

Estimation of Information Flow in Multivariate Financial Time Series

Final Project Report

Artur Muratov, Ruslan Tasmukhanov, Diana Zhuldubayeva, Yerulan Kongrat

Department of Computer Science, Nazarbayev University, Astana, Kazakhstan

Adviser: Prof. Siamac Fazli

Spring 2026

Abstract—This report is a final report of a two-semester senior project that is related to the estimation and exploitation of directional information flow in multivariate time series. The work has provided a modular, computing-based framework that links 4 classical measures of connectivity, Granger causality (GC), transfer entropy (TE), phase slope index (PSI), and magnitude-squared coherence, to a sliding-window analysis engine, which in turn builds on that result to a computational-scale backtesting system based on LASSO-based cross-asset return forecasting. The main conclusion is that a small statistical signal is convertible into lucrative daily trading performance via the regime filtering, whereas the same signal fails at the one-minute frequency when realistic transaction costs are used.

Index Terms—directional information flow, Granger causality, transfer entropy, phase slope index, coherence, LASSO, cryptocurrency forecasting, backtesting

I. EXECUTIVE SUMMARY

The report is the final presentation of a two-semester senior project that was about the estimation and exploitation of directional information flow in multivariate time series. The project provided a computing based modular framework that incorporates four classical measures of connectivity, Granger causality (GC), transfer entropy (TE), phase slope index (PSI), and magnitude-squared coherence, into a sliding-window analysis engine, and extends the analysis into a full scale backtesting platform based on LASSO based cross asset returns forecasting.

The main issue is that time series correlation is unable to differentiate the process that affects another one. The project will establish and prove the methods, which help to quantify the direction of information transfer and evaluate whether directional signs can be used in real economic terms in cryptocurrency futures markets.

The main quantitative results are as follows. An average cumulative gain of +8.8% across all the 11 major cryptocurrencies over the 11 months of December 2024–March 2026 using a regime-filtered LASSO forecasting strategy ($SMA40+fc>0$) on daily Binance futures data, compared to a cumulative return of −43.7% on a passive buy-and On ETH, which is the best-performing, the strategy has a yield of +244.4% and a Sharpe ratio of +171% and the worst is only −21.9%. This approach to one-minute candles does not work at all, and the forecast-return correlation is reduced to 0.098 and the transaction cost of 9.77% of the gross return consumes all the gross return,

which proves again that the cross-asset information flow signal is a daily-frequency phenomenon.

The project shows that a weak statistical signal (around $\sim 52\%$ directional accuracy) can be turned into a lucrative trading approach via regime filtering and that this can be generalised to a wide range of cryptocurrency assets. The resulting codebase represents a transparent, reproducible and extensible base of future learning-based extensions.

II. INTRODUCTION

Complex systems are usually analyzed as multivariate time series where individual components interact in a manner that cannot be described by pairwise correlation. Correlation gauges the degree to which co-movement exists, but does not give any data on the direction of co-movement. The issue of directional information flow, whether historical values of one process decrease the uncertainty of the future of the other, past their own prediction, or not, is an issue in the domain of neuroscience, as well as the domain of quantitative finance.

Directed connectivity between the regions of the brain is involved in the coordination of cognitive tasks in electroencephalography (EEG) studies [1]. Price-relevant information spreads among assets in financial markets with quantifiable delays: some assets predictably lead other assets, and vice versa, these predictable lead-lag relationships can be used to make predictions [2]. Each of the settings allows a quantitative degree of directional influence that can be used to develop more mechanistic knowledge and make more principled decisions.

The following problem is the focus of this project: provided a collection of multivariate time series, we can weaken the estimate of the direction and magnitude of information flow between components, and can such estimates produce economically significant signals? The process is in two stages. During the first semester, the team developed a modular analysis framework with GC, TE, PSI and coherence that used a sliding-window architecture, tested on synthetic causal systems, and used on cryptocurrency price data. The framework was expanded in the second semester with a rigorous LASSO-based forecasting pipeline, full-scale backtesting engine on 11 assets with a horizon of 1,549 days and systematic exploration of 15 formal hypotheses of the structure and exploitability of the information flow signal.

The rest of this report has the following structure. Section II outlines the major constraints that impact the project. Section III presents the prior work, which is relevant and puts the methodology into perspective. The system architecture and approach to computation are outlined in Section IV. Section V gives an account of the project implementation in the two semesters. The results of the evaluations are given in Section VI. Section VII ends with a findings and future directions discussion.

III. IDENTIFICATION OF CONSTRAINTS

A variety of constraints in the economic, ethical, technical, and sustainability aspects influenced the design, implementation and deployment of this system.

A. Economic Constraints

Experiments were done on commodity hardware (12 cores) and cloud compute budgets were not available. This limited the search of hyperparameters that were possible and the number of assets that could be run simultaneously. The entire daily backtest (11 coins, 12 LASSO models at a time, 165 parallel tasks) took about 34 minutes to execute and thus larger grid searches of window sizes and lag orders were constrained. On the data side, all the data of the cryptocurrencies was obtained at no cost using the Binance public API, which ruled out the use of more expensive data feeds of institutions with low latency or ticks.

