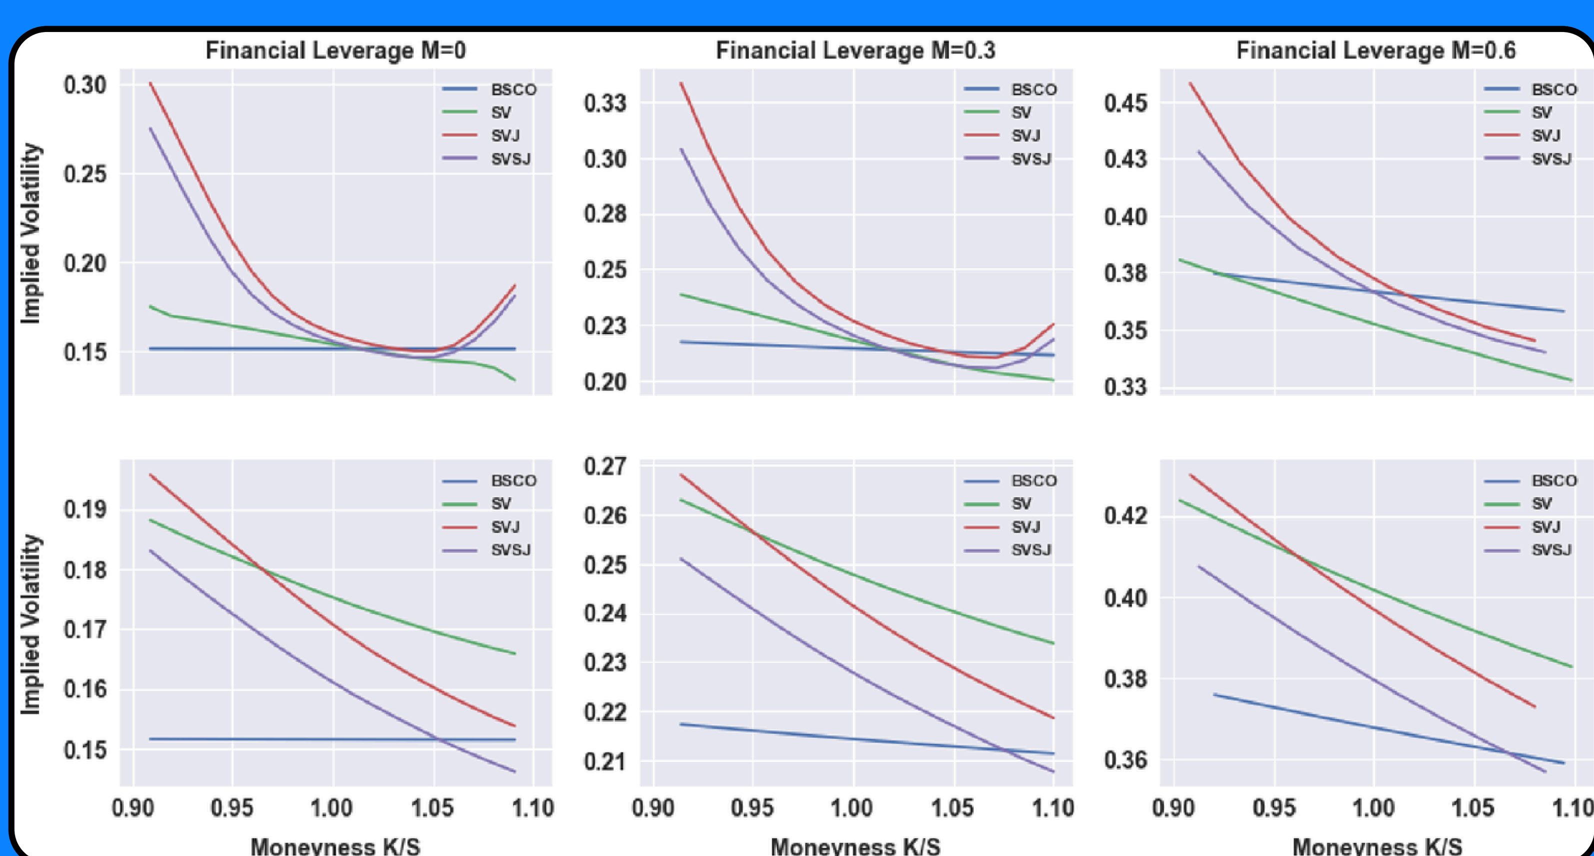

A new structural option model uses deep neural networks to price compound equity options in 0.01 seconds, mapping option data to asset value, variance, leverage, and default risk. Stochastic volatility is pivotal. Jumps matter at low leverage like IBM, while AAL's high leverage plus SV explains skew. Option-implied PDs match CDS for AAL but not IBM. An intermediary constraint index explains over 30 percent of that gap. The approach delivers real-time, equity-based PD and leverage.

Derivative Modeling, Option trading and Volatility Modeling

Computational Finance

Distributional Adversarial Attacks and Training in Deep Hedging

Guangyi He|Tobias Sutter|Lukas Gonon
Imperial College London University of St. Gallen

Deep hedging models break under tiny distribution shifts. The authors fuse distributional adversarial training with deep hedging over a Wasserstein ball, proposing WPGD and WBPGD attacks and a sensitivity-based worst case shift from optimal transport. They coin distributional FGSM and PGD for time series and extend to price and variance via weighted norms. Large experiments and S&P 500 tests show attacks raise losses. Adversarial training improves performance, notably halving loss in Heston with 5,000 paths.

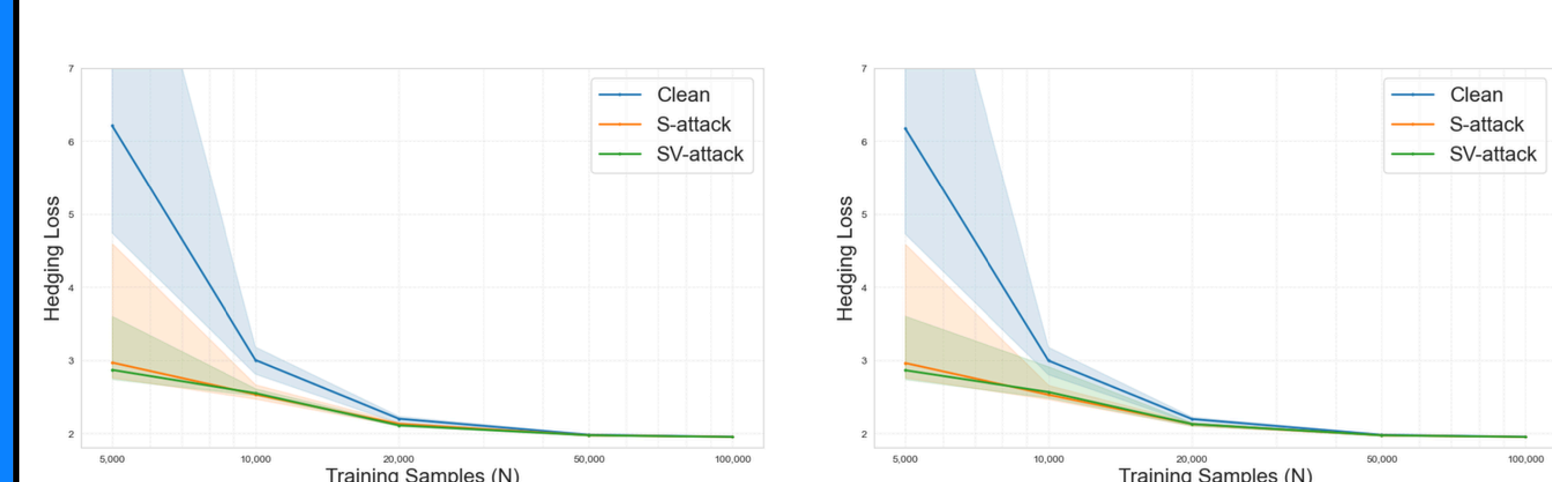

(a) Out-of-sample performance (b) Out-of-distribution performance

Derivative Modeling, Option trading and Volatility Modeling

Deep Reinforcement Learning in Trading Strategies

ZANISTA

Marketron Through the Looking Glass: From Equity Dynamics to Option Pricing in Incomplete Markets

