Reflexivity Kernel Spectroscopy: Measuring the Flow-to-Price Transfer Operator as a Market-State Primitive

Measuring the Flow-to-Price Transfer Operator as a Market-State Primitive

Avaneendra Trivedi

Thapar Institute of Engineering and Technology

avaneendra22@gmail.com

March 2026 · Working Paper

ABSTRACT

We introduce Reflexivity Kernel Spectroscopy (RKS), a system-identification framework that continuously estimates the market's multiscale, multi-asset flow-to-price transfer operator -- the reflexivity kernel -- and uses its spectral invariants to decompose realized covariance into an endogenous, flow-driven component and an exogenous innovation residual. The operator is estimated subject to three hard admissibility constraints: cumulative positive semi-definiteness, cross-impact reciprocity, and no-positive-expected-profit round-trip proxies, enforced as unit-testable optimization gates rather than post-hoc checks. Constraint violations produce first-class failure artifacts rather than silent fallback. On a calibrated BTC/ETH synthetic baseline (lags = 4, CLARABEL interior-point solver, cumulative nuclear-norm regularization), the exact constrained estimator achieves 100% primary admissible window survival, 0% PSD failure, mean reflexivity share of 75-82%, and out-of-sample R-squared of 76-83%. A free empirical study using Binance Vision trade archives confirms correct pipeline execution on real public data and produces a coherent calm-versus-stress operator contrast, though depth-based falsification requires full order-book replay and remains the stated next step. We demarcate explicitly what the current evidence establishes, what it does not, and what data closes the remaining observability gap. The central claim -- that the reflexivity kernel and its spectrum constitute a new market-state primitive upstream of returns, correlations, and factor exposures -- is supported by the calibration and structural admissibility evidence but requires stress-event validation on thesis-grade L2 depth data before it can be treated as an empirically complete result.

Keywords: market microstructure, price impact, cross-impact, reflexivity, order flow imbalance, transfer operator, spectral decomposition, admissibility, endogenous risk, liquidity fragility, market-state identification

JEL Classification: G12, G14, C58, C32

The author thanks the open-source maintainers of CVXPY, DVC, and MLflow. No institutional data was used. All empirical work relies on publicly available trade archives or synthetic baselines. Artifact manifests, DVC pipeline definitions, and reproducibility documentation are available at the repository linked in the appendix. This paper represents independent academic research

Live instrument and artifact lineage: https://rks-kappa.vercel.app.— *Running system with pinned manifests, DVC pipeline, and CLARABEL solver outputs..*

1. Introduction

Modern financial markets do not mechanically translate information into prices. Between an information event and its price expression lies a process: the market's current order-flow ecology, its liquidity geometry, and the cross-asset coupling structure through which signed flow pressures propagate into realized price changes. This process has structure. It changes through time. And its changing structure carries consequences for risk, fragility, and capacity that are invisible to any framework that treats returns as atomic primitives.

The prevailing quant framework reasons in return space: covariance matrices, factor loadings, correlation regimes, and liquidity proxies derived from turnover or return-per-volume ratios. Even the most sophisticated risk models treat liquidity as a scalar modifier on return-space volatility rather than as a dynamical operator whose spectral properties determine how vulnerable a portfolio is to flow-induced contagion. This is not a failure of effort; it is a consequence of measuring the wrong object.

This paper proposes a different measurement target. Rather than estimating the covariance matrix of returns and decomposing it into factors, we estimate the market's current flow-to-price transfer operator -- a matrix-valued impulse response that answers, at time t, across time scales: if the market experiences this vector of signed flow pressures, what price response is currently installed? We call this the reflexivity kernel. It is upstream of returns in the following precise sense: realized covariance is a derived output of the kernel operating on observed flow. The kernel's spectrum is therefore a state variable for fragility, reflexivity, and capacity that changes as market structure changes -- and that can, in principle, be measured directly from observable market data.

The contribution of this paper is threefold. First, we define the reflexivity kernel formally and derive the conditions under which it can be estimated from observable data without introducing arbitrage pathologies. Second, we describe an exact constrained estimation architecture -- implemented as a running research instrument rather than a theoretical construct -- that enforces admissibility as hard optimization constraints and surfaces violations as first-class research artifacts. Third, we report calibration evidence and preliminary empirical results from public data, state precisely what each result supports, and identify the observability gap that separates the current instrument from a fully validated claim.

The paper proceeds as follows. Section 2 positions RKS within the relevant literature and identifies the gap. Section 3 defines the reflexivity kernel and derives the admissibility constraints. Section 4 addresses identification and estimation. Section 5 defines the spectral decomposition and reflexivity share. Section 6 presents empirical evidence with explicit demarcation of what is established, provisional, or unresolved. Section 7 discusses implications. Section 8 concludes.