B. Ethical Constraints

The main ethical hazard is that of misinterpretation of causal results. Directional tests like Granger causality test the existence of statistical predictability as opposed to causality. Exaggerating this fact may cause the practitioners to give structural significance to what are conditional correlations, which may be confounded by some other factors. The team thus confined itself to predictive directionality in all conclusions and time and again separated statistical causality and stronger mechanistic assertion.

The second issue that could be of concern is the misuse of trading signals. Weak predictive edge based algorithmic trading strategies are subject to overfitting, misuse, or application in inappropriate circumstances. Because of this reason, the study focuses on transaction costs, out-of-sample analysis and robustness checks as opposed to reporting gross returns.

C. Data and Legal Constraints

The financial information used in this project is freely available on Binance, open API with terms allowing the use of the information in research. No material, no proprietary and no non-public information was utilized. Similarly, there is no personal data or use of human subjects or patient records in the project and that is why no privacy sensitive processing pipeline was needed. These legal restrictions made data management easier, but did not eliminate responsible reporting and reproducibility.

D. Technical and Sustainability Constraints

The framework has been executed in Python 3.13 with the aid of common open-source packages like NumPy, SciPy, pandas, scikit-learn and matplotlib. In this decision, maintainability and reproducibility is better, but there are practical limits on the performance at runtime. The forecasting problem is of high dimension: when the model of the assets of interest is that of p lags and is stateful with a number of state variables of number of state variables M , the generic predictive model can be expressed in

$$\hat{r}_{i,t+1} = \beta_0 + \sum_{j=1}^d \sum_{\ell=1}^p \beta_{j,\ell} r_{j,t-\ell+1} + \sum_{k=1}^m \gamma_k z_{k,t}, \quad (1)$$

where $r_{i,t+1}$ is the next-period return of asset i and z_k is a state auxiliary variable. Memory usage and repeated model fitting are the primary computation bottlenecks, as d and p grow. In sustainability terms, the project was specifically biased towards interpretable and computationally light-weight approaches instead of large-model training, which is energy intensive.

E. Manufacturability and Scalability

The sliding-window LASSO pipeline has a quadratic scaling to the number of assets due to the number of new targets and new cross-asset lagged predictors with each new asset. The 11 assets at 7 lags and 4 transformed variables per asset to form the 308 input features. It is still feasible on commodity hardware, but requires either more effective dimensionality reduction or distributed execution to naively scale to dozens of assets. This means that the current system can be produced and replicated as a research prototype, though not as an engineered platform at the scale of production, i.e. a large-universe.

IV. BACKGROUND AND RELATED WORK

A. Directional Information Flow Measures

The directional influence between time series quantification problem has been tackled through two paradigms: model-based and information-theoretic.

The model-based method, which is the basis of the approach, is Granger causality, formalised by Granger [3] in 1969 in the article Granger causality. One process, X , Granger-causes another process, Y when lagged values of process X can give statistically significant information about future values of process Y that is not already present in lagged values of process Y . In the case of restricted and unrestricted autoregressive models, the standard test statistic is

$$F = \frac{(RSS_r - RSS_u)/m}{RSS_u/(T - k)}, \quad (2)$$

where RSS_r and RSS_u are the sum of squares of the residual of the restricted and unrestricted models respectively, m is the number of restrictions, T is the sample size, and k is the number of unrestricted parameters..

Transferred entropy, which was introduced by Schreiber [4] is an information-theoretic analogue, model-free and nonlinear in nature. To embed past vectors of $x_t^{(\ell)}$ and $y_t^{(k)}$,

$$TE_{X \rightarrow Y} = \sum p(y_{t+1}, y_t^{(k)}, x_t^{(\ell)}) \log \frac{p(y_{t+1} | y_t^{(k)}, x_t^{(\ell)})}{p(y_{t+1} | y_t^{(k)})}. \quad (3)$$

This value is an indication of the decrease in the uncertainty of Y_{t+1} that occurs due to the knowledge of the past of X given the past of Y . Empirical TE is estimated by discretising continuous signals into bins in practice.

Magnitude-squared coherence is a measure of frequency-dependent linear interaction between two signals:

$$\gamma_{XY}^2(f) = \frac{|S_{XY}(f)|^2}{S_{XX}(f)S_{YY}(f)}, \quad (4)$$

, with $S_{XX}(f)$ and $S_{YY}(f)$ auto-spectral densities and $S_{XY}(f)$ the cross-spectral density. Coherence is zero to one, and is symmetric, so is useful as a baseline but not sufficient in itself in the directional inference.

Directionality in the frequency domain is restored by the phase slope index, introduced by Nolte *et al.* [5], which takes advantage of the phase of the complex coherency, $CXY(f) = SXY(f) / \sqrt{SXX(f)SYY(f)}$:

$$\text{PSI}_{X \rightarrow Y} = \sum_k w_k \Im[C_{XY}(f_k + \Delta f) C_{XY}^*(f_k)], \quad (5)$$

, where $\Im\{\cdot\}$ and $\cdot^*$ denote complex conjugation. Having a positive PSI means that in the net spectral sense, X is a cause of Y .