I Halperin|A Itkin
Fidelity Investments New York University

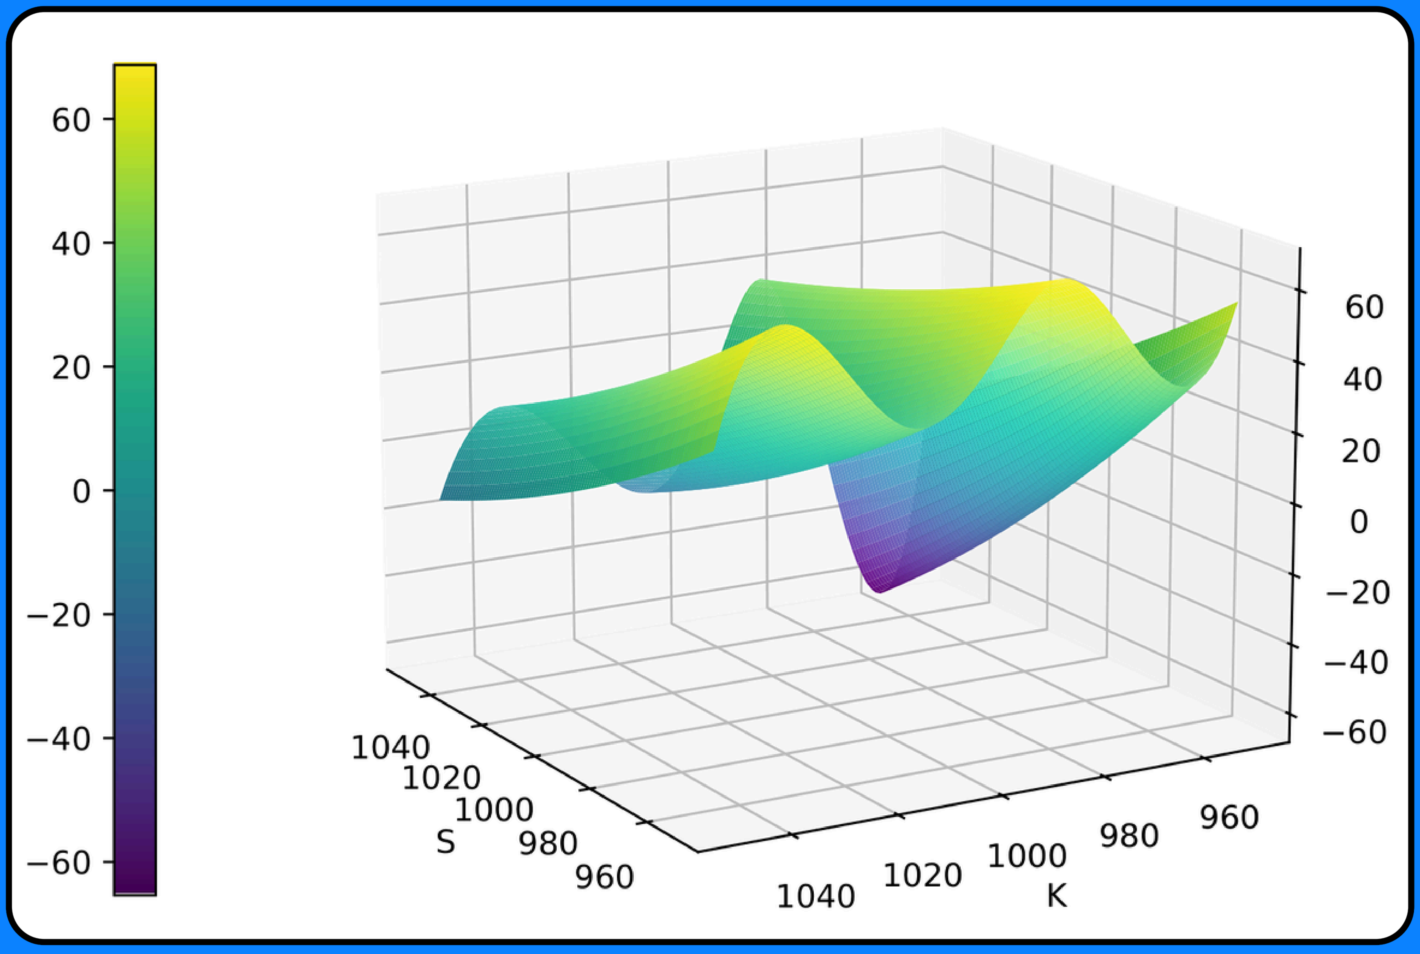

The paper extends the Marketron model to options in incomplete markets with a utility pricer. Nonlinear HJB becomes a Volterra integral via a Duhamel method and is solved with Gaussian RBFs, or by operator splitting and Cole Hopf. SPX calibration hits 3 to 8 percent and prices 49 options in 0.7 seconds. The market price of risk is closed form and state dependent. Results show Black Scholes departures, rough paths, and mixed multi maturity fits.

Derivative Modeling, Option trading and Volatility Modeling | Computational Finance

The Discrete-Time 4-Factor Path-Dependent Volatility Model: Calibration under P and Q

Julien Guyon|L?o Parent
Institut Polytechnique de Paris Inria New York University

Authors present a discrete time path dependent volatility model with fat tailed shocks that fits SPX and VIX. Under Q it matches VIX with r squared 0.96 in sample and 0.87 out of sample and captures V shaped SPX smiles plus a wing. Under P it beats benchmarks and a three factor wins out of sample. P and Q align. A mixed loss stabilizes memory by 1.5 to 3 times, leaving smile and futures gaps.

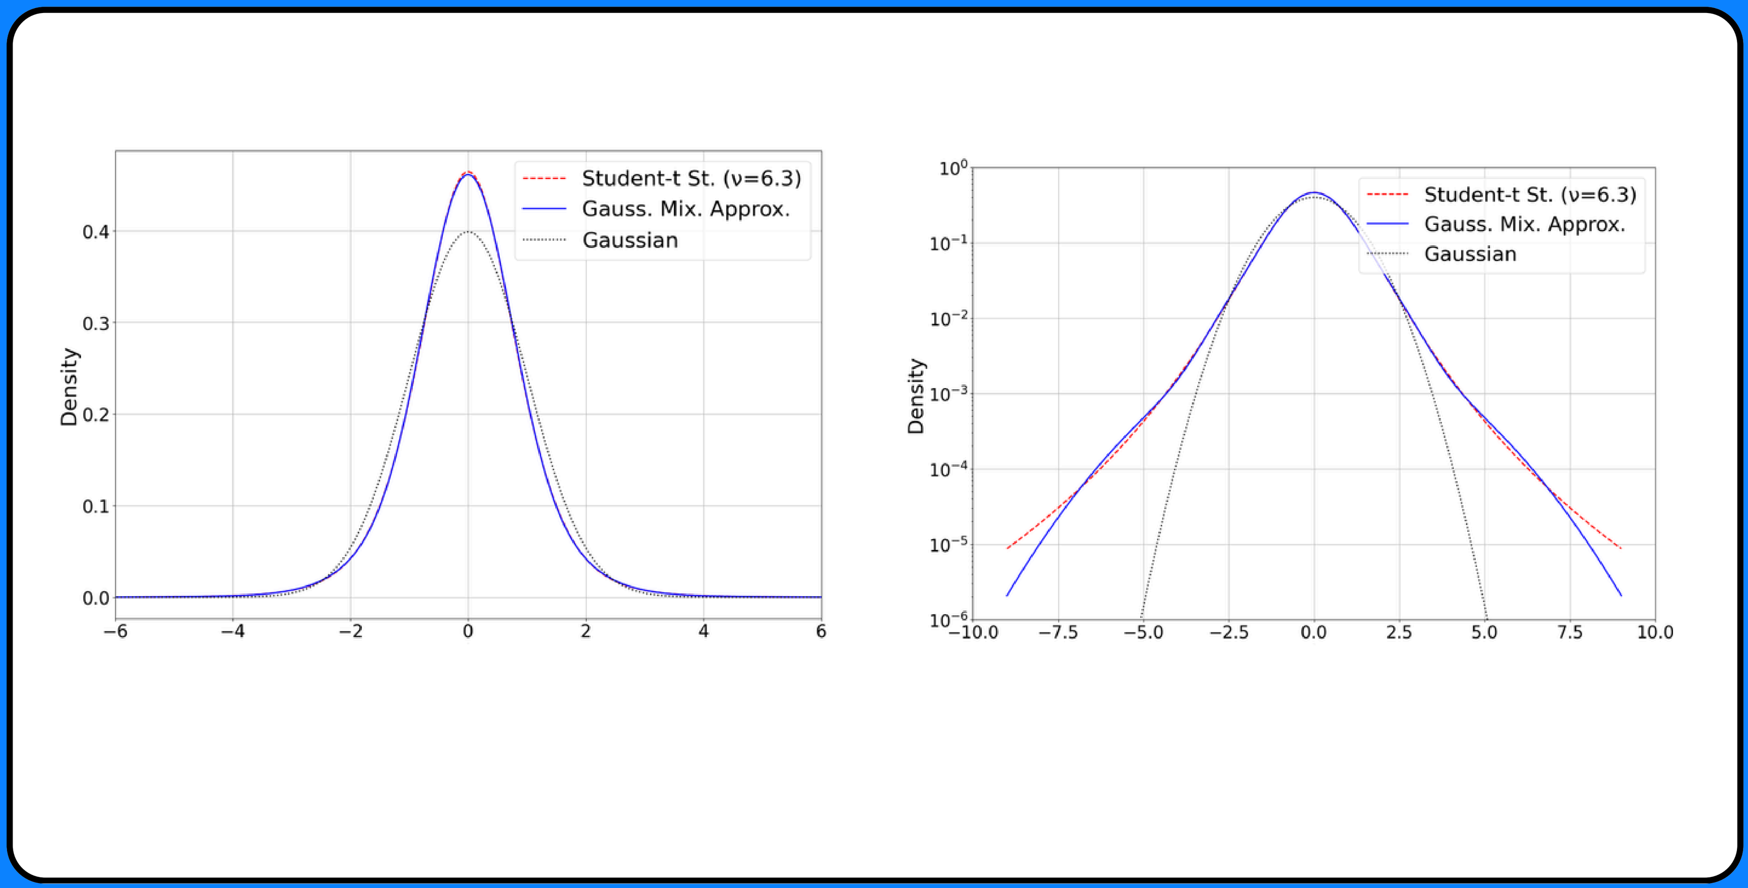

Derivative Modeling, Option trading and Volatility Modeling | Quantitative Finance

Ambiguity-Averse Deep Hedging with Feature Clustering

Adam C. Jones|Blanka Horvath|Christoph Reisinger|Ben Wood|Lianjun Bai|Amira Akkari
University of Oxford J.P. Morgan