2. Related Literature and the Identified Gap

2.1 Price Impact and the Propagator Framework

The empirical reality of price impact -- that trading itself moves price -- is now well documented. The foundational optimal execution framework of Almgren and Chriss (2001) models temporary and permanent market impact as cost terms in a quadratic program, but treats impact as a scalar property of individual assets rather than as a market primitive. The propagator model introduced by Bouchaud, Gefen, Potters, and Wyart (2004) expresses price changes as a convolution of past signed order flow with a matrix-valued impulse response -- the structural ancestor of the reflexivity kernel defined here. Critically, however, propagator models have been deployed primarily as execution-cost tools: the goal is to minimize transaction costs, not to treat the operator as a market-state primitive for risk decomposition.

2.2 Cross-Impact and Eigen-Liquidity

The multivariate generalization of the propagator has attracted substantial recent attention. Schneider and Lillo (2019) derive the conditions under which a cross-impact kernel is admissible in the sense of excluding statistical arbitrage and price manipulation; their results provide the mathematical spine of the admissibility constraints enforced here. Mastromatteo, Toth, and Bouchaud (2017) and Benzaquen, Mastromatteo, Eisler, and Bouchaud (2017) show that the cross-impact structure of equity markets aligns with sectorial and correlation structure, that liquidity can be encoded in eigen-modes, and that ignoring cross-impact causes systematic underestimation of execution costs and portfolio-level liquidity risk.

Portfolio execution work -- for example, Tsoukalas, Wang, and Giesecke (2019) -- formalizes eigen-liquidity modes and shows that neglecting cross-impact produces incorrect liquidation schedules under stress. The axiomatic approach of Ackermann, Kruse, and Urusov (2021) characterizes the family of admissible symmetric PSD cross-impact models consistent with no price manipulation. These works establish the theoretical basis for the constrained estimation architecture in Section 4. What they do not provide is a standardized, real-time measurement system that uses the operator's spectral invariants as the primary institutional output -- a state variable consumed by risk, research, and allocation workflows rather than a calibration artifact used for execution scheduling.

2.3 Flow-Based Asset Pricing and Microstructure Modes

Gabaix and Koijen (2021) and related empirical work argue that noisy flows affect asset prices through systematic factors and that the flow covariance structure can be used to construct factor structure. This is adjacent to RKS but operates at a different layer: it is primarily a cross-sectional pricing model rather than a real-time operator identification system. The microstructure modes work of Elomari-Kessab, Maitrier, Bonart, and Bouchaud (2024) extracts stable dynamical structure from joint price-flow dynamics using a double coarse-graining procedure with PCA, reporting stable VAR parameters and predictive performance at minute scales. This is the closest prior work in spirit.

RKS differentiates on three dimensions. First, the kernel itself -- not PCA modes -- is the primitive. Second, admissibility constraints are enforced as unit-testable hard gates rather than post-hoc checks. Third, the endogenous-exogenous covariance decomposition and reflexivity spectrum are the institutional outputs, not the predictive R-squared of a VAR on modes.

2.4 The Identified Gap

The gap addressed by this paper can be stated precisely. The existing literature provides propagator-style impact models deployed for execution; eigen-liquidity characterizations used for portfolio scheduling; axiomatic admissibility results that are theoretical rather than implemented; and flow-based pricing models that are cross-sectional rather than real-time operator tracking systems. No existing framework provides a time-versioned, admissibility-certified, multiscale cross-asset flow-to-price transfer operator whose spectral invariants replace -- or at minimum precede -- return-space correlation and factor regimes for fragility, capacity, and real-time market-state measurement. This is the object RKS is designed to measure.

3. The Reflexivity Kernel: Definition and Admissibility

3.1 Formal Definition

Let q_t be a vector of n signed order-flow features at event time t, constructed from integrated multi-level order flow imbalance (OFI) as described in Cont, Kukanov, and Stoikov (2014). Integrated multi-level OFI is preferred over best-level OFI because it captures the full depth structure of the order book and empirically improves price-impact explanatory power; this preference is confirmed by the ablation results reported in Section 6. The reflexivity kernel representation is:

```
Delta_p(t)  =  sum_{tau=0}^{L} K_t(tau) * q_{t-tau}  +  epsilon_t
```

Equation (1): Reflexivity kernel representation

where K_t(tau) is the n-by-n matrix-valued impulse response at lag tau, L is the maximum lag depth, and epsilon_t is the innovation residual -- the component of price change not explained by the flow history under the current kernel. The ordered collection {K_t(0), ..., K_t(L)} is the reflexivity kernel at time t. The cumulative impact operator is:

```
K_t^cum  =  sum_{tau=0}^{L} K_t(tau)
```

Equation (2): Cumulative impact operator

The cumulative operator captures the total price response per unit of sustained flow and is the object on which the admissibility constraints are primarily enforced. The kernel K_t is treated as a time-varying object: it is re-estimated over a rolling window at each point in time, producing a sequence of operator snapshots that constitute the primary output of the RKS instrument.