B. Cross-Asset Information Flow in Finance

Directional flow measures have a considerable literature with regard to their financial use. Lo and MacKinlay [2] have reported systematic lead/lag correlations between large- and small-cap stocks, and subsequent connectiveness research formalised the spillovers and structure of financial systems in terms of networks and connectivity, including connectivity/network effects, as well as interrelations amongst events and agents [6], [7]. The high-frequency, 24/7 nature of trading in cryptocurrency markets is a natural setting to investigate the dynamics of the dissemination of information, yet the very structure of the market also beckons arbitrage that can be executed quickly, eliminating short-lived opportunities.

C. LASSO Regression for Return Forecasting

LASSO (Least Absolute Shrinkage and Selection Operator), first proposed by Tibshirani [8], optimises the residual sum of squares with an added penalty on the coefficient vector which is based on the ℓ_1 norm:

$$\hat{\beta} = \arg \min_{\beta} \left\{ \frac{1}{2n} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1 \right\}. \quad (6)$$

. The sparsity created by the use of the ℓ_1 penalty can automatically perform variable selection, by fixing irrelevant features to zero and in the process creating sparsity in the model. This fits well to cross-asset forecasting, where there are

numerous lagged features, which are only weakly informative, and naive dense models easily overfit.

D. Positioning of This Work

The project has three contributions in comparison to the previous work. First, it offers a single, unified, modular framework capable of calculating a number of directional measures in one sliding-window engine. Second, it relates such measurements to a stringent forecasting and backtesting pipeline instead of halting at the visual analysis of descriptive measurements. Third, it records good and bad results: daily-frequency directional signal can be utilized under regime filtering, but minute-level transfer cannot be used after realistic costs. Such balance is needed since it explains the boundaries of the proposed solution that has been based on computing.

V. PROJECT APPROACH

A. System Architecture Overview

The system is designed to be a modular pipeline with four main layers: (1) data ingestion and preprocessing, (2) directional connectivity estimation, (3) strategy and backtesting, and (4) evaluation and visualisation. The interfaces of each layer are clean to the other, allowing them to be replaced or extended without altering the rest of the system.

The following data flows. Loading Modality-specific loaders take raw OHLCV (open, high, low, close, volume) data and available on the Binance public API and do simple cleaning (aligning timestamps and handling missing values). The resulting normalised and stationarised signals, which are the log returns, the log-differenced volume, the Parkinson volatility and the Amihud illiquidity are sent to the sliding-window engine. Each window contains all the four connectivity estimators (GC, TE, PSI, coherence) applied to all the ordered pairs of series and the directional scores are stored in a structured form by a window start time, pair and method. These scores are inputted into two downstream layers which are parallel to each other: the analytical layer (cosine similarity, heatmaps, best-window detection) and the strategy layer (LASSO forecasting, signal generation, backtesting).

B. Data Preprocessing

In the case of financial data, the raw closing prices, P_t are converted to log returns,

$$r_t = \log P_t - \log P_{t-1}. \quad (7)$$

. This levels variance, eliminates price-level effects and makes the stationarity assumption more justifiable. The first differences of log volume are computed to get volume changes. Parkinson volatility, a range based estimator is calculated as

$$\sigma_{P,t}^2 = \frac{1}{4 \ln 2} \left(\ln \frac{H_t}{L_t} \right)^2, \quad (8)$$

. with H_t and L_t the high and low prices of the period. Amihud illiquidity ratio

$$\text{ILLIQ}_t = \frac{|r_t|}{V_t}, \quad (9)$$

where V_t is the volume of trading. Time alignment, with no window-level computation, asynchronous data collection can introduce spurious causality unless time alignment is done, by having all assets have a shared time index.

C. Sliding Window Engine

Time series of financial and neural data are hardly ever stationary in the long run. All connectivity measures are calculated in a sliding-window framework to obtain time-varying directional structure. Considering a pair of series and parameters (W, S) , consecutive windows are

$$\mathcal{W}_n = \{t_n, t_n+1, \dots, t_n+W-1\}, \quad t_{n+1} = t_n + S. \quad (10)$$

In each window, all the four measures (GC, TE, PSI, and coherence) are calculated in both the $X \rightarrow Y$ and $Y \rightarrow X$ directions. The result is a time-stamped series of directional scores of every method direction combination. This representation allows analyzing directional dominance with time resolution, directly comparing methods, and integrating with the real-time streaming mode.

D. LASSO Forecasting Pipeline

Each target asset j has 12 LASSO models estimated: a single model (28 features: the lags of asset j itself), ten pairwise models (56 features: the lags of asset j and the lags of a single other asset), and one full model (308 features: the lags of asset j and the lags of all 11 assets). The parameter of the penalty of the type of the ℓ_1 penalty, denoted by λ is chosen through five-fold cross-validation (a rolling window of 714 days). The model is re-estimated on the last 714-day window on a test day, resulting in a purely out-of-sample forecast sequence.