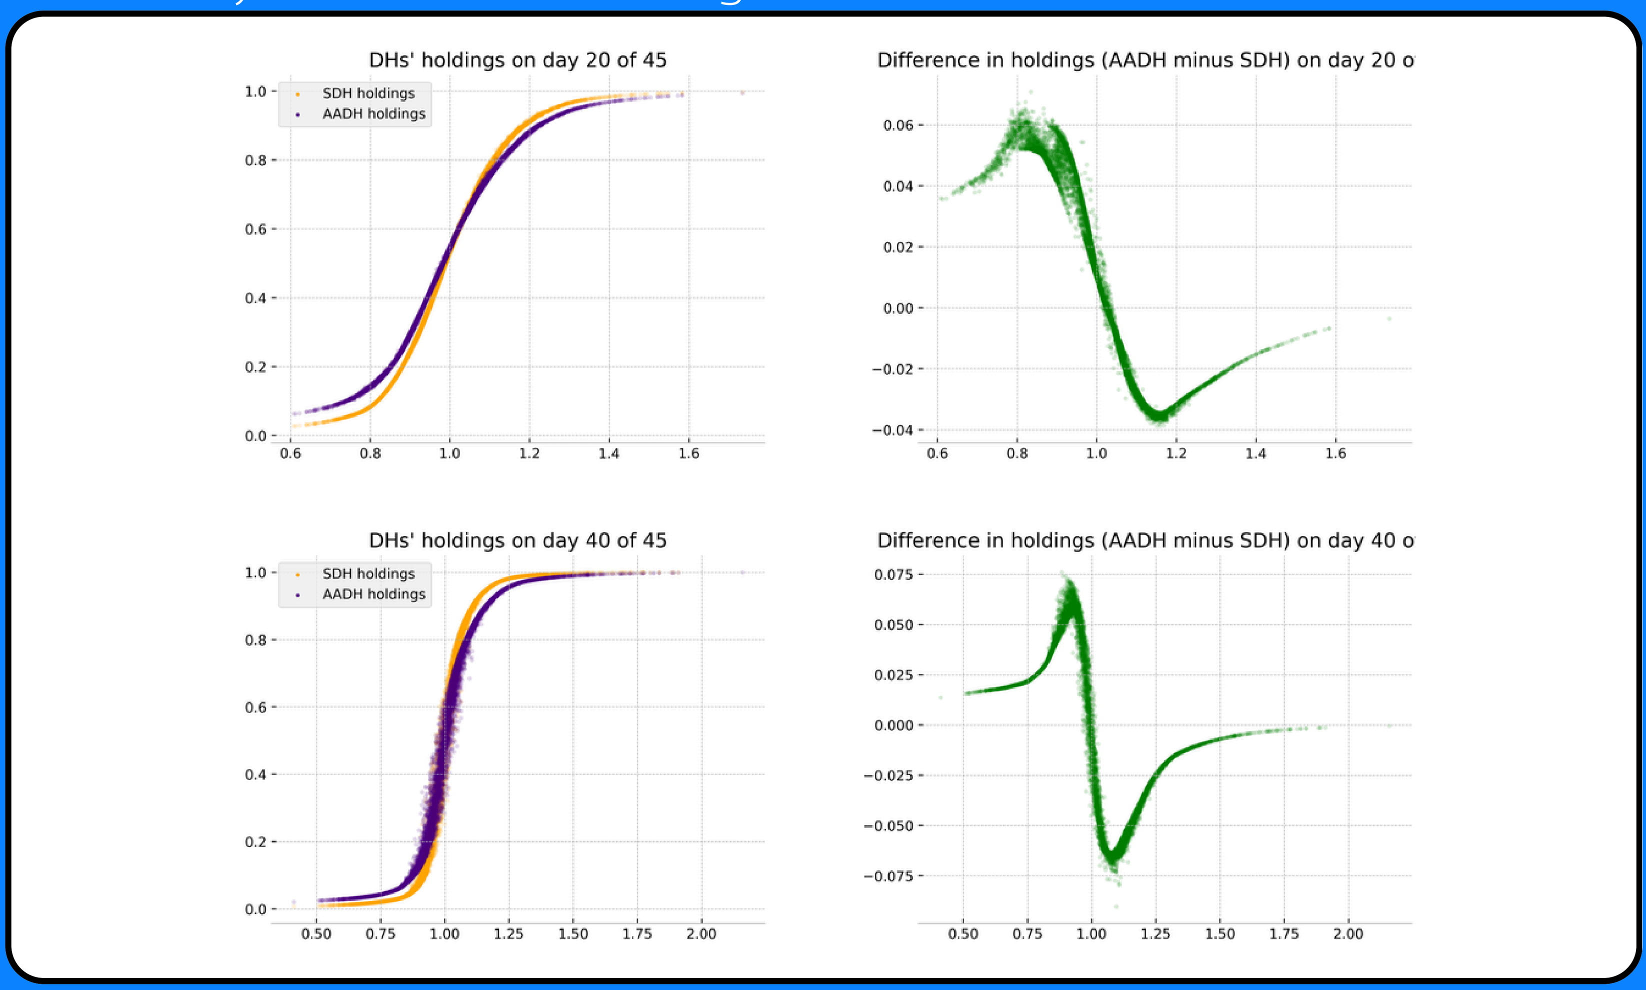

The paper proposes Ambiguity Averse Deep Hedger, which trains hedging policies to withstand distribution shifts by clustering paths on features like volatility, autocorrelation, and drawdowns, then aggregating cluster risks with an entropic ambiguity measure. Without altering generators or charging premiums, it cuts risk by 54 percent. It beats deep hedging on out of distribution tests across forward-starts, vanilla calls, and COVID shock, but lags in distribution. Robustness depends on feature choice, clusters, and ambiguity level.

Derivative Modeling, Option trading and Volatility Modeling | Machine Learning for Option Pricing

ZANISTA

Price Pressures from Daily Mutual Fund Flows: Does it pay to go with the flow?

Matti Suominen|Eeli Tuovinen

Aalto University

Google Scholar SSRN

Daily mutual fund flows move large cap U.S. stocks. Both heavy inflow and heavy outflow pressures precede positive 20 day abnormal returns, with impact 1 to 5 days pre report from early trading. Using Morningstar flows and CRSP data, researchers find reversals after outflows and continuation after inflows. Two 20 day strategies earn up to 7.4 percent alpha. Aggregate flows predict 20 day excess returns with 3.9 percent R squared. Coverage, assumptions, frictions temper results.

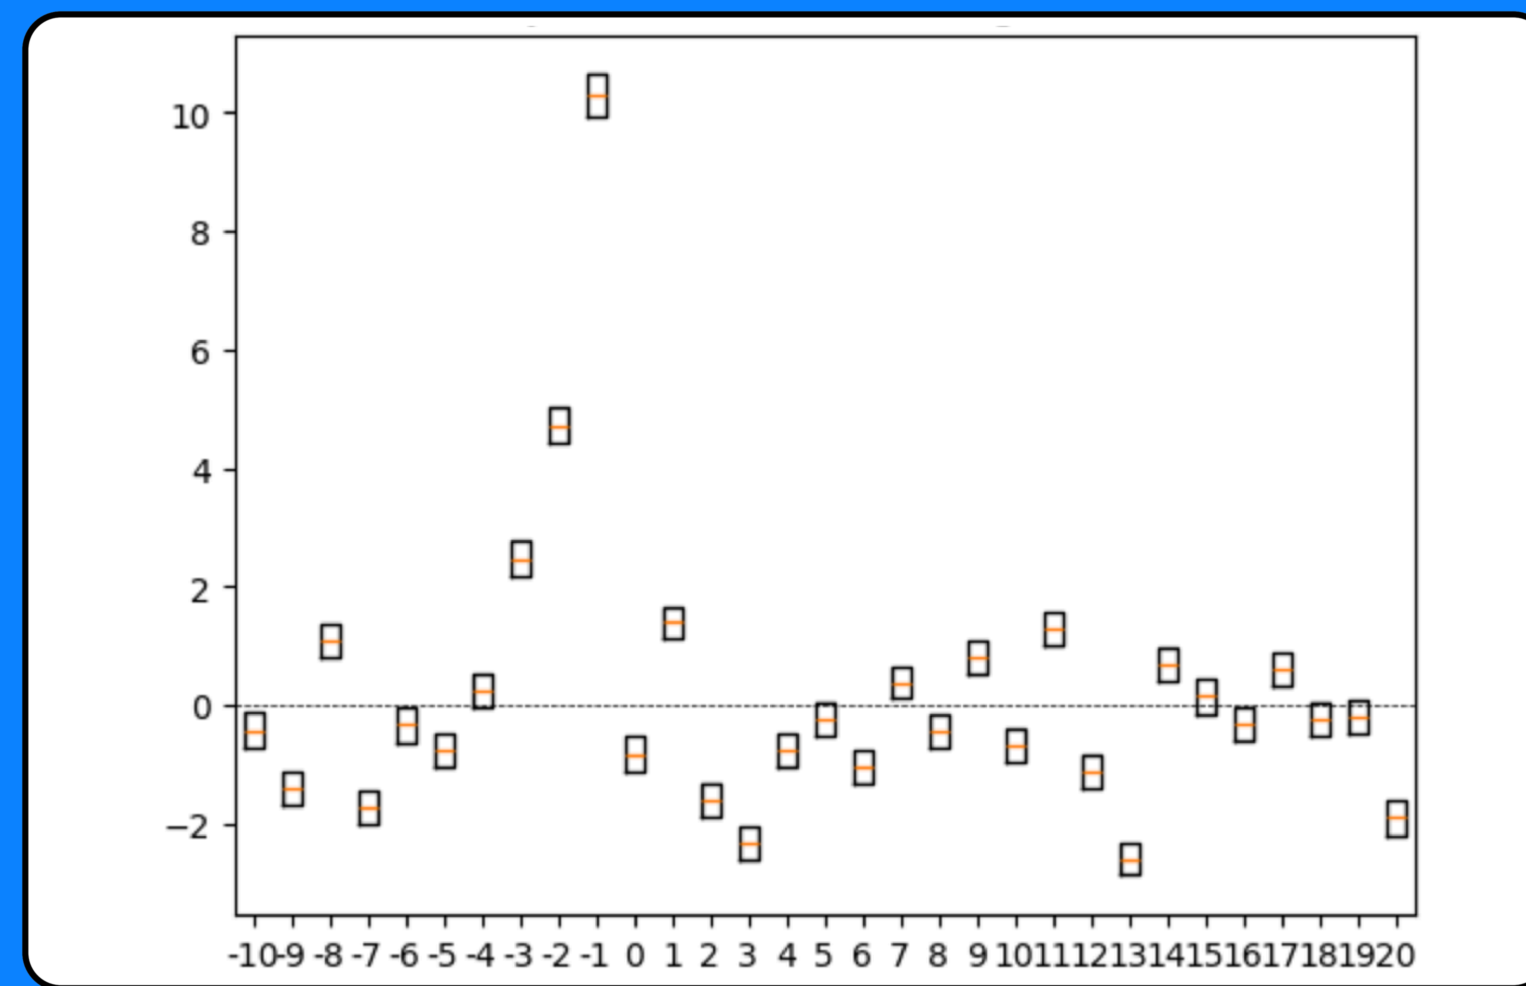

Investments and mid to low frequency equity and ETF trading strategies

Investments and mid to low frequency equity and ETF trading strategies

Cointegration-based pairs trading: identifying and exploiting similar exchange-traded funds