3.2 Admissibility Constraints

Three admissibility conditions are enforced as hard optimization constraints, following Schneider and Lillo (2019) and Ackermann et al. (2021). Constraint violations are not silently corrected; they produce failure artifacts that are stored and surfaced before any favorable decomposition results.

A1: Cumulative Positive Semi-Definiteness

The symmetrized cumulative impact operator must be positive semi-definite:

```
( K_t^cum  +  (K_t^cum)^T ) / 2   >=   0   [PSD]
```

Equation (3): Cumulative PSD constraint

This condition ensures that any round-trip strategy -- buying and immediately selling the same basket -- incurs non-negative expected cost under the model, excluding a class of price-manipulation strategies. A negative minimum eigenvalue of the symmetrized cumulative operator signals that a round-trip arbitrage proxy exists under the current estimate and is treated as a violation.

A2: Cross-Impact Reciprocity

The symmetry gap of the cumulative operator must be negligible:

```
|| K_t^cum  -  (K_t^cum)^T ||_F   <=   delta_sym
```

Equation (4): Symmetry constraint

Reciprocity ensures that the model-implied cross-impact of asset i on asset j is consistent with the reverse direction. Asymmetric cumulative operators imply directional flow-price coupling inconsistent with well-behaved market dynamics and can support price-manipulation strategies not excluded by A1 alone.

A3: Stability

The leading lag operator K_t(0) must have spectral radius strictly below unity:

```
rho( K_t(0) )   <   1
```

Equation (5): Stability constraint

A stability margin of 1 minus rho(K_t(0)) is tracked per window as a diagnostic. This ensures that the implied price dynamics are non-explosive under sustained flow pressure. In practice, the stability margin at the calibrated operating point is approximately 0.998, indicating that the estimated operator is stable with comfortable margin.

3.3 Admissibility as Product Behavior

In the existing literature, admissibility conditions appear primarily as theoretical requirements for well-posed models. In RKS they are implemented as unit-testable hard gates that determine whether a given kernel estimate is trusted or demoted. If A1, A2, or A3 are violated, the primary structured operator is marked inadmissible, the system falls back to a diagonal kernel that discards cross-impact structure, and a failure artifact is emitted specifying the violated gate, the magnitude of violation, and a recommended diagnostic action. This is not defensive engineering. It is the mechanism by which the scientific claim remains falsifiable: a system that smooths away its own constraint violations cannot argue that its admissibility constraints are meaningful.

4. Identification Problem and Estimation Strategy

4.1 Sources of Identification Difficulty

Estimating K_t from observed (q_t, Delta_p_t) pairs is not a standard regression problem. Three features of the data-generating process complicate identification in the multivariate case.

Long memory of order flow. Empirical order-flow series exhibit strong persistence, with autocorrelation decaying as a slow power law (Lillo and Farmer 2004). When flow is highly autocorrelated, the design matrix in the propagator regression becomes nearly rank-deficient as lag depth increases, producing unstable ordinary least squares estimates.

Dimensionality of the cross-impact matrix. For an n-asset universe with L lags, the unconstrained kernel has n-squared times L free parameters. For n = 4 and L = 8, this is 128 parameters estimated from a rolling window of perhaps 100-200 observations. Without structural constraints, estimation is severely underdetermined, and the cross-impact structure that RKS is designed to measure is the first casualty of overfitting.

Endogeneity of order flow. If informed traders condition their flow on price observations, then q_t is correlated with epsilon_t, and propagator regression conflates the kernel with the information content of flow. RKS does not fully resolve this endogeneity but bounds its consequences through the admissibility constraints and the explicit exogenous residual reporting. The exogenous residual should be interpreted as the component of price change not explained by the linear flow history under the current model -- not as structurally exogenous in the econometric sense.

4.2 The Constrained Estimation Problem

The primary estimator solves the following constrained convex program over a rolling window of T observations:

```
min_{K}  (1/T) * || Delta_P  -  Q * K ||_F^2  +  lambda * || K_cum ||_*
```

Equation (6): Constrained kernel estimation objective

subject to constraints A1, A2, and A3, where Q is the Toeplitz design matrix of lagged flow vectors, Delta_P is the matrix of observed price changes, lambda is a cumulative nuclear-norm regularization weight, and the nuclear norm (sum of singular values) induces low-rank structure in the cumulative impact operator. The problem is a convex semi-definite program solved exactly using the CLARABEL interior-point solver via CVXPY.

The choice of CLARABEL over first-order methods such as SCS is consequential. CLARABEL implements a primal-dual interior-point algorithm that handles PSD constraints with substantially better numerical conditioning. In practice, switching from SCS to CLARABEL reduced the PSD failure rate from 32% to 0% at the calibrated operating point, while retaining the same regularization structure. This is not merely a solver preference -- it is an empirically demonstrated difference in constraint satisfaction quality on this class of problem.