One of the most important implementation findings is related to the forecast threshold. The initial approach defines a constant significant threshold of 0.0025: only in case of the exceedance of the value of the said significant threshold by the value of the $|\hat{r}_{t+1}|$ we take positions. This threshold applied to our data would result in 0 trades in all 11 coins since LASSO regularisation tends to compress the forecast magnitudes around 0. The standard deviation of actual returns divided by forecast standard deviation is about $30\times$ and this makes a constant standard deviation of actual returns non-transferable across time and assets. We thus substitute it with an adaptive threshold per asset which is the 25th percentile of the absolute non-zero forecast distribution, calibrated on the training set:

$$TH_j = Q_{0.25}(|\hat{r}_{j,t+1}| : |\hat{r}_{j,t+1}| > 0). \quad (11)$$

This is so that about 75 percent of non-zero forecasts above the threshold are created producing a significant trading frequency but still eliminating the weakest signals.

E. Trading Strategies

Four strategies are evaluated. Buy-and-hold maintains a constant long position and serves as the passive benchmark. Paper “All” trades the raw full-model forecast: long when $fc_t > TH$, short when $fc_t < -TH$, and flat otherwise. SMA40+fc>0 adds a 40-day simple moving average regime filter: the strategy goes long only when the asset price is above its 40-day SMA and the forecast is positive, and short only when price is below the SMA and the forecast is negative. When regime and forecast disagree, the position is flat. Formally,

$$\text{pos}_{i,t} = \begin{cases} +1, & \hat{r}_{i,t+1} > 0 \text{ and } P_{i,t} > SMA_{40}(P_i)_t, \\ -1, & \hat{r}_{i,t+1} < 0 \text{ and } P_{i,t} < SMA_{40}(P_i)_t, \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

The Agree+Profit+fc>TH strategy requires three simultaneous conditions for entry: at least 8 of 12 LASSO model variants must agree on direction, the trailing 10-day average of forecast times realised return must be positive, and $|fc_t|$ must exceed the adaptive threshold. Transaction costs of 0.08% round-trip (Binance futures taker fee) are applied at every position change, so net returns are evaluated as

$$R_{i,t} = \text{pos}_{i,t-1} r_{i,t} - c |\text{pos}_{i,t} - \text{pos}_{i,t-1}|. \quad (13)$$

F. Multi-Measure Integration and Validation

To evaluate the consistency of the four directional measures, time-resolved profile of the four directional measures are normalised and compared through cosine similarity between all pairs of assets:

$$\cos(a, b) = \frac{a^\top b}{\|a\|_2 \|b\|_2}. \quad (14)$$

By summing these similarities over all pairs of methods, we obtain a method-by-method agreement array. Whenever methods are highly consistent within a specific window the confidence of the identified directional pattern is increased. The framework further determines, per method and pair, the window that gives the highest directional score (“best detection window”), so that visualisation can be made of what price dynamics contribute to the largest measured flow.

G. Tools and Third-Party Components

Python 3.13 is used to implement the system. NumPy and SciPy are used as core numerical computation. scikit-learn (LassoCV) provides the LASSO estimator and cross-validation tools, and pandas offers the time-indexed data structures of the sliding-window engine, and matplotlib creates all the figures. Parallelisation between assets and variants of a model is done with Python multiprocessing module, with 12-core utilisation. It ingests the data through the Binance public klines REST API and caches all the obtained data in a local storage to facilitate reproducible re-runs. The Butterworth low-pass filter applied in the meta-analysis is a part of the SciPy signal processing, run in causal (and therefore look-ahead biased) mode (sosfilt) to avoid look-ahead biases in feature construction.

H. Team Roles

Unless there was a particular division of responsibility between the four main system components. The data pipeline was headed by Diana Zhuldubayeva: API integration, pre-processing, alignment of timestamps, and caching infrastructure. The directional connectivity measures (GC, TE, PSI, coherence) and sliding-window estimation engine were implemented under the guidance of Yerulan Kongrat. Backtesting engine (position management, transaction cost modelling, performance metric computation, and the one-minute frequency experiment) was headed by Ruslan Tasmukhanov. LASSO predictive architecture was headed by Artur Muratov: feature engineering, rolling-window model-fitting, adaptive threshold-calibration, and the multi-coin parallel execution system. Experimental design, analysis and report writing were shared by all members.

VI. PROJECT EXECUTION

A. Semester 1: Framework Construction

The initial semester was devoted to the framework of theoretical and computational basis. The team has started by surveying the literature on directional connectivity and have found GC, TE, PSI, and coherence as the four main measures, which have complementary sensitivity profiles: GC can be used to detect linear predictive structure, TE is used to detect nonlinear information transfer, PSI is used to detect spectral phase-based directionality, and coherence is used as a symmetric baseline. An implementation and validation of the sliding-window engine was done on synthetic autoregressive systems having known causal structure, and then on real data.

The first experiments were devoted to BTC-ETH pair where one-minute cryptocurrency data are utilized. The initial directional scores profiles were obtained as a result of a full-day connectivity analysis using 15-minute sliding windows. These initial findings showed a common problem that the four measures tended to concur in high-magnitude windows but to disagree in low-magnitude windows, and how to combine their signals in a reliable way. Cosine similarity framework—a comparison of normalised time-series of method-based directional scores—was created directly in response to this difficulty.

The team completed the work at the end of Semester 1, with a report containing the sliding-window connectivity engine, first tests with synthetic and real data, and a suite of visualisation tools such as time-resolved directional profiles, heatmaps, and best-detection-window plots. The final report and presentation of Semester 1 suggested the concrete plan of Semester 2: backtesting of the GC/TE/PSI-based trading signals implementation, the introduction of deep learning sequence models (temporal CNNs, RNNs, transformers) as the benchmark, systematic search of hyperparameters.

