White paper

Real-time flow-based predictions of SEC Form 13F trades:

Accuracy measurements of LSEG's real-time Trading Flow data and subsequent 13F flows

LSEG DATA & ANALYTICS

EXPONENTIAL

Contents

Contents 1
Summary 2
Research Motivation 2
Methodology Details 3
Results 5
Conclusion 8

Summary

Trading Flow provides a minute-by-minute, cross-asset view of trading activity, segmented by investor type, and in some cases, trading strategy. Powered by LSEG's premium data and Exponential Technology's (XTech) expertise in advance analytics and trading behaviour, it decodes real-time order book information to quantify cumulative net buying and selling activity for both institutional and retail investors. By using **Trading Flow,** in this white paper, we explore how researchers classified every trade in real-time, applying advanced inference algorithms to distinguish institutional from retail activity.

It is shown that for the S&P 500 universe, using this cumulative net institutional flow metric one can predict the direction of the aggregate purchase/sales of stock reported on future SEC Form 13F filings with a 65.5% directional accuracy (DA) (86% confidence) and 34.5% cross sectional information coefficient (IC) (88% confidence). Then, isolating only those S&P500 stocks that exhibit a 95% confidence level in DA and IC the sub-sample of 222 and 289 stocks respectively demonstrate a 71.1% DA (>99% confidence) and 45.1% IC (>99% confidence). The same comparisons were tested across the market cap spectrum and sectors and demonstrated robust and consistent performance in predicting future SEC Form13F aggregate net investment flows. Remarkably, Energy, Communications Services, Consumer Stables and Materials were the top sectors as measured by DA and IC. The cumulative retail flow shows no statistical relationship to the future SEC Form13F filings with 48.8% DA (54% confidence) and 0.4% IC (59% confidence). The study confirms with a "ground-truth" SEC Form 13F dataset that **Trading Flow** is effective at decoding orderbook dynamics in real-time into institutional and retail flow.

Research motivation

Modern equity markets churn through billions of shares every day, yet the consolidated tape tells us nothing about who is behind each trade. Portfolio managers, stock analysts, agency execution desks, proprietary traders and risk managers are therefore forced to speculate about institutional rotation, retail exuberance, or market-maker inventory changes from price and volume alone, blurring the very forces that drive stock outperformance, market regime shifts and liquidity shocks.

Regulatory filings do provide transparency, but at a glacial pace: SEC Form 13F holdings arrive as much as 45 days after quarter-end, which could be 45-165 days after a trade occurs--long after markets have digested the information. For alpha signal engineers and risk managers, waiting that long is tantamount to learning last season's playbook during a live game.

Trading Flow narrows this gap by analysing every trade in real-time and inferring if it belongs to an institutional trader, market-maker, or retail trader. Those tagged trades are then cumulated into 1-minute buckets and can power short term momentum alerts and flow-based momentum intraday signals. When those same flows are cumulated over days and weeks they can power forecasting over investment horizons of days to years. Cross-sectional analysis of this data can provide colour on secular trends and act as an early warning of sector rotation trades, index rebalancing and asset allocation trades. This is an exciting proposition, but we need to understand how accurate the XTech inference algorithms are at measuring aggregate flow by investor type by comparing it to a correctly labelled, so-called "ground truth" dataset. One such "ground truth" dataset is the SEC Form 13F filings which taken together represent a significant portion of actual institutional flow in the prior quarter. With the 13F dataset we can evaluate how closely the overall Trading Flow inference predicts actual reported aggregate flow of large institutional investors. This approach provides a statistical estimate of the accuracy of the XTech real-time inference algorithms at the level of the largest trades. If there is a strong alignment between Trading Flow and 13F changes, then it supports the efficacy of the inference algorithm and also suggests it is possible to create a myriad of strategies related to trading along with these large investors, capturing reversion signals at the conclusion of their trades and even predicting the very public release of the SEC Form 13F itself. Below we present a straightforward framework to compare XTech's institutional flow with 13F changes and quantify its accuracy.

Methodology details

2.1 "Ground Truth" – Form 13F filings

Form 13F filings are treated as "ground truth," but in reality, they offer only an imperfect snapshot of institutional holdings. Institutional investors under $100M are excluded entirely, and even for large managers the reported market value is marked to quarter-end closes, not the actual execution prices. Trading Flow, by contrast, classifies all institutional trades in real-time, based on the prices and quantities at which transactions occurred.