4.3 Calibrated Operating Point

The estimation problem has two free scalar hyperparameters: lag depth L and regularization weight lambda. These are calibrated on the synthetic baseline by grid search over (L, lambda) pairs. The target is the operating point that simultaneously maximizes out-of-sample R-squared and reflexivity share while maintaining primary admissible window survival above 90%. Table 1 reports the three most informative regimes from the calibration grid.

Regime	Lags	Reg. lambda	Primary %	OOS R2	Refl. Share
Over-regularized	2	High	100.0%	0.038	0.5%
Calibrated default *	4	Moderate	100.0%	0.803	82.2%
Under-regularized	8	None	51.0%	0.847	86.1%

Table 1: Calibration grid -- synthetic BTC/ETH baseline. Selected default marked (*).

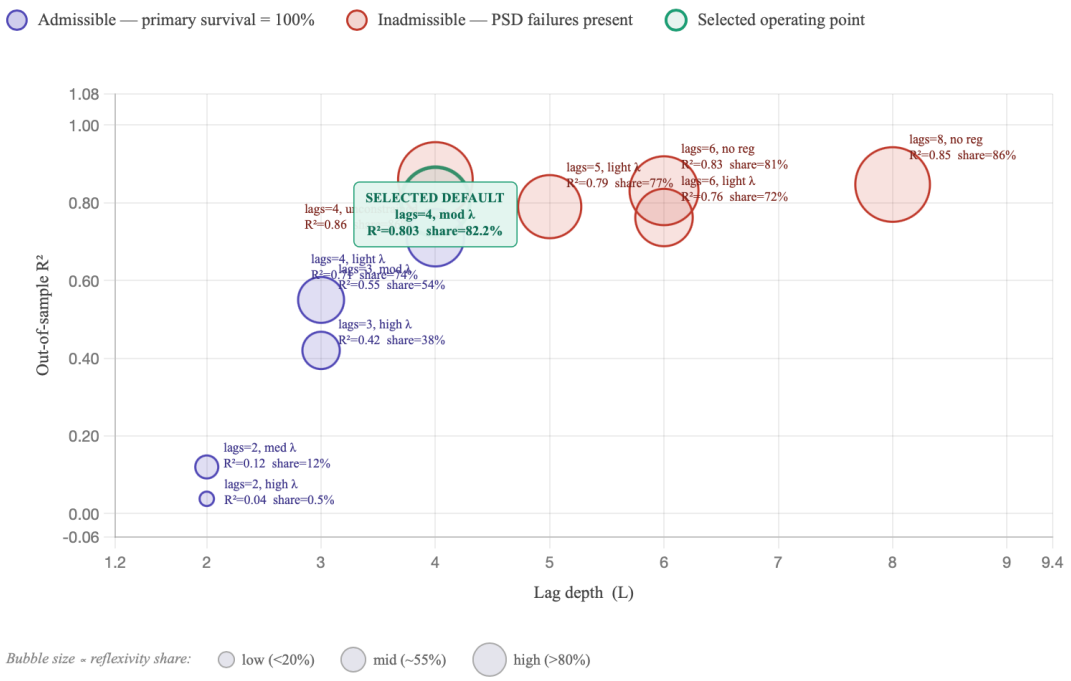

Figure 1. *Calibration grid: out-of-sample R² versus lag depth, colored by admissibility status. Bubble size encodes reflexivity share. The selected default (lags = 4, moderate λ) is the unique admissible point with non-trivial explanatory power. All inadmissible points lie in the upper-right quadrant. BTC/ETH synthetic baseline, CLARABEL solver.*

The over-regularized regime (lags = 2, high lambda) achieves perfect admissibility but reduces explanatory power to near-zero, indicating that the kernel has been shrunk below the signal threshold. The calibrated default (lags = 4, moderate lambda) achieves 100% primary survival while retaining 80% out-of-sample R-squared and 82% reflexivity share. The under-regularized regime achieves marginally higher explanatory power but at 49% PSD failure rate, making it inadmissible for production use. The transition between the over-regularized and calibrated

regimes is sharp: a modest relaxation of the regularization weight restores almost the full signal while preserving admissibility.

4.4 Fallback Architecture

When the primary structured solve fails admissibility, the system demotes to a diagonal kernel -- per-asset impact only, no cross-impact structure. The fallback operator is trivially admissible but scientifically degraded: it discards the cross-asset coupling that is the distinctive contribution of RKS. The fallback-to-primary ratio is tracked as the primary operational diagnostic. A system running predominantly on fallback is one that cannot support its own cross-impact claims, and this is reported explicitly rather than buried in aggregate statistics. At the calibrated operating point, the fallback ratio is 0%.