B. Semester 2: Pivots and Key Decisions

The Semester 2 roadmap had its major revision during the initial month. The deep learning aspect had been given low priority to a more stringent empirical study of the LASSO-based cross-asset forecasting framework that provided a more

manageable and directly comparable baseline. There were two aspects behind this decision. To begin with, the computational infrastructure was not in place to support training transformer models over sequence lengths of 1,549 days of 11 assets during the project time frame. Second, and more significantly, the team realised that the very basis of the LASSO replication had a critical flaw: the fixed forecast threshold of 0.0025 as dictated by the reference methodology led to zero trades when applied to the 20222026 data period, since the regularisation of LASSO compresses the magnitudes of forecasts to values that are significantly smaller than the critical forecast threshold.

This was the initial significant deliverable of Semester 2, in terms of correcting this flaw. The adaptive threshold calibration algorithm in Section IV-D, has been created, tested on the training set and proven to be required to get a trade happening. This observation was more broadly methodologically important: as it showed that fixed thresholds are not transferable between assets and market regimes.

The second crucial choice was to experiment with the strategy at a frequency of one minute. Semester 1 framework was set up based on intraday BTC-ETH data, so the question of the day-frequency LASSO signal transferring to shorter timescales was natural. The one-minute experiment that they controlled provided a resounding yes to this question: a limited gross advantage was exceeded by the frequency of regime changes and real world transaction costs.

The third execution problem was how the computational cost of the 165 parallel LASSO tasks to execute the complete 11-coin backtest could be managed. The early sequential code took more than eight hours, per complete run, and it was not possible to conduct iterative experimentation. The computation was restructured by the team with the multiprocessing module of Python and breaking the computationally intensive (in comparison) “All” model into smaller sub-tasks, and finally the computation time was reduced to approximately 34 minutes per full daily execution.

C. Collaborative Process

The project team was a mini research engineering unit as opposed to isolated contributors. Decisions on design were also made collectively, the issues faced during implementation were addressed as a team and failures were reported instead of being concealed. This was particularly relevant during Semester 2, during which a few of the initially promising concepts, such as fixed thresholds, meta-learning gates and some other trading heuristics, were not as strong as they should have been. The disheartening outcomes changed focus to more solid simplifications, especially the regime-filtered strategy which currently constitutes the core empirical finding of the paper.

VII. EVALUATION

A. Forecast Quality and Model Comparison

On 1,156 out of sample days (November 2019–December 2022), directional accuracies of all 12 LASSO model variants, solo, in 10 pairwise combinations, and “All”, have values




Fig. 1. Model comparison over the full out-of-sample period in terms of MSE and directional accuracy.




Fig. 2. Rolling MSE and rolling directional accuracy of the aggregate model. Error spikes are concentrated at particular events and regimes.

ranging between 51.4% and 51.9%, as reported in Fig. 1. The MSE values are also similar (around 0.00146 all the way), so the addition of cross-asset features does not enhance predictive performance of the solo model significantly.

The LASSO regularisation shrinks the magnitudes of forecasts to about a third of the actual magnitude of the standard deviation of the returns (forecast std ≈ 0.00127 , actual return std ≈ 0.03820). This compression is not a sign of failure of the model but rather a part of it, the trading advantage lies in the sign of the forecast and not its strength. The rolling directional accuracy in Fig. 2 vary significantly with time, swining along the 50% percent line with intervals when seven-day accuracy is over 80% percent and times when it is below 20% percent. This large temporal variance is one point of note: the overall 51–52% accuracy is a composite of regimes where the model is locally far more or far less accurate.

B. Daily Strategy Performance

The evaluation of trading performance is organised at two levels: a detailed ETH case study and a broader cross-asset generalisation across the full 11-coin universe.

1) *ETH Single-Asset Deep Dive*: The most complete single-asset analysis was conducted on ETH using the 10-coin LASSO feature set. Results across four strategies are presented in Table I.

The SMA40+fc>0 will take the -41.6% passive loss and change it into a $+244.4\%$ cumulative gain of 286 percentage points of alpha but will only be in the market 49% of the time. The Sharpe ratio improves from -56.6% to $+171.0\%$ and the maximum drawdown is reduced from -62.2% to -21.9% . The triple-gated entry requirement of the Agree+Profit strategy is overly restrictive: it invests capital on just 36 percent of days

TABLE I
ETH STRATEGY PERFORMANCE COMPARISON (DEC. 2024–MAR. 2026
TEST PERIOD, 10-COIN LASSO FEATURES).

Metric	Buy & Hold	Paper “All”	SMA40+fc>0	Agree+Profit+fc>TH
Cumulative return	-41.6%	-33.2%	$+244.4\%$	-63.7%
Sharpe ratio	-56.6%	-51.7%	$+171.0\%$	-182.9%
Max drawdown	-62.2%	-50.1%	-21.9%	-67.6%
Days in market	100%	75%	49%	36%
Long / Short	100/0%	64/11%	33/16%	30/6%
Win rate	49.5%	48.9%	54.2%	41.9%

and has a win rate of just 41.9 percent resulting in a worse outcome than the passive benchmark.