The result is that 13F and Trading Flow will never line up perfectly in scale or timing. Still, despite its blind spots, 13F remains the accepted benchmark for measuring institutional activity. To anchor this benchmark, we constructed quarterly SEC Form 13F trade value snapshots for each S&P 500 constituent and defined the series:

(i) $\Delta V_{i,q}^{13F}$: the quarter-over-quarter change in the reported *market value* of security *i* held by all 13F filers during quarter q.

2.2 "Forecast" – XTech real-time institutional flow

If 13F filings are the slow-moving, delayed benchmark, Trading Flow is the fast-moving, real-time forecast. It uses proprietary inference algorithms to classify every trade in real-time as institutional, retail, or market-making, and aggregates the results into minute-level granularity. Alongside flow values, the system also records trade counts, though these are not used in the present study.

To compare with 13F, which is quarterly by construction, we collapse this rich real-time feed into quarterly aggregates and treat it as the "forecast" series for the benchmark series $\Delta V_{i,q}^{13F}$. For implementation, we drew XTech data directly from XTech's Unifier API.

At the daily level, institutional net flow for stock *i* on trading day *d* is defined as:

$$\Delta \boldsymbol{v}_{i,d} = \boldsymbol{buy}_{i,d}^{(inst)} - \boldsymbol{sell}_{i,d}^{(inst)} \tag{1}$$

where $buy_{i,d}^{(inst)}$ and $sell_{i,d}^{(inst)}$ are the total market value classified as institutional buys and sells on trading day *d*. To bring this to the quarterly horizon, we sum across trading days in quarter *q*:

$$\Delta \boldsymbol{V}_{i,q}^{xtech} = \sum_{d \in \mathcal{D}_q} \Delta \boldsymbol{v}_{i,d} \tag{2}$$

We adopt standard quarter-end dates - March 31, June 30, September 30, and December 31 - and restrict coverage to current S&P 500 constituents with one-to-one ticker alignment across datasets. The result is a quarterly institutional flow forecast that can be directly compared to 13F filings, enabling a statistical evaluation of the real-time inference against the delayed "ground truth".

The forecast $\Delta V_{i,q}^{xtech}$ captures linear relationships and directional accuracy with the benchmark $\Delta V_{i,q}^{13F}$ reasonably well. However, they cannot be used directly to measure magnitude errors without adjustment. Misclassification inevitably arises when tagging buys and sells, and these errors systematically bias the estimated flow downward relative to the true underlying flow.

This calibration step ensures that the adjusted signal $\Delta V'^{xtech}_{i,q}$ has approximately the same mean and variance as the observed 13F changes $\Delta V_{i,q}^{13F}$ historically. In effect, we place the XTech forecast and the 13F benchmark on the same statistical footing, allowing for a cleaner comparison of predictive accuracy.

2.3 Evaluation against 13F

To evaluate how well Trading Flow approximates institutional activity, we evaluate the calibrated signal $\hat{y}_{i,q} = \Delta V'^{xtech}_{i,q}$ against the reported 13F changes $y_{i,q} = \Delta V^{13F}_{i,q}$ for each stock *i* and quarter $q = 1, \ldots, n$. Accuracy is assessed along several complementary dimensions, each highlighting a different aspect of predictive alignment.

(i) Pearson correlation (linear association).
Pearson correlation measures the linear association between the calibrated flow signal and 13F holding changes.

$$\boldsymbol{\rho_i} = \frac{\sum_{q=1}^{n}(y_{i,q}-\bar{y}_i)(\hat{y}_{i,q}-\bar{\hat{y}}_i)}{\sqrt{\sum_{q=1}^{n}(y_{i,q}-\bar{y}_i)^2}\sqrt{\sum_{q=1}^{n}(\hat{y}_{i,q}-\bar{\hat{y}}_i)^2}} \quad \textbf{(5)}$$

(ii) Directional accuracy (hit rate).
Markets often care less about exact magnitudes and more about whether the signal gets the direction right. The hit rate captures the proportion of quarters in which the forecasted and reported flows share the same sign:

$$\boldsymbol{hit_i} = \frac{1}{n}\sum_{q=1}^{n}\mathbb{I}\{\boldsymbol{Sign}(\boldsymbol{y_{i,q}}) = \boldsymbol{sign}(\boldsymbol{\hat{y}_{i,q}})\} \quad \textbf{(6)}$$