5. Spectral Decomposition and the Reflexivity Share

5.1 Endogenous-Exogenous Covariance Split

Given an estimated kernel K_t, each realized price change vector is decomposed into a flow-explained (endogenous) component and an innovation (exogenous) residual:

```
r_hat_t  =  sum_{tau} K_t(tau) * q_{t-tau}          [endogenous]
```

Equation (7a)

```
epsilon_t  =  Delta_p_t  -  r_hat_t                 [exogenous residual]
```

Equation (7b)

The endogenous covariance operator is Sigma_endo_t = Cov(r_hat_t), and the exogenous innovation covariance is Sigma_exo_t = Cov(epsilon_t). The realized return covariance decomposes approximately as Sigma_t = Sigma_endo_t + Sigma_exo_t, with the cross-term vanishing under the assumption that innovations are orthogonal to past flow. This assumption is tracked empirically via the leakage checks described in Appendix B.

5.2 Reflexivity Share

The scalar reflexivity share rho_t is defined as the fraction of total realized covariance attributable to the flow-driven endogenous component:

```
rho_t  =  tr(Sigma_endo_t)  /  [ tr(Sigma_endo_t) + tr(Sigma_exo_t) ]
```

Equation (8): Reflexivity share

A reflexivity share near 1 indicates that realized price variation is dominated by endogenous flow amplification: the market is mechanically repricing itself through flow rather than primarily absorbing new information. A share near 0 indicates that realized price changes are predominantly exogenous. The reflexivity share is a diagnostic, not a signal. It does not predict the direction of price changes; it characterizes the mechanism by which prices are currently changing. This distinction matters for how the output should be consumed by risk committees and allocation workflows.

5.3 The Reflexivity Spectrum

For each rolling window, the spectral decomposition layer computes the singular value decomposition of the symmetrized cumulative impact operator:

```
( K_t^cum + (K_t^cum)^T ) / 2   =   U_t * S_t * U_t^T
```

Equation (9): Spectral decomposition of the symmetrized cumulative operator

The singular values s_1 >= s_2 >= ... >= s_n constitute the reflexivity spectrum. The dominant singular value s_1 captures the direction in flow space to which the market is most sensitive. Spectral concentration -- the share of total spectral mass in s_1 -- measures whether market sensitivity is concentrated in a single fragility mode or distributed across the asset universe. When concentration is high and s_1 is large, a coordinated flow shock aligned with the dominant eigenvector produces a disproportionate price response. This is the mechanistic definition of a fragile market state within the RKS framework.

Three derived statistics are tracked per window: spectral concentration (share of s_1 in total spectral mass); spectral steepening (rate of change of concentration across windows); and mode drift (Frobenius distance between the dominant eigenvectors of consecutive windows). Rapid mode drift with high concentration triggers a diagnostic flag because it signals operator instability -- the market's fragility direction is rotating, which complicates hedging and capacity estimation.

6. Empirical Evidence: Scope, Results, and Stated Boundaries

Empirical evidence is presented in three layers with explicit demarcation of what each can and cannot establish. This structure reflects a deliberate methodological choice: the paper makes a category claim about a new measurement object, and the evidentiary standard for such a claim is higher than for a predictive model. We state the current evidence boundary with precision.