Google Scholar Springer

Kezhong Chen|Constantinos Alexiou

Cranfield University

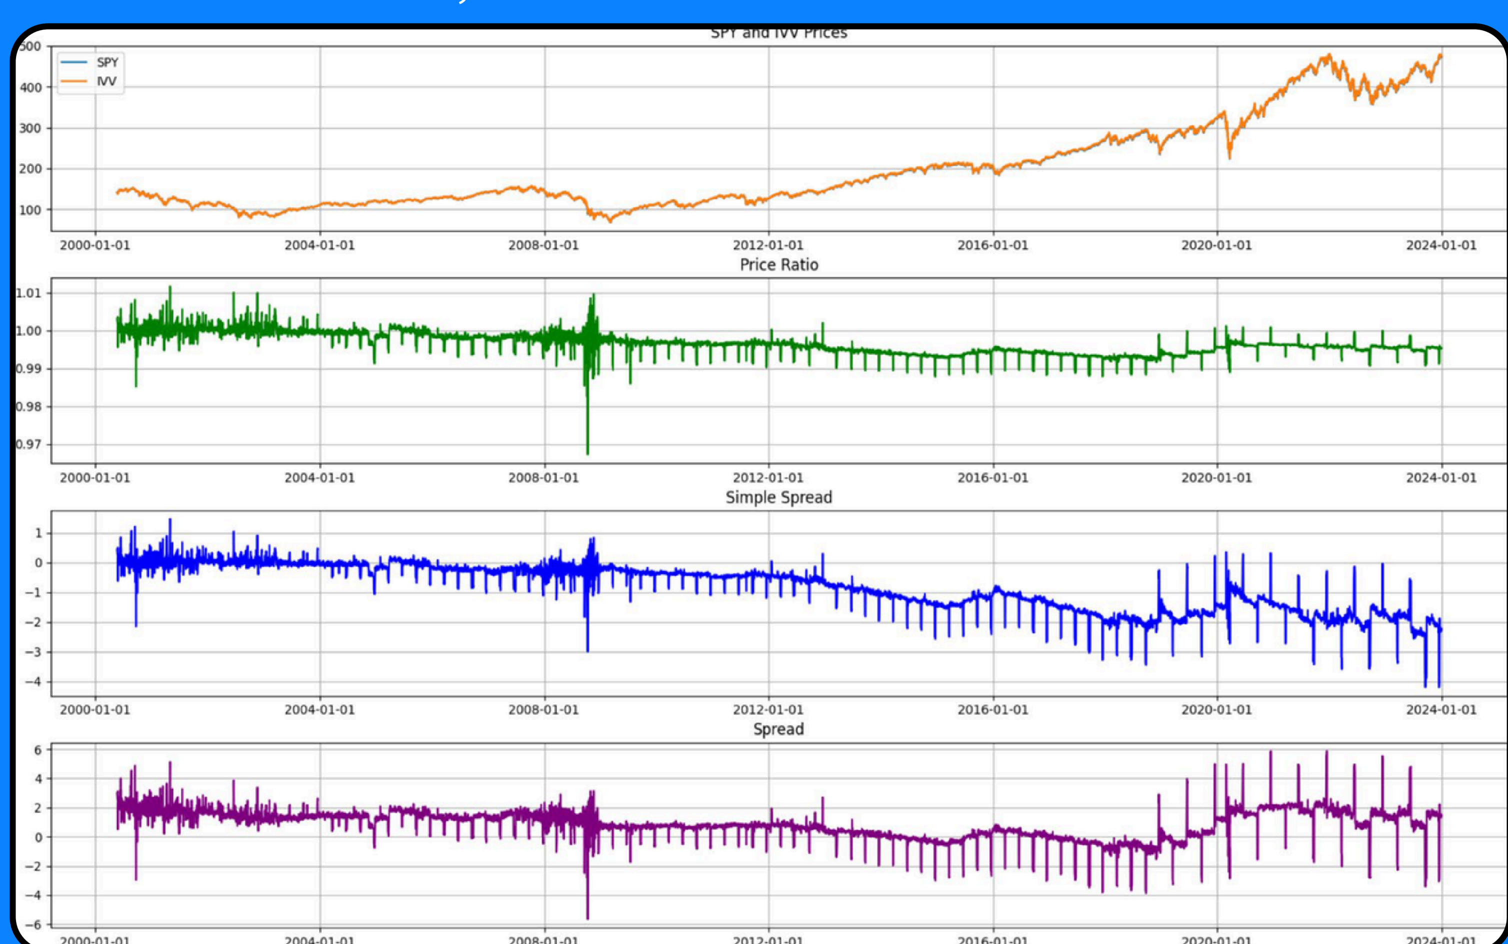

Chen and Alexiou find ETF pairs trading based on cointegration works in brief stable spells, then weakens. Using 30 pairs from 2000 to 2024, ADF screens show lower entry thresholds raise trades, profits and Sharpe, but also volatility and drawdowns. SPY with IVV excels. QQQ with XLK disappoints. A regime engine selects ADF leaders and uses VIX to switch to momentum, yielding negative correlation, Sharpe 0.81, alpha 12.8, and 1.22 in 50 50 SPY blend.

Investments and mid to low frequency equity and ETF trading strategies

Investments and mid to low frequency equity and ETF trading strategies

Levered Single-Stock ETFs

Hendrik Bessembinder
WP Carey School of Business, Arizona State University

Google Scholar SSRN

This paper analyzes levered single-stock ETFs (LSS-ETFs), comparing their returns to a simple leverage-times-stock-return benchmark. It finds that both hypothetical and actual LSS-ETFs underperform this benchmark, with losses attributed to daily rebalancing costs and frictions (fees, trading costs). The study also highlights cost asymmetries between long and inverse funds and examines the likelihood of extreme negative returns.

Table 3: Actual and Frictionless Returns to Levered Single Stock ETFs, July 2022 to June 2025

Panel A: Weekly Outcomes

		ETF Outcomes								Underlying Stock Outcomes	
Num Obs	Leverage	Mean Benchmark Return: With Weekly Rebalance (Expression 4)	Mean Benchmark Return With Daily Rebalance (Expression 5)	Mean Actual ETF Return	Median Actual ETF Return	Robust Skewness, Actual ETF Return	Mean Rebalance Cost	Mean Frictional Cost	Mean Total Cost vs. Exp (4)	Mean Stock Return	Std of Daily Returns
4050	All	0.348%	0.301%	0.198%	-0.145%	0.114	0.047%	0.102%	0.149%	0.741%	2.546%
2443	Long	1.193%	1.179%	1.054%	0.481%	0.194	0.014%	0.125%	0.140%	0.700%	2.472%
1607	Short	-0.938%	-1.035%	-1.102%	-0.930%	0.016	0.097%	0.067%	0.164%	0.804%	2.660%

Panel B: Monthly Outcomes

		ETF Outcomes								Stock Returns	
Num Obs	Leverage	Mean Benchmark Return: With Monthly Rebalance (Expression 4)	Mean Benchmark Return With Daily Rebalance (Expression 5)	Mean Actual ETF Return	Median Actual ETF Return	Robust Skewness, Actual ETF Return	Mean Rebalance Cost	Mean Frictional Cost	Mean Total Cost vs. Exp (4)	Mean Stock Return	Std of Daily Returns
960	All	1.446%	0.999%	0.569%	-1.171%	0.261	0.448%	0.429%	0.877%	3.057%	2.752%
579	Long	4.889%	4.629%	4.096%	2.246%	0.244	0.259%	0.533%	0.793%	2.894%	2.680%
381	Short	-3.784%	-4.518%	-4.790%	-3.968%	0.232	0.734%	0.272%	1.005%	3.305%	2.862%