(iii) Error magnitude (RMSE).
To quantify the typical size of forecast errors in magnitude, we compute the root mean squared error:

$$\boldsymbol{RMSE_i} = \sqrt{\frac{1}{n}\sum_{q=1}^{n}(\boldsymbol{\hat{y}_{i,q}} - \boldsymbol{y_{i,q}})^2} \quad \textbf{(7)}$$

(iv) Proportional deviation (sMAPE).
Because scale varies widely across stocks, we also measure proportional error using the symmetric mean absolute percentage error, which is both scale-free and robust to sign:

$$\boldsymbol{sMAPE_i} = \frac{1}{n}\sum_{q=1}^{n}\frac{2|\hat{y}_{i,q}-y_{i,q}|}{|\hat{y}_{i,q}|+|y_{i,q}|} \quad \textbf{(8)}$$

For Pearson correlation ρ_i and hit rate hit_i, we compute t-statistics and p-values to test against the null of random guessing. Under this null, the expected correlation is zero and the hit rate is 50%.

By contrast, $RMSE_i$ and $sMAPE_i$ measure raw error magnitudes and lack a well-defined null distribution. Instead, following common practice in financial forecast evaluation, we report performance relative to simple benchmark models:

(i) Theil U_1 — RMSE relative to a random walk forecast

$$\boldsymbol{U_{1,i}} = \frac{\boldsymbol{RMSE_i(forecast)}}{\boldsymbol{RMSE_i(random)}} \quad \textbf{(9)}$$

Values $U_1 < 1$ indicate improvement over the random walk baseline.

(ii) Theil U_2 — RMSE relative to the naïve “previous value” forecast:

$$U_{2,i} = \frac{RMSE_i(forecast)}{RMSE_i(prev)} \tag{10}$$

Here again, values $U_2 < 1$ indicate superiority to the naïve benchmark.

(iii) Relative sMAPE vs. zero forecast:

$$r^0_{sMAPE_i} = \frac{sMAPE_i(forecast)}{sMAPE_i(zero)} \tag{11}$$

(iv) Relative sMAPE vs. previous value forecast:

$$r^{prev}_{sMAPE_i} = \frac{sMAPE_i(forecast)}{sMAPE_i(prev)} \tag{12}$$

Results

Figure 3.1—Predictive Performance by Sector (Average Pearson Correlation & Hit Rate)