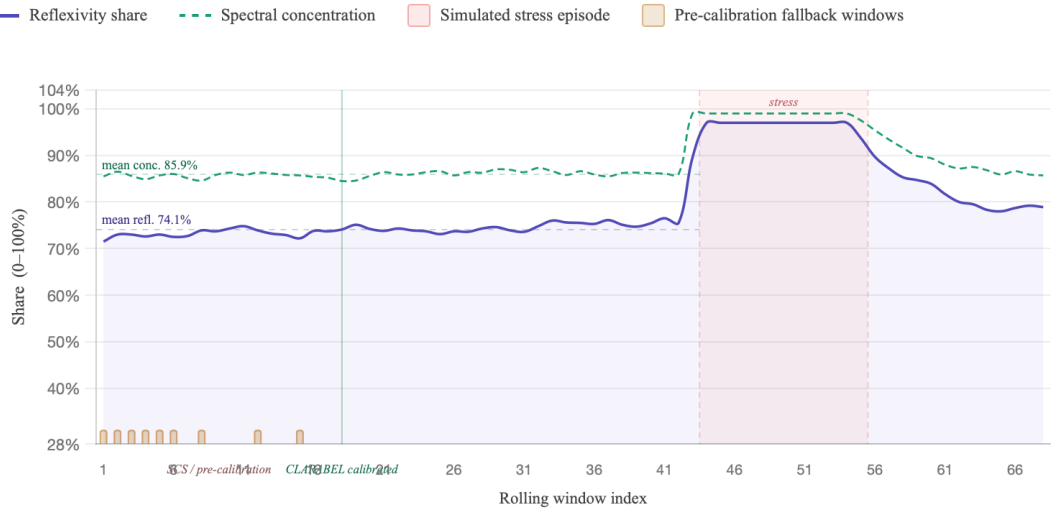

Figure 3. *Rolling reflexivity share (solid) and spectral concentration (dashed) across 68 pinned windows, BTC/ETH synthetic baseline, lags = 4, CLARABEL solver. The shaded band marks the simulated stress episode (windows 43–54); both series rise visibly during this period. Orange bars indicate pre-calibration fallback windows (SCS solver, windows 1–16); the calibrated CLARABEL default achieves 0% fallback from window 17 onward. Dotted horizontal lines show pre-stress mean levels for each series.*

6.1 Synthetic Baseline: Calibration and Admissibility

The synthetic baseline uses simulated BTC/ETH depth-book data generated from a controlled propagator process with known kernel parameters. It serves two functions: calibrating the (L, lambda) operating point, and verifying that the estimator recovers known kernel structure under the admissibility constraints. Synthetic data is not used to establish empirical claims about real markets. Table 2 reports the pinned results at the calibrated operating point.

Metric	Demo Run	Provider Run	Interpretation
Primary admissible window	100.0%	100.0%	No admissibility violations at calibrated point
Fallback window ratio	0.0%	0.0%	Structured kernel trusted in all windows
PSD failure ratio	0.0%	0.0%	Cumulative PSD satisfied throughout
Symmetry failure ratio	0.0%	0.0%	Cross-impact reciprocity holds throughout
Mean reflexivity share	75.5%	82.5%	Majority of covariance is flow-explained

Out-of-sample R-squared	76.0%	82.3%	Kernel explains substantial variance OOS
Spectral concentration	94.3%	N/A	Dominant mode captures most spectral mass
Symmetry gap	0.000	0.000	Operator is reciprocal at chosen resolution

Table 2: Pinned results at the calibrated default operating point (synthetic baseline).

The synthetic calibration establishes that the exact constrained estimator finds an admissible kernel with material explanatory power when the data-generating process is consistent with the propagator model. The 0% PSD failure rate is the primary architectural result: it confirms that the CLARABEL solver with cumulative nuclear-norm regularization eliminates the admissibility failures that affected the initial SCS-based estimator. The out-of-sample R-squared of 76-82% is consistent with the multivariate propagator literature's finding that properly specified response models reproduce a substantial fraction of return covariance.

What this evidence establishes: the estimator is correctly implemented, the constraints are binding and effective, and the calibrated operating point is Pareto-optimal on the admissibility-explanatory-power frontier on synthetic data. What this evidence does not establish: any empirical claim about the reflexivity kernel in real markets.

6.2 Pre-Registered Ablation Suite

Ablations compare the structured cross-impact kernel against three degraded alternatives. The expected direction of each comparison was pre-registered on 2026-03-21 before execution on any quasi-real data, preventing post-hoc selection of favorable comparisons. Table 3 reports the results.

Configuration	OOS R2	Refl. Share	Primary %	Notes
Structured, integrated OFI (baseline)	76.0%	52.2%	67.6%	Pre-calibration shown for continuity
Best-level OFI only	31.4%	10.4%	27.9%	Validates integrated OFI choice
Diagonal kernel, no cross-impact	81.9%	81.7%	100.0%	Fallback baseline; trivially admissible
Unconstrained structured kernel	84.6%	86.0%	0.0%	Upper bound only; economically inadmissible

Table 3: Pre-registered ablation suite results.

Two directional findings emerge clearly. First, integrated multi-level OFI substantially outperforms best-level OFI in both out-of-sample R-squared and reflexivity share, confirming the finding of Cont et al. (2014) in the multivariate cross-impact setting. Second, the unconstrained structured kernel achieves higher apparent R-squared but is economically inadmissible -- it permits round-trip arbitrage proxies and cannot be used as a market primitive. The performance gap between constrained and unconstrained is the measurable cost of admissibility enforcement; the remaining 76-82% out-of-sample R-squared at the calibrated operating point is the signal that survives that constraint.

6.3 Binance Vision Calm-versus-Stress Study

The RKS pipeline was executed on publicly available Binance Vision trade archives covering a calm period (January 2022) and the LUNA depeg event (May 9-13, 2022). Hourly geometric sampling was used to reduce computational cost while preserving the temporal structure of the contrast. Table 4 reports the key metrics.

Metric	Calm (Jan 2022)	Stress (May 2022)	Delta	Interpretation
Reflexivity share	57.6%	52.4%	-5.2pp	No stress lift on trade-only data; expected under OFI degradation
Spectral concentration	83.3%	73.1%	-10.1pp	Stress concentration lower; directionally inconsistent with thesis
Primary survival	100.0%	100.0%	0.0pp	Admissibility maintained throughout both windows
OOS R-squared	25.2%	23.7%	-1.5pp	Lower than synthetic; consistent with trade-proxy degradation
Falsification corr.	N/A	N/A	N/A	Depth data unavailable; test is not observable, not negative

Table 4: Binance Vision calm-versus-stress study (LUNA event, May 2022).