2) *Cross-Asset Generalisation: All 11 Coins*: The SMA40+fc>0 strategy was applied across all 11 cryptocurrencies using the full 11-coin LASSO feature set. Results are presented in Table II.

The strategy produces positive alpha on 10 of the 11 coins, with an average return of $+8.8\%$, compared with an average buy-and-hold return of -43.7% , for an average alpha of $+52.5\%$. The four coins with positive absolute returns (ETH, LINK, BNB, and ADA) are associated with assets that exhibit stronger cross-asset predictability. The only coin that performs worse than buy-and-hold is TRX, which followed a long, unidirectional upward trend over the test period; as a result, the strategy’s short signals were counterproductive. The equity curves for each of the 11 coins under the SMA40+fc>0 strategy are shown in Fig. 3.

C. The Role of Regime Filtering

The SMA40 filter improves directional accuracy by concentrating trades in regimes where the LASSO forecast is systematically more reliable. The mechanism is straightforward: in uptrends (price above SMA40), the LASSO’s positive forecasts tend to be correct at a higher rate (approximately 54–55%) than in choppy or downtrending markets. The filter eliminates trades in regimes where the model’s edge is absent or negative, effectively converting a marginally predictive model into a profitable strategy.

Figure 4 shows forecast directional accuracy decomposed by regime. For most coins, long trades (uptrend and positive forecast) and short trades (downtrend and negative forecast) both exhibit accuracy above 50%, while the ungated model hovers near the 50% baseline. The resulting 2–4 percentage point improvement in filtered accuracy is sufficient to generate substantial profits when applied consistently with low per-trade transaction costs.

The SMA40 filter is enhanced to have directional accuracy by focusing trades in the regimes where LASSO forecast is more consistently accurate. The process is simple, when in an uptrend (price above SMA40), the LASSO gives positive predictions, it is more likely to be accurate (approximately 54–55%) than in choppy or falling markets. The filter filters out trades in regimes where the model has an edge that is negative or zero, then transforms a marginally predictive model into a profitable strategy.

TABLE II
SMA40+fc>0 STRATEGY PERFORMANCE ACROSS 11 CRYPTOCURRENCIES. TEST PERIOD: DEC. 2024–MAR. 2026.

Coin	B&H	SMA40+fc>0	Sharpe	Max DD	Days In	Alpha
ETH	−41.6%	+181.4%	+147.4%	−21.9%	48%	+223.1%
LINK	−63.1%	+61.8%	+65.1%	−28.0%	46%	+124.9%
BNB	−10.2%	+10.1%	+20.0%	−32.5%	49%	+20.3%
ADA	−74.4%	+6.3%	+7.7%	−38.1%	50%	+80.7%
SOL	−57.6%	−0.4%	−0.6%	−34.3%	42%	+57.2%
BTC	−31.7%	−18.0%	−66.4%	−32.6%	40%	+13.7%
DOGE	−71.4%	−28.9%	−48.7%	−59.3%	35%	+42.5%
ETC	−70.6%	−32.0%	−50.6%	−63.4%	63%	+38.6%
LTC	−47.4%	−38.4%	−66.7%	−55.2%	51%	+9.0%
XRP	−41.9%	−39.7%	−92.6%	−49.9%	31%	+2.2%
TRX	+28.7%	−5.9%	−20.0%	−30.5%	50%	−34.6%
AVG	−43.7%	+8.8%	—	−40.5%	—	+52.5%

Figure 4: Equity Curves Across All 11 Coins
SMA40 + fc>0 Strategy




Fig. 3. Equity curves across the 11-coin universe for the regime-filtered strategies.

Figure 4 indicates the accuracy of the forecast direction by regime. With most coins, both long trades (uptrend and positive forecast) and short trades (downtrend and negative forecast) are accurate above 50%, but the ungated model is close to the 50% mark. The 2–4 percentage-point increase in filtered accuracy leads to enough profit to accumulate to significant profits when implemented on a regular basis with minimal transaction costs per trade.

D. Frequency Scalability: The 1-Minute Experiment

To test Hypothesis H9—the information flow signal generalises to shorter time scales—the SMA40+fc>0 strategy was

scaled to one-minute candles on seven days of data (10,672 bars per coin). The feature engineering, model architecture and strategy logic remained the same as the daily version with the parameters being scaled proportionally (rolling window: 1,440 bars; SMA period: 40 minutes; maximum lag: 7 minutes). Table III shows the results.

The plan goes awry. There are three compounding failure modes as indicated by the diagnostic decomposition. To begin with, the forecast-return correlation decreases at frequency of one-minute (0.098) compared to a frequency of daily (around 0.15), which implies that the signal of cross-asset information flow is not present at this frequency. The directional accuracy




Fig. 4. Forecast directional accuracy by regime per coin. Regime filtering improves conditional accuracy.

TABLE III
STRATEGY PERFORMANCE AT ONE-MINUTE FREQUENCY (ETH, 7-DAY WINDOW, 32-HOUR TEST PERIOD).