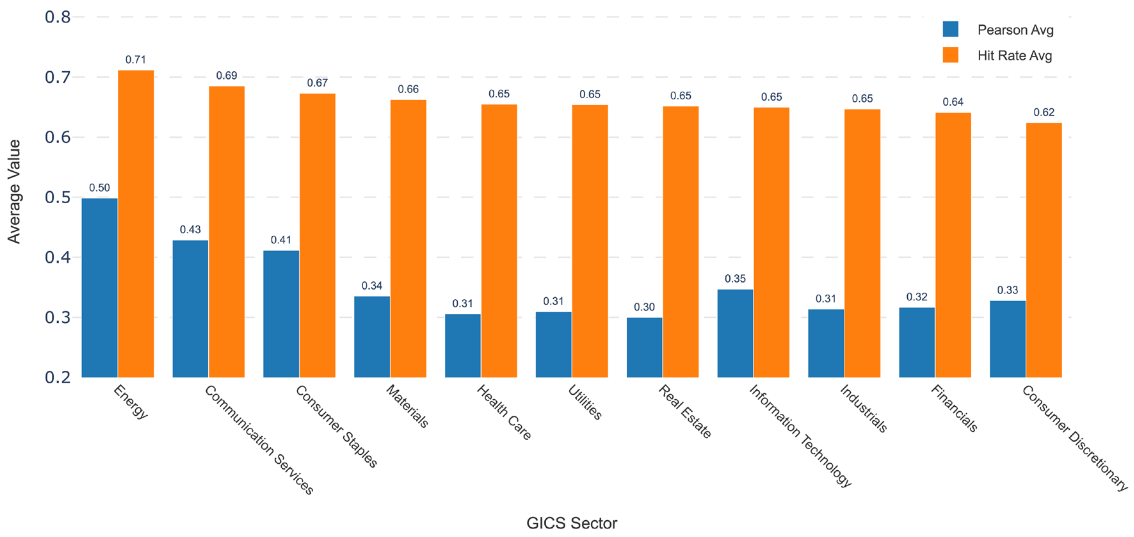

In Table 3.1, the cross-sectional averages are unconditional and therefore conservative, since they include names where the classifier carries little signal. Restricting the universe to stocks that are statistically significant at the 5% level (95% confidence) yields materially higher performance: the average hit rate increases to **0.711** (mean p-value **0.0090 < 1%),** and the average Pearson correlation rises to **0.451** (mean p-value **0.0079 < 1%**). Importantly, this 5% significance screen still covers a large share of the universe: 222 names are significant on hit rate (48.79%) and 289 on Pearson (64.22%) as reported in Table 3.2. This conditioning confirms that the signal delivers strong and statistically reliable alignment with 13F changes.

Figure 3.1 reports average Pearson correlation (ρ) and hit rate (h) by GICS sector. Directional accuracy is robust, with h exceeding 0.60 in nearly every sector and clustering between 0.62 and 0.71, while linear association remains broadly consistent with ρ in the 0.30–0.40 range. Energy (h = 0.71, ρ = 0.46), Communication Services (h = 0.64, ρ = 0.38), and Consumer Staples (h = 0.67, ρ = 0.35) exhibit the strongest forecast performance compared to other sectors. The observed heterogeneity likely reflects structural differences in institutional participation, index-related and ETF-related rebalancing activity, and the frequency of corporate actions. Importantly, despite these sources of cross-sector variation, the institutional signal consistently provides statistically meaningful directional information across the entire sectoral landscape for future 13F filings.

Figure 3.2 - Relationship between Log Market Capitalisation and Predictive Performance (Pearson & Hit Rate)

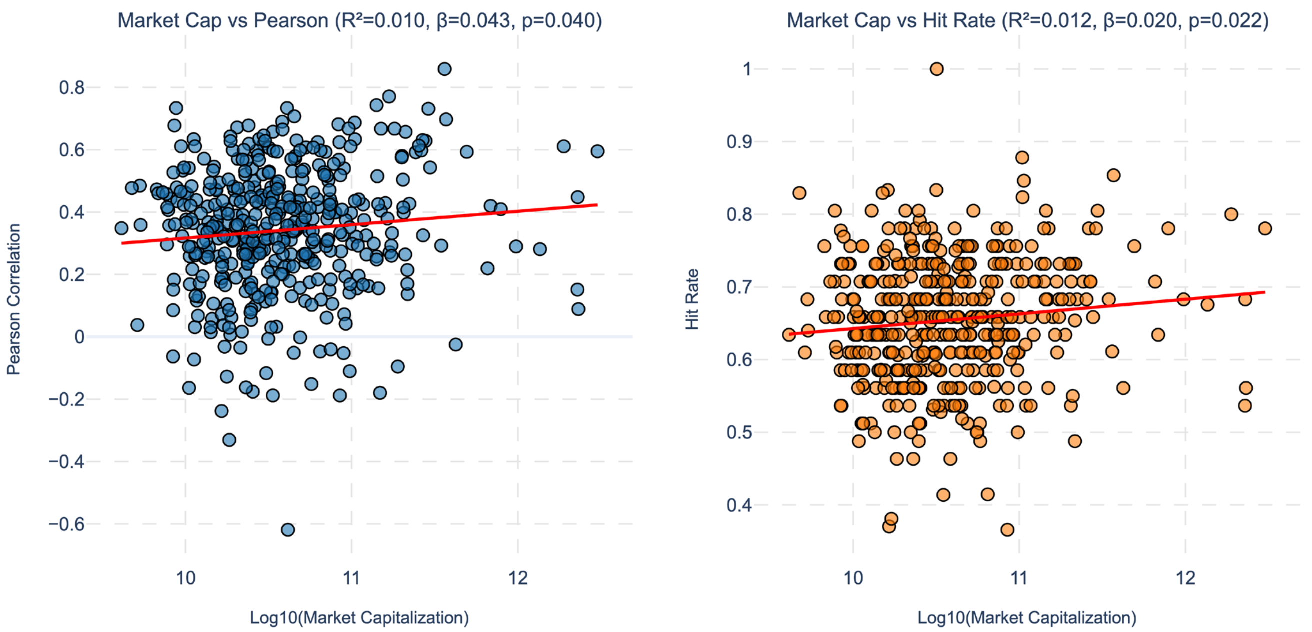

Figure 3.3—Quarterly 13F vs. Trading Flow (Raw, Calibrated and Real-time): Ticker-PG