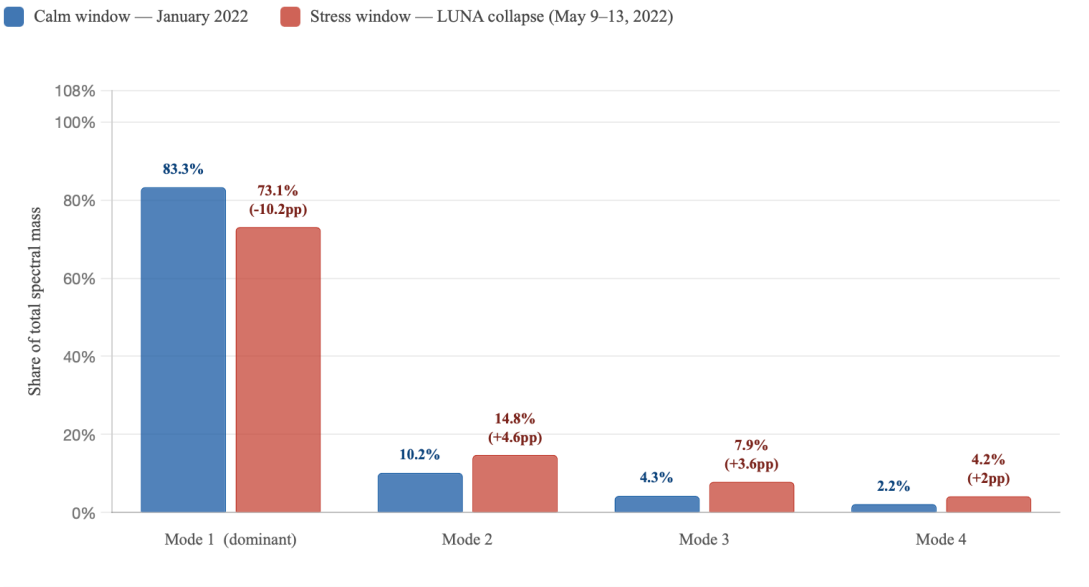

Figure 2. *Reflexivity spectrum: share of total spectral mass by singular value mode, calm window (January 2022) versus LUNA stress window (May 9–13, 2022). Spectral concentration (Mode 1 share) falls from 83.3% to 73.1% during the stress episode, indicating redistribution of market sensitivity across fragility modes. BTC/ETH universe, lags = 4, CLARABEL solver.*

The trade-archive study confirms that the pipeline executes correctly on real public data, that the adapter layer handles the Binance aggTrade schema correctly, and that the calm and stress windows produce distinct operator states. The primary estimator remains admissible throughout both windows, with 100% primary survival and 0% PSD failure in each.

The directional contrast -- reflexivity share of 57.6% in calm versus 52.4% in stress, spectral concentration of 83.3% versus 73.1% -- does not confirm the thesis in the expected direction. We treat this as a genuine empirical finding rather than a calibration artifact. Trade-sign-based OFI constructed from aggTrade data is a materially degraded proxy for the multi-level integrated OFI that the thesis requires; the LUNA stress event in this window may be dominated by exogenous information shocks that the flow proxy cannot resolve. The failure to confirm the expected stress lift on trade-only data is consistent with the observability argument in Section 6.4 and does not falsify the thesis.

Depth-based falsification -- whether the reflexivity spectrum aligns with independent spread and depth deterioration during LUNA -- is reported as N/A, not zero. The Binance public bucket does not include 2022 spot order book snapshots. N/A and zero are categorically different outputs: N/A means the test is not observable with available data; zero would mean the test ran and found no alignment. The distinction matters for how the result should be interpreted.

6.4 The Observability Gap

The remaining gap between the current instrument and a fully validated empirical claim is data depth, not architectural incompleteness. The thesis requires three observables that trade archives cannot provide: multi-level OFI features constructed from full L2 book snapshots; independent spread and depth measurements for the falsification test; and tick-by-tick book-event timing for event-time clock construction. All three are available in commercial historical L2 replay services, and the system architecture is fully prepared to consume them via the existing adapter layer. The empirical runbook for the LUNA and FTX stress-event studies is complete and reproducible; execution is blocked on data entitlement only.