Investments and mid to low frequency equity and ETF trading strategies

Investments and mid to low frequency equity and ETF trading strategies

ZANISTA

Alpha In Indian Mutual Fund Returns

Raj Mehta|Azain Jaffer

Morgan Stanley Bridgewater Associates Yale University

Google Scholar SSRN

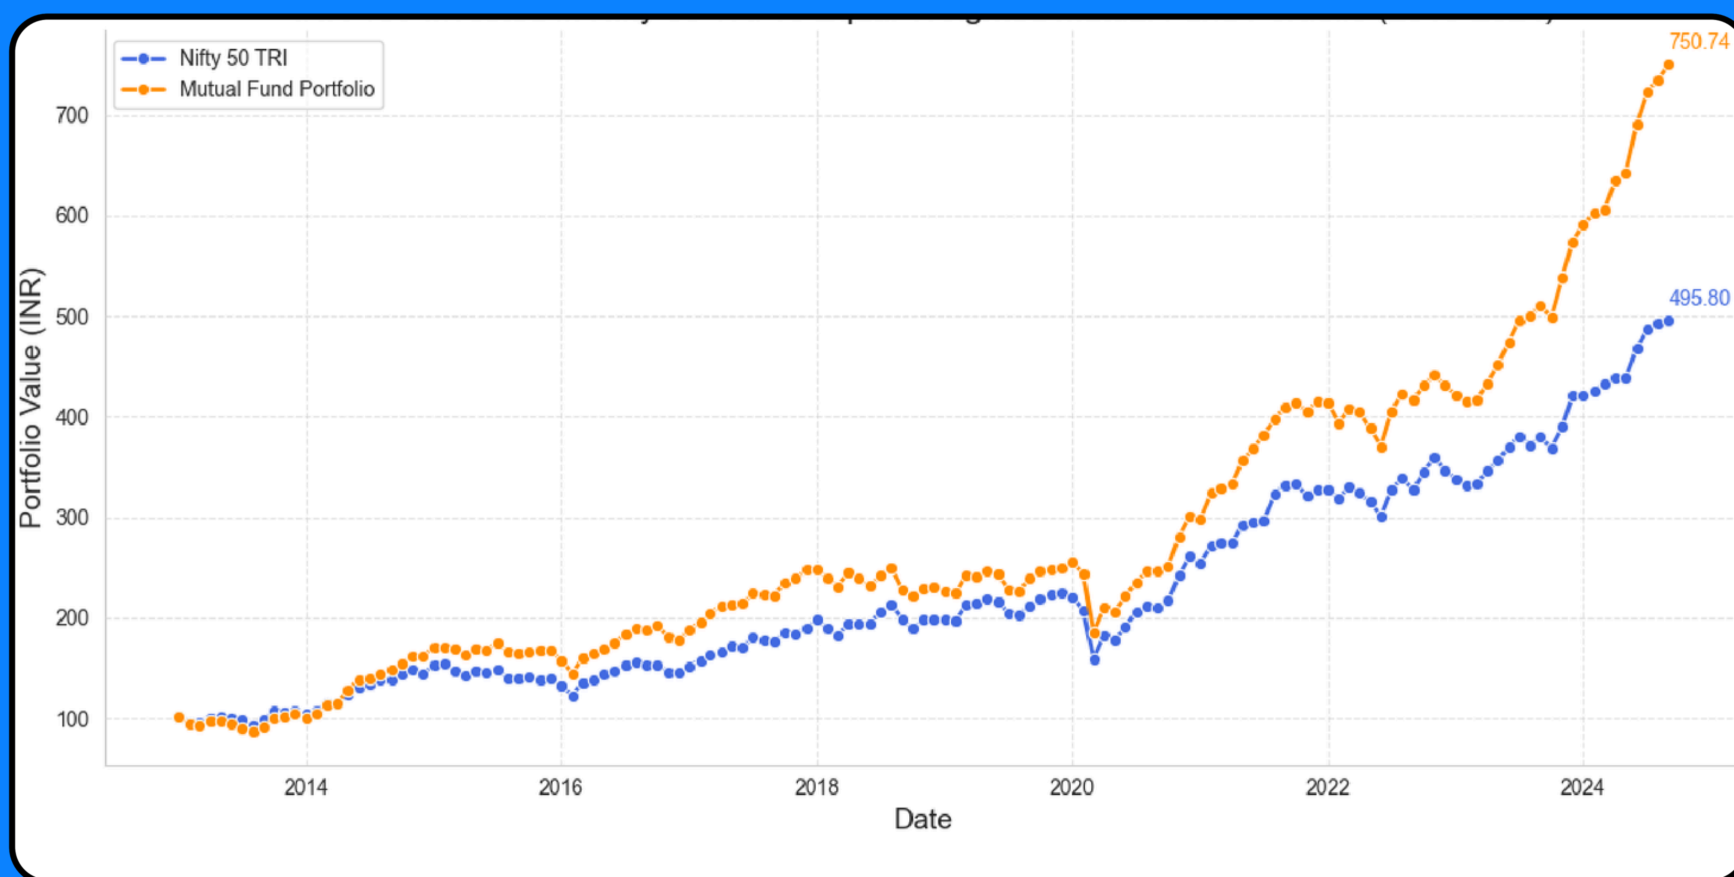

A study of 425 Indian equity funds from 2013 to 2024 finds no net alpha after fees and factors. Even large funds fail to beat factor benchmarks. Using Fama French, Carhart, and a profitability factor EBME with bootstrap tests across AUM, it shows managers load on profitability but shun value and momentum. The top 7 to 15 percent show alpha, which fees erase. HML is weak in India. Bigger funds dodge blowups without delivering alpha.

Investments and mid to low frequency equity and ETF trading strategies

Investments and mid to low frequency equity and ETF trading strategies

Robust Econometrics for Growth-at-Risk

Tobias Adrian|Yuya Sasaki|Yulong Wang

International Monetary Fund Vanderbilt University Syracuse University

Google Scholar arXiv

The paper challenges a core Growth at Risk assumption: constant tail heaviness. The standard two step method fixes the Pareto exponent, making expected shortfall and long rise constant multiples of the quantile. A new approach lets the tail index vary, with $v(x) = \exp(x'\beta)$ tail index regression and data driven thresholds. Simulations and 1893 to 2016 U.S. data show tighter, less biased estimates, near nominal 5 percent coverage, and crises: old -10.2 percent versus new -4.6 percent.