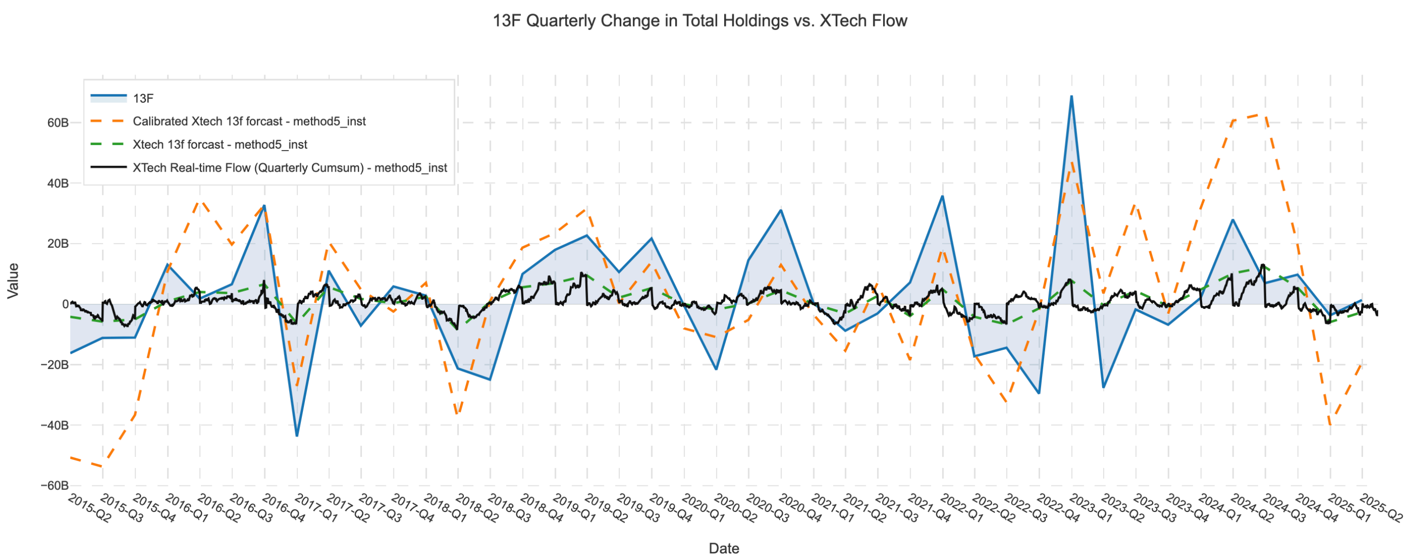

We further assess whether predictive accuracy varies systematically with firm size. Figure 3.2 relates forecast performance to log market capitalisation by plotting sector-level observations alongside fitted linear regressions. Both correlations are weak but statistically significant: larger firms exhibit slightly higher Pearson correlations ($\rho = 0.10$, $p = 0.040$) and modestly stronger hit rates ($\rho = 0.12$, $p = 0.022$). These results suggest that the institutional signal is preserved across the capitalisation spectrum, but alignment with 13F benchmarks is somewhat stronger in large-cap stocks. This size tilt is consistent with the higher liquidity, greater indexation, and heavier institutional ownership typical of larger firms, which may facilitate cleaner classification and stronger predictive linkage. Importantly, however, the presence of statistically meaningful signal across all size deciles indicates that the methodology captures institutional activity broadly, rather than being narrowly concentrated in either mega-caps or small-caps.

Figure 3.4—Quarterly 13F vs. Trading Flow (Raw, Calibrated and Real-time): Ticker-OXY