Strategy	Cum. Return	Sharpe	Win Rate
Buy & Hold	+1.80%	+850%	50.5%
Paper All	-6.29%	-3016%	47.6%
SMA40+fc>0	-7.25%	-4674%	46.3%

declines to 50.3% which is statistically the same as a coin toss. Second, the gross positive return created by the strategy is small of +2.79% but the 40-minute SMA changes regime 4.5 times an hour resulting in 252 position changes in 32 hours. This is estimated to bring about 9.77% in transaction costs at 0.08% round-trip cost per change which completely nullifies the advantage. Third, the magnitudes of the forecasts (mean ≈ 0.0000016) are about $500\times$ smaller than the standard deviation of the actual returns (0.00079) and as a result the LASSO is virtually making a forecast of zero at each step.

The experiment gives a mechanistic reason as to why the daily signal exists: at daily frequency, information flows between crypto assets along liquidity and sentiment contagion channels along with portfolio rebalancing channels, which act on multi-hour to multi-day timescales. At a frequency of one minute, the algorithmic arbitrage imposes an almost instant price synchronization of correlated assets, removing any lagging arbitrage. The additional Hypothesis H12 the Sharpe ratio is monotonically increasing between one minute and 60 minutes of decision horizon, is convergent evidence that this interpretation is correct.

E. Meta-Learning and Forecast Error Analysis

A second line of inquiry hypothesized that with a second-layer classifier, it would be possible to predict when the LASSO forecast would be directionally correct so that the trades could be gated on a confidence signal (H5). To predict the probability of a directional accuracy in the next seven day window (over 55%) a logistic regression meta-classifier was fitted on 10 causal features: rolling MSE at 7/14/30-day windows, 7-day and 14-day Parkinson volatility, volatility percentile rank, forecast disagreement across the 12 model forms, low-pass filtered forecast

The meta-classifier had an AUC of 0.493 on the validation set and 0.466 on the held-out test set- lower than the 0.5




Fig. 5. Meta-classifier feature importances. The ranking is interpretable, but the classifier itself remains ineffective.

random baseline. The feature importance analysis (Fig. 5) indicates that there is a strongly negative coefficient of short-term volatility: volatility in the short-term is associated with a lower forecast accuracy, which is in line with the MSE volatility-regime analysis indicating that high-volatility periods have mean MSE approximately $45\times$ higher than low-volatility periods (Fig. 6).

Although the meta-classifier has the correct intuition (volatility predicts forecast quality), it cannot transform it into a per-day trading gate of use. The regime-level trend (trending versus choppy, which is proxied by SMA40 filter) is much more predictive than any combination of day-level features. The monthly MSE heatmap singles out the largest single failure event as March 2020 with MSE of 99.9×10^{-4} , approximately 12 times greater than the average monthly value. The 9 of the 11 MSE peaks identified are found in bear market or high-volatility episodes and this gives solid evidence that the LASSO model that is trained on the relatively calm historical windows is systematically miscalibrated in crisis periods.

Final synthesis is given by the three-panel low-pass filter analysis. Regime-level trend direction is captured in the low-pass filtered BTC return. The LASSO forecast has a low-pass filtered forecast which traces a positive bias that changes slowly across the 2020–2022 period and then becomes negative in late 2022. The low-pass filtered MSE, the resulting figure of which features 11 local maxima highlighted, indicates that error concentration follows local spikes of volatility with a short lag. Collectively, these panels describe the information structure that the SMA40 filter takes advantage of: at the time the low-pass filtered return and the low-pass filtered forecast are aligned and the MSE is minimal, the strategy can be most beneficial.

F. Summary of Hypothesis Testing

Table IV summarises the 15 formal hypotheses tested and their outcomes.

VIII. CONCLUSION AND POSSIBLE FUTURE WORK

A. Summary of Findings

This project provided an all-encompassing framework to predict directional information flow in multivariate time series




Fig. 6. MSE regime analysis showing that forecast errors concentrate in high-volatility conditions.




Fig. 7. Monthly average MSE heatmap across time and asset.

and used the model on the cryptocurrency futures markets. The main results are as following.

Cross-asset information flow exists but is weak. It has a weak predictive relationship between cryptocurrency returns, and LASSO regression can recover a directional signal with a moderate out of sample prediction with about 52% accuracy. But on its own, this signal cannot be used to make profitable trades without further filtering.

Regime filtering is the key ingredient. The SMA40+fc>0 strategy (only trades when market trend and LASSO predict agree) has an average +8.8% return in 11 coins versus −43.7% with buy-and-hold and provides positive alpha on 10 of 11 assets. The filter raises the conditional directional accuracy of about 52% to about 54–55% by removing counter-trend trades, where LASSO is more likely to be inaccurate. Such a 2–4 percentage point increase is enough to reach the profitability level in the case of transaction costs less than 0.1%.

The signal is a daily-frequency phenomenon. At a one-minute frequency, the forecast-return correlation is reduced to 0.098 (near-random), and all gross return is eaten up by transaction costs in excess of regime switching (4.5 SMA switches per hour). This signal increases monotonically with decision horizon, which validates that the information spread of cross-asset in cryptocurrency markets is on a day-scale, rather than a minutes-scale.