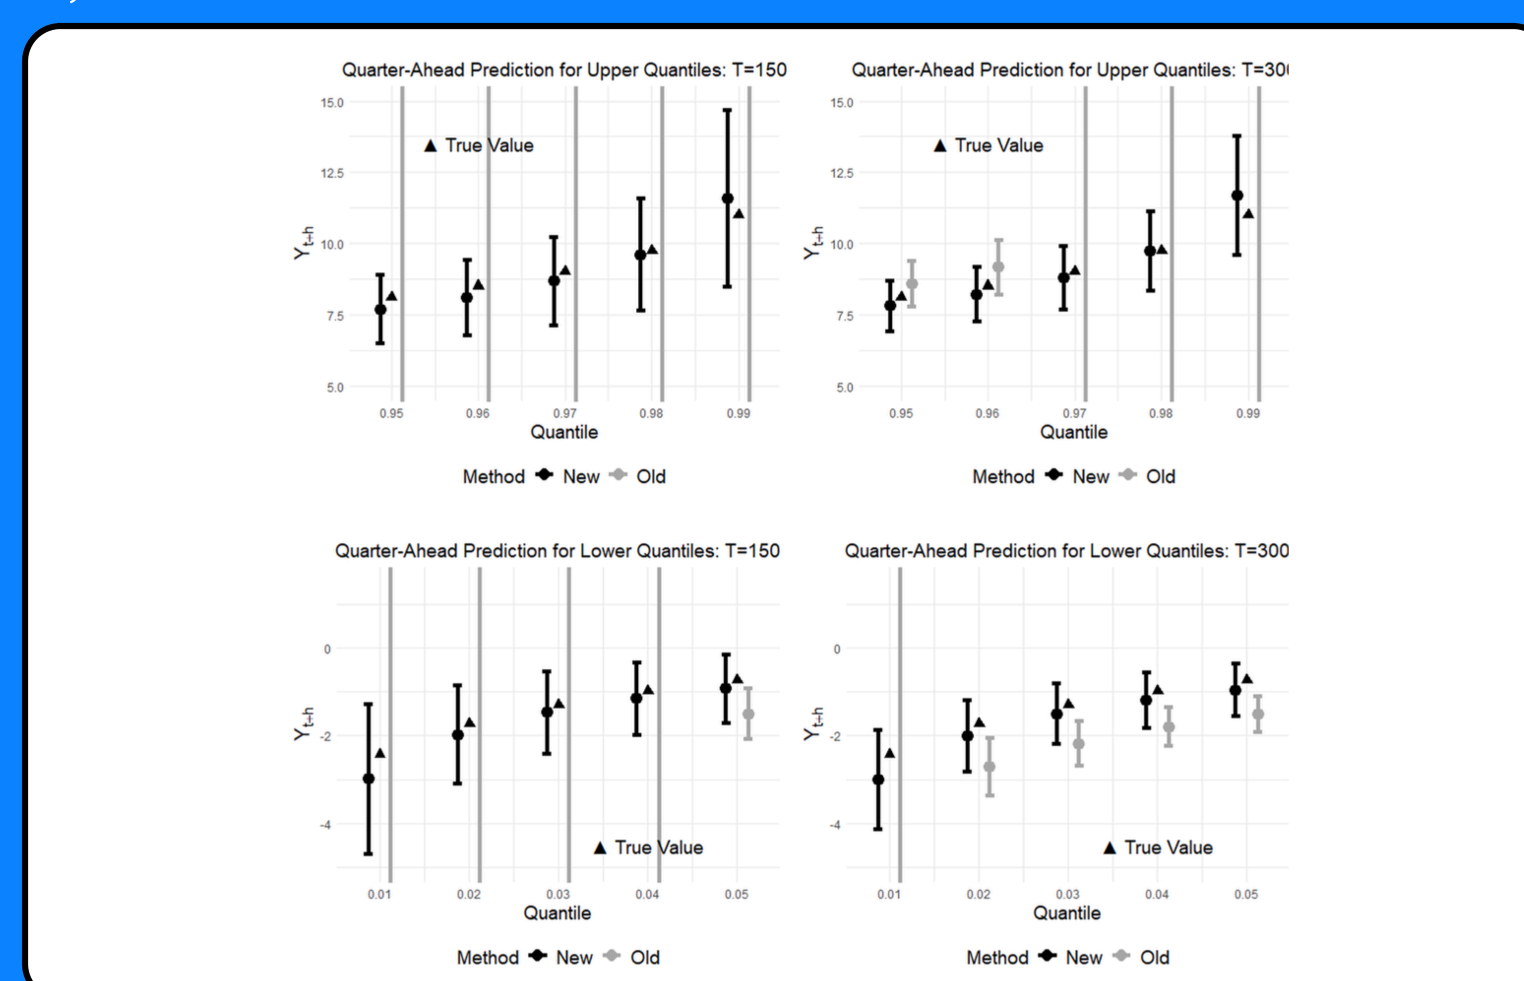

Economics and Macroeconomics

econometric

Combining a Large Pool of Forecasts of Value-at-Risk and Expected Shortfall

JW Taylor|C Wang

University of Oxford The University of Sydney

Google Scholar arXiv

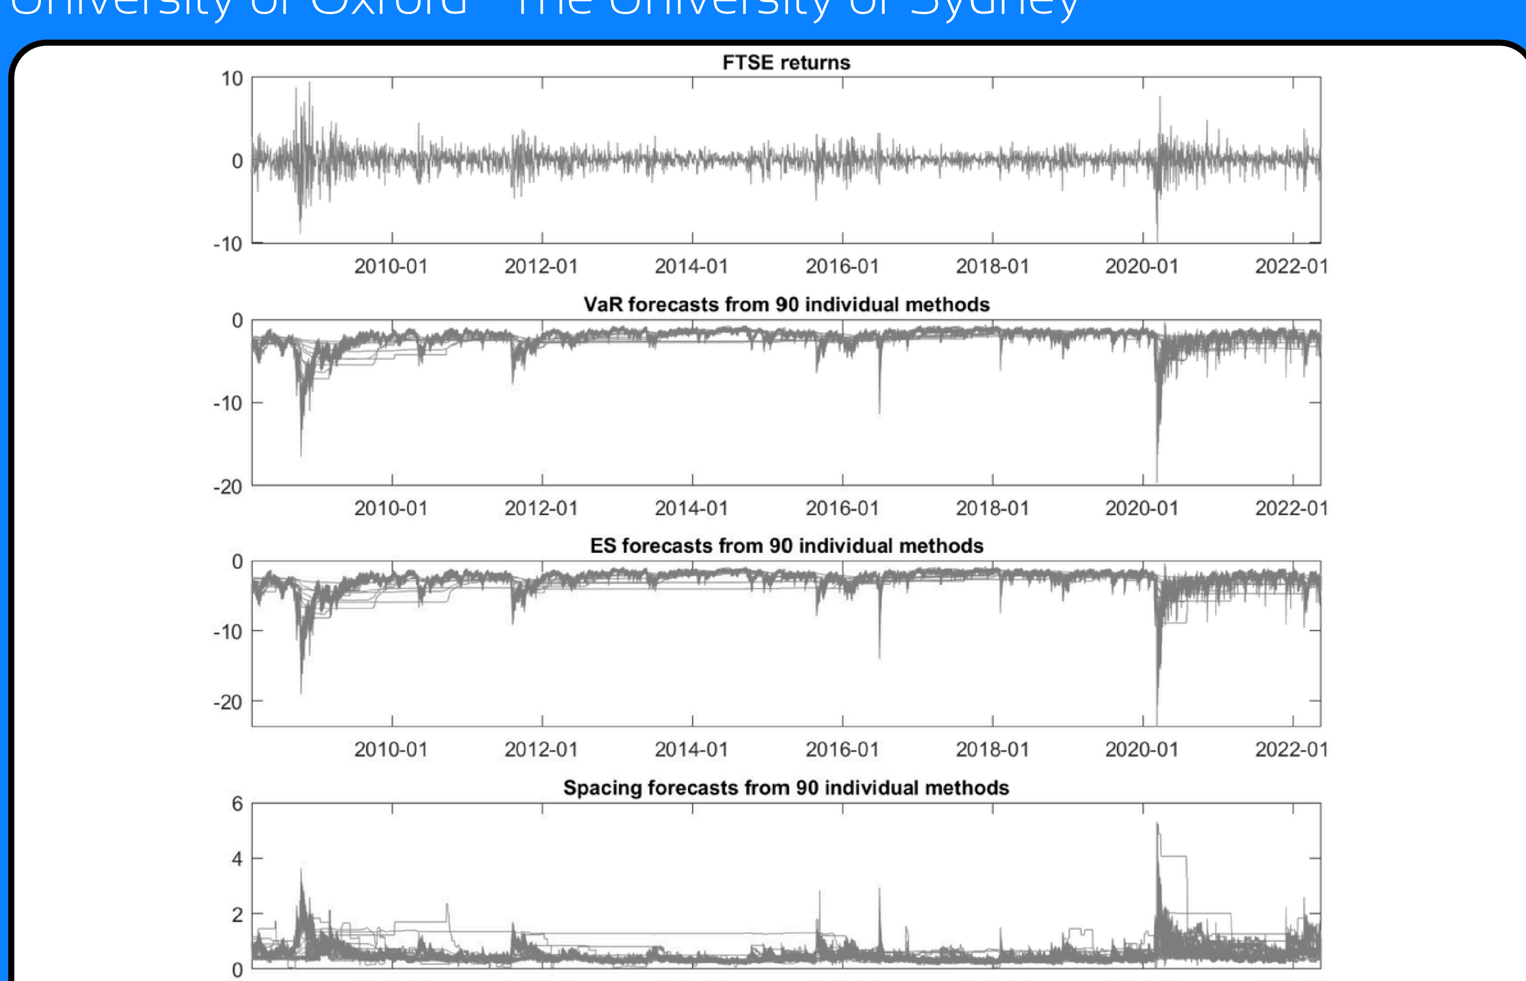

Combining risk forecasts beats choosing one model. Using 90 VaR and ES models on six indexes, the study finds three winners: a higher trimmed mean, probability averaging using skew t fits from VaR and ES pairs, and performance based weights with regularization. Guided by Fissler Ziegel scoring, combinations push tails to correct mild risk and aid stability. Averaging combiners works. With shorter windows, probability averaging and relative score weights shine. Limits include 2.5 percent focus.

Quantitative Risk Management

Quantitative Finance

ZANISTA

CTBench: Cryptocurrency Time Series Generation Benchmark

Yihao Ang|Qiang Wang|Qiang Huang|Yifan Bao|Xinyu Xi|Anthony K. H. Tung|Chen Jin|Zhiyong Huang

National University of Singapore Harbin Institute of Technology

Researchers release CTBench, a benchmark for crypto time series generation using a 452-token dataset from 2020 to 2024. It tests models on Predictive Utility and Statistical Arbitrage across 13 metrics. No model dominates. Diffusion models minimize forecast error yet trail in profits. COSCI-GAN excels in bull markets. TimeVAE shines in stable regimes. Fourier-Flow is consistent. LS4 and KoVAE lead stat arb. FIDE is safest but unprofitable. Fees of 0.03 percent cut Sharpe and punish turnover.

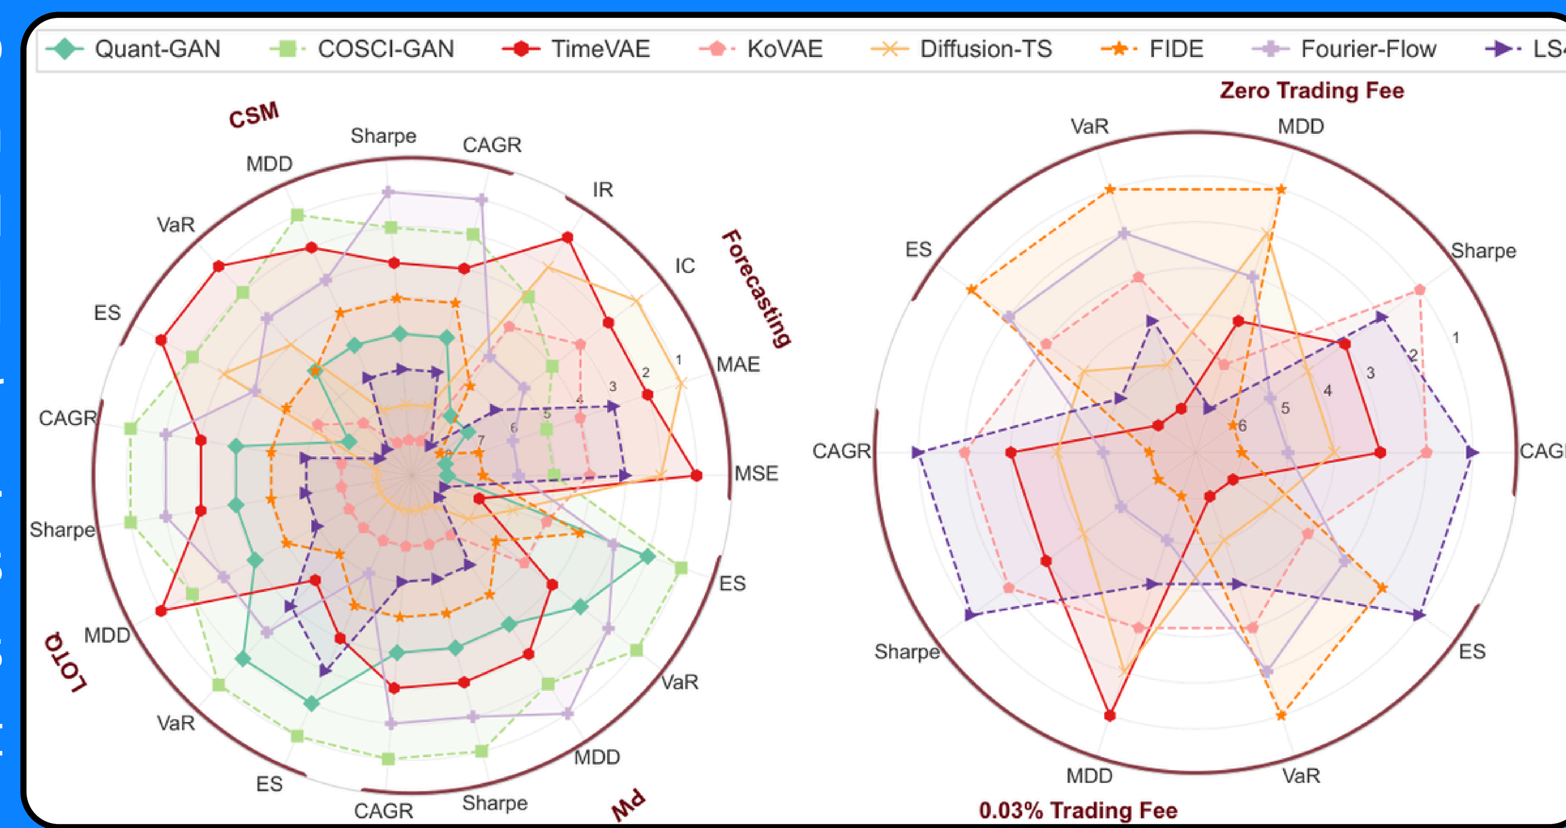

Cryptocurrencies and Decentralized Finance (DeFi)