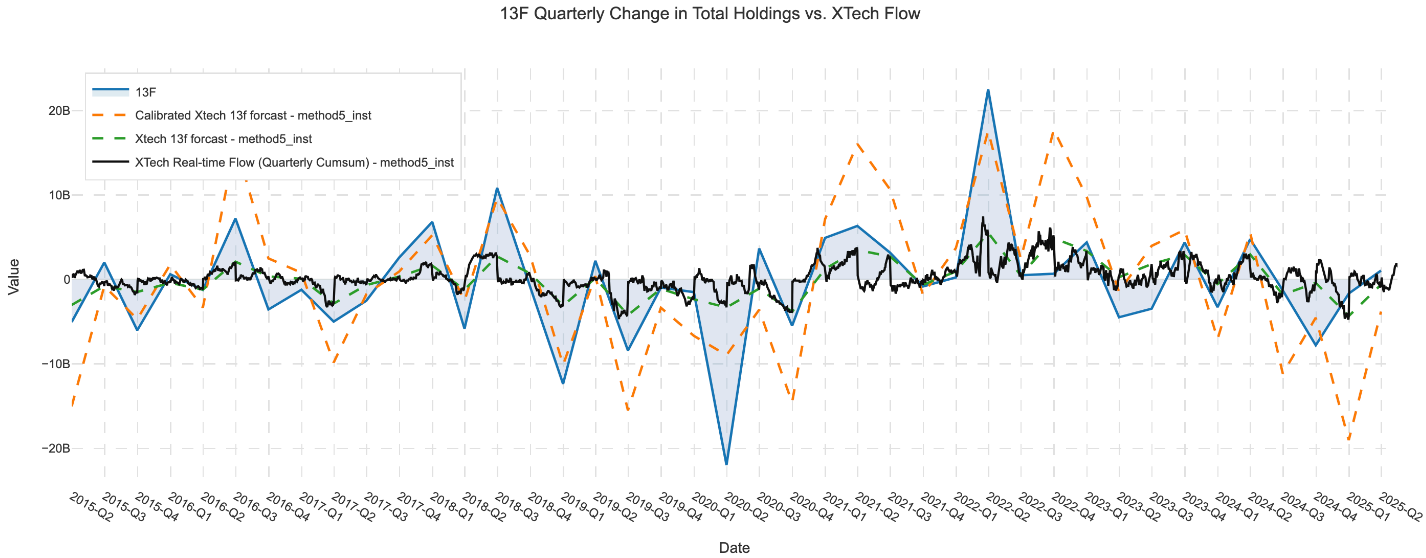

Figure 3.3 and 3.4 overlays quarterly 13F changes with Trading Flow, shown both as the raw within-quarter cumulative series and as the calibrated series. The calibrated series rescales the forecast to match the 13F mean and variance while preserving the real-time turning points, thereby improving sign coherence and making magnitudes directly comparable. In contrast, the real-time “quarterly cumsum” resets at the start of each quarter and accumulates daily institutional flow, with peaks and troughs that typically precede the 13F print—evidence that the classifier is capturing the build-up and unwind of institutional activity as it happens. Together, the two perspectives provide both a real-time and a statistically aligned view of institutional positioning.

This dual lens has several important implications: the cumulative series can serve as a nowcast of quarter-to-date institutional demand, trading strategies can exploit flow inflections to trade with or fade institutional activity and risk managers can monitor crowding and liquidity exposure by tracking sustained one-sided flows. Finally, the calibrated series enables direct comparability to 13F filings, making it possible not only to validate the inference but also to forecast the likely sign and scale of public disclosures ahead of time.

Conclusion

This study delivers the first large-scale benchmark that rigorously validates real-time investor flow signals against the definitive ground truth of SEC Form 13F filings. Leveraging a decade of S&P 500 history, we demonstrate that **Trading Flow** – specifically its institutional classifier – accurately decodes order book dynamics into cumulative institutional activity, achieving directional accuracy as high as **71.1%** with a cross-sectional information coefficient of **45.1%**, **both at >99% confidence levels**. These results decisively exceed random chance, confirm robustness across sectors and market caps, and, most importantly, establish real-time institutional trading flow as a measurable structural force rather than a mysterious idiosyncratic artifact.

By contrast, as expected, retail flow shows no predictive value (48.8% DA, 0.4% IC), underscoring the precision of XTech's classification in isolating institutional behaviour from retail flow. Sector-level analysis further highlights Energy, Communication Services, Consumer Staples, and Materials as exhibiting the strongest predictive signals, providing additional practical levers for strategy design in those sectors.

Crucially, XTech's framework bridges the temporal blind spot of regulatory reporting: what 13F filings reveal only with multi-month delay can now be inferred in near real-time from order book dynamics. This capability transforms institutional flows from a retrospective disclosure into a forward-looking signal, enabling applications in quantitative trading, tactical sector allocation, and flow-aware execution and risk management.

By grounding real-time inference in verified institutional outcomes, this work establishes a new standard for flow-based analytics — one that enhances market transparency and equips investors with a timely, statistically validated edge.