Adaptive threshold calibration is necessary. The set threshold of the reference methodology (0.0025) would give zero trades across all 11 assets in the 2024 to 2026 period. Magnitudes of LASSO forecasts are non-stationary and dependent on assets; they cannot be transferred between historical periods or particular assets using static thresholds that are calibrated to historical periods.

Simple filters outperform complex models. A meta-classifier of logistic regression which is trained on 10 features achieves



Fig. 8. Low-pass trend analysis for BTC returns, forecasts, and error concentration.

TABLE IV
SUMMARY OF THE 15 FORMALLY TESTED HYPOTHESES AND THEIR OUTCOMES.

H#	Hypothesis	Result
H1	Cross-asset Granger causality exists	Partial support
H2	Raw LASSO forecasts are profitable	Not supported
H3	SMA40 regime filter makes LASSO profitable	Strongly supported
H4	Model consensus (8/12 agreement) improves signals	Partial (inferior to SMA40)
H5	ML meta-models predict forecast accuracy	Not supported (AUC ≤ 0.51)
H6	Hurst exponent identifies trading regimes	Not supported (destructive)
H7	Position sizing / vol targeting improves Sharpe	Not supported (alone)
H8	Stop-loss mechanisms improve outcomes	Modest support
H9	Information flow signal exists at one-minute frequency	Not supported (strongly rejected)
H10	GC/TE/PSI/coherence predict direction at one-minute frequency	Not supported (after costs)
H11	Auto-detecting the leader improves pairwise trading	Mixed
H12	Longer decision horizons improve stability	Supported
H13	Correlation regimes affect profitability	Modest support
H14	Ridge outperforms LASSO for single-asset prediction	Supported
H15	Adaptive thresholds are necessary	Strongly supported

AUC of at most ≤ 0.51 —worse than random—on the task of predicting day-level forecast accuracy. The regime-level pattern that the SMA40 captures is strong in the situations when the per-day pattern cannot be distinguished by noise.

B. Future Work

There are a number of suggestions of future research based on the findings and limitations of the project.

Extending to deep learning predictors. The Semester 1 plan involved training directional score sequence temporal CNNs, RNNs and transformers. This is a logical extension: the GC/TE/PSI profiles generated using the framework are time-indexed structured features that are very appealing to sequence models. Daily-frequency LASSO strategy offers a tangible benchmark by which deep learning predictors are estimated in the identical backtesting framework, with the same assumptions on the cost of transactions.

Improving threshold calibration. The adaptive 25th-percentile threshold that is utilized in this study is a simple remedy to a basic non-stationarity issue. Additional more principled techniques such as forecast standardisation via rolling volatility, quantile-regression thresholds or conformal prediction intervals could enhance the stability and generalisability of the trading gate.

Portfolio-level optimisation. The strategy at hand considers each asset separately. This could be enhanced by a portfolio-level strategy based on inputs of directional flow scores to a mean-variance or risk-parity optimiser with perhaps the cross-asset correlation regime detection identified in H13, to take advantage of diversification across the 11 assets at the same time.

EEG application. The EEG was the second application domain used in the Semester 1 design. It can be fully applied to multichannel neural recordings and the analysis (sliding-window GC/TE/PSI/coherence) would confirm the applicability of the technique to real EEG data, not just to financial data.

REFERENCES

- [1] R. W. Thatcher, E. Palmero-Soler, D. M. North, and C. J. Biver, "Intelligence and EEG measures of information flow: efficiency and homeostatic neuroplasticity," *Scientific Reports*, vol. 6, p. 38890, 2016.
- [2] A. W. Lo and A. C. MacKinlay, "When are contrarian profits due to stock market overreaction?" *The Review of Financial Studies*, vol. 3, no. 2, pp. 175–205, 1990.
- [3] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969.
- [4] T. Schreiber, "Measuring information transfer," *Physical Review Letters*, vol. 85, pp. 461–464, 2000.
- [5] G. Nolte, A. Ziehe, V. V. Nikulin, A. Schlögl, N. Krämer, T. Brismar, and K.-R. Müller, "Robustly estimating the flow direction of information in complex physical systems," *Physical Review Letters*, vol. 100, p. 234101, 2008.
- [6] F. X. Diebold and K. Yilmaz, "On the network topology of variance decompositions: Measuring the connectedness of financial firms," *Journal of Econometrics*, vol. 182, no. 1, pp. 119–134, 2014.
- [7] M. Billio, M. Getmansky, A. W. Lo, and L. Pelizzon, "Econometric measures of connectedness and systemic risk in the finance and insurance sectors," *Journal of Financial Economics*, vol. 104, no. 3, pp. 535–559, 2012.
- [8] R. Tibshirani, "Regression shrinkage and selection via the lasso," *Journal of the Royal Statistical Society: Series B*, vol. 58, no. 1, pp. 267–288, 1996.
- [9] Y. Amihud, "Illiquidity and stock returns: cross-section and time-series effects," *Journal of Financial Markets*, vol. 5, no. 1, pp. 31–56, 2002.
- [10] M. Parkinson, "The extreme value method for estimating the variance of the rate of return," *Journal of Business*, vol. 53, no. 1, pp. 61–65, 1980.

[1] R. W. Thatcher, E. Palmero-Soler, D. M. North, and C. J. Biver, "Intelligence and EEG measures of information flow: efficiency and