Investments and mid to low frequency equity and ETF trading strategies

Optimal Fees for Liquidity Provision in Automated Market Makers

Steven Campbell|Philippe Bergault|Jason Milionis|Marcel Nutz

Columbia University   Université Paris Dauphine-PSL   Category Labs

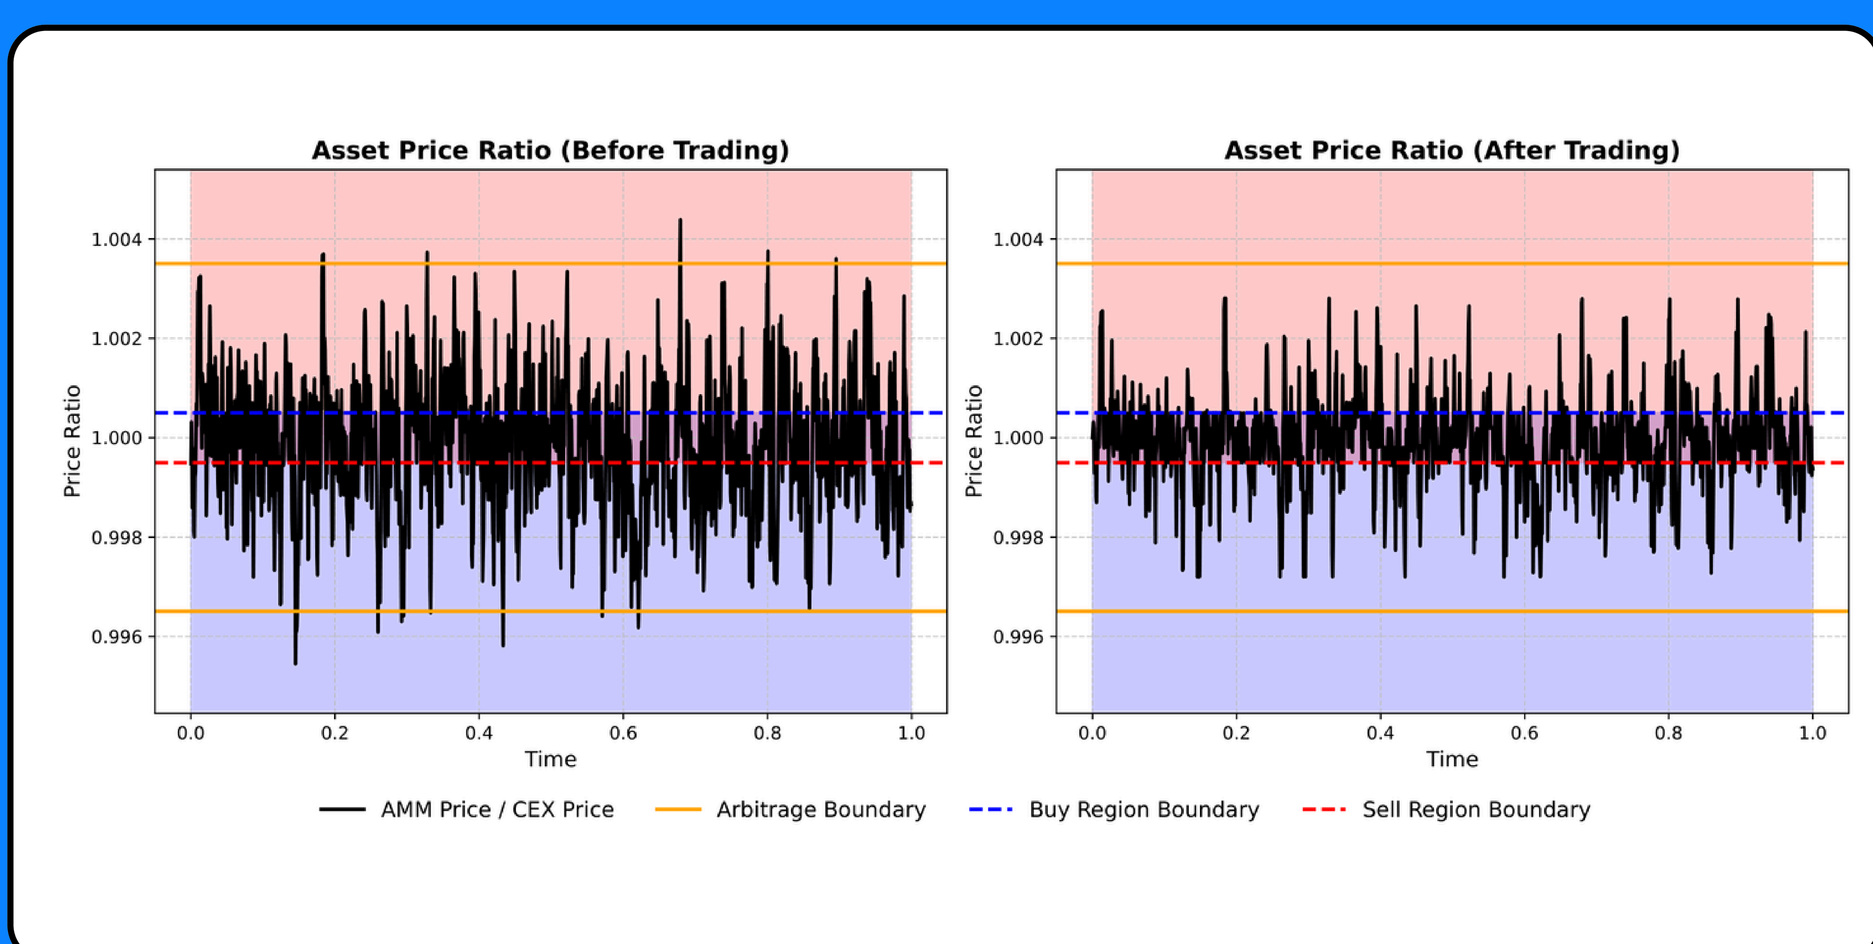

The study finds passive LPs earn more by pricing below CEXs. Optimal AMM fees track 70 to 80 percent of CEX all in cost, stay stable in normal markets, and spike during turmoil. A simulator and CPMM model with routing, arbitrage, and hedged PnL explain profit bands. ETH USDC January 2025 data show CEX cost 6.6 to 11.8 bps, volatility 2.6 percent per day, and 30 to 5.4 bps cuts boosted flow and PnL meaningfully.

Cryptocurrencies and Decentralized Finance (DeFi)

Trading and Market Microstructure

Enhancing Cryptocurrency Sentiment Analysis with Multimodal Features

Chenghao Liu|Aniket Mahanti|Ranesh Naha|Guanghao Wang|Erwann Sbai

The University of Auckland   Queensland University of Technology

Study finds social media mood moves crypto differently. TikTok's video sentiment sparks short speculative bursts, led by Dogecoin. Twitter's tracks longer market trends like Bitcoin. Combining both lifts price and volume forecasts by up to 20 percent. TikTok alone boosts short horizon DOGE prediction 35 percent. Stablecoins barely respond. A multimodal framework fuses Llama 3, MiniGPT4 Video, Whisper, and series tools. Dogecoin transmits volume spillovers. Limits include influencer bias, English focus, and missing causal identification.

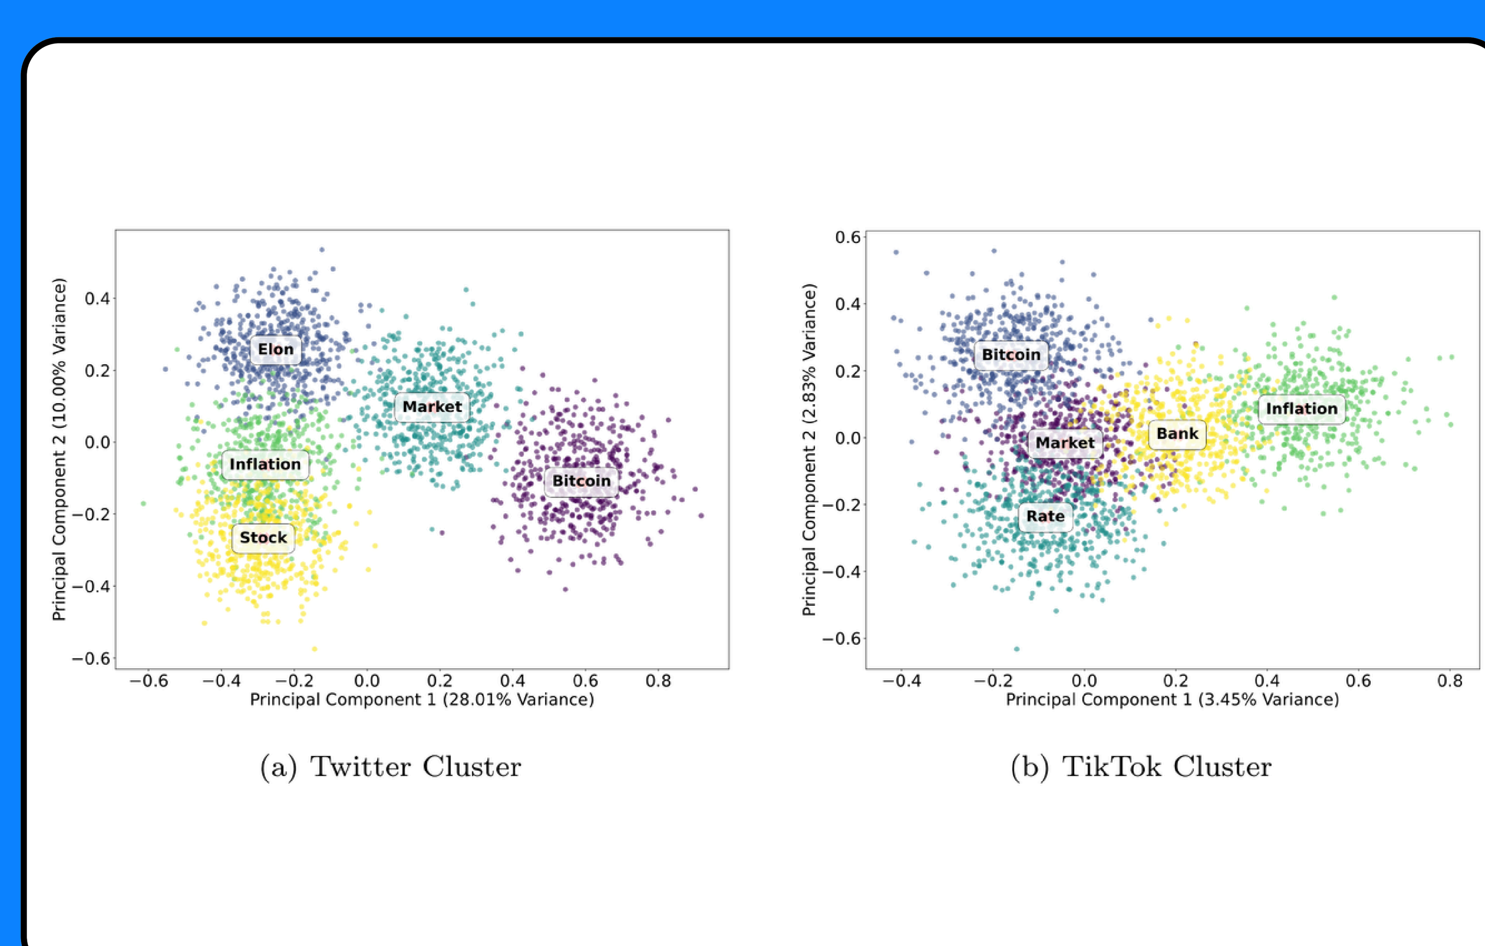

Cryptocurrencies and Decentralized Finance (DeFi)

Sentiment Analysis using AI in Trading

Can LLMs Identify Tax Abuse?

Andrew Blair-Stanek, Nils Holzenberger, Benjamin Van Durme
University of Maryland Carey School of Law, Johns Hopkins University, Telecom Paris - Institut Polytechnique de Paris

	Analysis Verification		Goal Verification with analysis		Goal Verification without analysis		Adversarial Step (without analysis)	
	viable	correct	viable	correct	viable	correct	viable	correct
o3	193/193	166/193	36/36	29/36	35/36	26/36	12/36	0/36
claude-4	193/193	170/193	36/36	33/36	35/36	22/36	20/36	0/36
gemini-2.5	185/193	157/193	34/36	27/36	34/36	31/36	14/36	0/36

Table 1: Results of the first three tests. Analysis verification runs for each of the 193 analysis steps. Goal verification and adversarial step goal-failure verification run for all 36 strategies. The adversarial steps are designed to make the strategy not viable, so a model that performs as well as a domain expert would find 0/36 viable and 0/36 correct.

This study evaluates large language models (LLMs) in identifying and generating U.S. tax-minimization strategies. While performance was generally mediocre, one model invented a novel, viable strategy, demonstrating LLMs' potential to assist tax agencies in combating sophisticated tax abuse by reasoning over complex legal texts.

Miscellaneous

A Comprehensive Survey of Time Series Forecasting: Concepts, Challenges, and Future Directions

Mingyue Cheng, Zhiding Liu, Xiaoyu Tao, Qi Liu, Jintao Zhang, Tingyue Pan, Shilong Zhang, Panjing He, Xiaohan Zhang, Daoyu Wang, Jiahao Wang, Enhong Chen
University of Science and Technology of China

This comprehensive survey examines time series forecasting, covering statistical models, deep learning approaches, preprocessing techniques, and emerging trends like transfer learning and foundation models. It addresses key challenges such as noise, non-stationarity, and long-term dependencies, while highlighting applications in healthcare, finance, and environmental science.

TABLE 1: Overview of the metrics widely evaluated in point-based forecasting and probabilistic forecasting tasks.

Task	Metric	Formula	Advantages	Disadvantages
Point-based Forecasting	MAE	$\frac{1}{n}\sum_{t=1}^{n}\lvert y_t-\hat{y}_t\rvert$	Intuitive and robust to outliers	Insensitive to large errors
	MSE	$\frac{1}{n}\sum_{t=1}^{n}(y_t-\hat{y}_t)^2$	Penalizes deviations in predictions	Sensitive to outliers
	SMAPE	$\frac{1}{n}\sum_{t=1}^{n}\frac{2\lvert y_t-\hat{y}_t\rvert}{\lvert y_t\rvert+\lvert\hat{y}_t\rvert}\times 100\%$	Standardized and works across scales	Sensitive to zeros and large range
	MAPE	$\frac{1}{n}\sum_{t=1}^{n}\left\lvert\frac{y_t-\hat{y}_t}{y_t}\right\rvert\times 100\%$	Intuitive and easily compute	Sensitive to small values
	MASE	$\frac{\frac{1}{n}\sum_{t=1}^{n}\lvert y_t-\hat{y}_t\rvert}{\frac{1}{n}\sum_{t=1}^{n}\lvert y_t-y_{t-1}\rvert}$	Scaled and comparable across datasets	Seasonality-sensitive reference-dependent
	OWA	$\sum_{t=1}^{n} w_t\cdot\lvert y_t-\hat{y}_t\rvert$	Flexible and weighted error handling	Requires careful weight selection and complex
Probabilistic Forecasting	CRPS	$\int_{-\infty}^{\infty}[F(z)-\mathbf{1}(z\geq y)]^2\,dz$	Interpretable and widely applicable	Single-variable only; Sensitive to outliers
	ρ-QL	$2(\hat{y}-y)\left(\rho I_{\hat{y}>y}-(1-\rho)I_{\hat{y}\leq y}\right)$	Ideal for quantile forecasting	Overfitting to asymmetric Data
	NLL	$-\log p_D^f(y)$	Rich gradient info and solid theoretical basis	Lack of intuitive interpretation
	VG	$\sum_{a,b}\left(\lvert y_a-y_b\rvert^p-\underset{x\sim\mathcal{D}^f}{\mathbb{E}}[\lvert x_a-x_b\rvert^p]\right)^2$	Describe spatial correlation	High computational complexity

Time Series Forecasting

Integrating Time Series into LLMs via Multi-layer Steerable Embedding Fusion for Enhanced Forecasting

Zhuomin Chen, Dan Li, Jiahui Zhou, Shunyu Wu, Haozheng Ye, Jian Lou, See-Kiong Ng
Sun Yat-Sen University, National University of Singapore

Table 1: Multivariate time series forecasting under a few-shot setting (10% train steps). All results are averaged from four different forecasting horizons: H ∈ {96, 192, 336, 720}. The best value for each metric is highlighted in red. A lower value indicates better performance.

	Ours		TimesNet		FEDformer		Autoformer		Stationary		ETSformer		LightTS		Informer	
	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
ETTh1	**0.451**	**0.463**	0.869	0.628	0.639	0.561	0.702	0.596	0.915	0.639	1.180	0.834	1.375	0.877	1.199	0.809
ETTh2	**0.376**	**0.411**	0.479	0.465	0.466	0.475	0.488	0.499	0.462	0.455	0.894	0.713	2.655	1.160	3.872	1.513
ETTm1	**0.371**	**0.399**	0.677	0.537	0.722	0.605	0.802	0.628	0.797	0.578	0.980	0.714	0.971	0.705	1.192	0.821
ETTm2	**0.283**	**0.337**	0.320	0.353	0.463	0.488	1.342	0.930	0.332	0.366	0.447	0.487	0.987	0.756	3.370	1.440
ECL	**0.236**	**0.334**	0.323	0.392	0.346	0.427	0.431	0.478	0.444	0.480	0.660	0.617	0.441	0.489	1.195	0.891
Weather	**0.229**	**0.270**	0.279	0.301	0.284	0.324	0.300	0.342	0.318	0.323	0.318	0.360	0.289	0.322	0.597	0.495
Traffic	**0.542**	**0.393**	0.951	0.535	0.663	0.425	0.749	0.446	1.453	0.815	1.914	0.936	1.248	0.684	1.534	0.811

This paper proposes Multi-layer Steerable Embedding Fusion (MSEF), a framework that integrates time series data into Large Language Models (LLMs) by injecting hierarchical time series representations and layer-specific steering vectors across all LLM layers. This prevents TS information loss in deeper layers, enabling superior few-shot forecasting with an average 31.8% MSE reduction over baselines.

Time Series Forecasting