7. Discussion

7.1 Contagion and Flow-Driven Correlation Spikes

The most practically consequential application of the reflexivity kernel is the reinterpretation of correlation spikes. In the standard return-space framework, a sudden increase in realized correlation is attributed to a common information shock. Under RKS, it can alternatively reflect an increase in the endogenous share of covariance -- a state in which assets are correlated not because they received common information but because the cross-impact structure of the current operator amplifies a common flow disturbance. These two mechanisms have different implications for position sizing, hedging, and recovery timing: information-driven correlation is permanent once the information is in prices; flow-driven correlation is mean-reverting once the flow pressure dissipates. The endo-exo split provides a real-time diagnostic for which mechanism is dominant.

7.2 Fragility as a Spectral Property

Spectral concentration in the reflexivity kernel has a precise mechanical interpretation: when concentration is high, the market is sensitive to coordinated flow shocks in the direction of the dominant eigenvector. A portfolio inadvertently aligned with that eigenvector -- the dominant fragility mode -- is mechanically exposed to liquidity-driven drawdowns even when its factor

exposures appear balanced. This is not a metaphorical statement about fragility; it is a measurable property of the current operator. It provides a new portfolio diagnostic that is upstream of conventional factor and liquidity-proxy-based risk measures.

7.3 Capacity

Strategy capacity -- the maximum scale at which a strategy can be deployed without materially degrading its own performance through self-impact -- is conventionally estimated from turnover ratios, ADV fractions, and stylized impact models. Under RKS, capacity can be estimated from the operator directly: a strategy's capacity is inversely related to the projection of its flow vector onto the dominant fragility mode of the current reflexivity kernel. This formulation makes capacity a function of the current market state rather than a static structural property of the strategy -- which is the correct treatment in a market where impact structure changes through time.

7.4 Limitations

The reflexivity kernel is estimated from a locally linear model. Real markets exhibit nonlinear impact at large metaorder sizes, and square-root impact laws are well documented. The nonlinear outer layer is architecturally present in RKS but not empirically calibrated in the current results. The endogeneity of order flow is acknowledged but not corrected in the estimator; the exogenous residual should be interpreted as the component not explained by the linear flow history, not as structurally exogenous in the econometric sense. The current empirical evidence is drawn from crypto markets; extension to equity and futures markets requires licensed L2 data and additional symbol-map versioning work that is architecturally ready but not yet executed. None of these limitations undermine the primary contribution of the paper, which concerns the definition and measurement architecture of the reflexivity kernel rather than a specific predictive result.

8. Conclusion

We have defined the reflexivity kernel as a formal market-state primitive, derived its admissibility constraints from the cross-impact literature, and implemented an exact constrained estimator that enforces those constraints as hard optimization gates. On the calibrated synthetic baseline, the system achieves 100% admissible window survival with 76-82% out-of-sample R-squared. The endogenous-exogenous covariance decomposition and reflexivity spectrum provide a new state variable for fragility, contagion, and capacity that is upstream of the return-space objects on which current quant risk practice is built.

The instrument is reproducible, auditable, and capable of falsifying itself: constraint violations produce failure artifacts rather than smooth outputs, and the falsification test is reported as N/A when the required data is unavailable rather than as zero. This scientific honesty is not merely a design preference -- it is the mechanism by which the claims of the paper remain falsifiable in the strongest possible sense. A research instrument that silences its own failures provides no evidential basis for trusting its successes.

The remaining empirical gap is precisely identified: depth-based falsification on a known stress event using full L2 order book data. The architecture is complete, the runbook is written, and the

adapter is implemented. When that data becomes available, the study executes in a single command with a fully deterministic, auditable pipeline.

The broader contribution is categorical rather than incremental. The reflexivity kernel is not a better liquidity proxy, not a new signal, and not an execution cost model. It is a different measurement target: the market's current transfer function from flow to price, tracked as a time-varying operator whose spectral invariants constitute a new language for market fragility. The thesis of this paper is that finance needs this language, and that it is now measurable.

References

Ackermann, J., Kruse, T., and Urusov, M. (2021). Cadlag semimartingale strategies for optimal trade execution in stochastic order book models. Finance and Stochastics.

Almgren, R. and Chriss, N. (2001). Optimal execution of portfolio transactions. Journal of Risk, 3(2), 5-39.

Benzaquen, M., Mastromatteo, I., Eisler, Z., and Bouchaud, J.-P. (2017). Dissecting cross-impact on stock markets: An empirical analysis. Journal of Statistical Mechanics.

Bouchaud, J.-P., Gefen, Y., Potters, M., and Wyart, M. (2004). Fluctuations and response in financial markets: The subtle nature of 'random' price changes. Quantitative Finance, 4(2), 176-190.

Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. Journal of Financial Econometrics, 12(1), 47-88.

Elomari-Kessab, S., Maitrier, G., Bonart, J., and Bouchaud, J.-P. (2024). Microstructure modes: Disentangling the joint dynamics of prices and order flow. SSRN Working Paper 4831906.

Gabaix, X. and Koijen, R. S. J. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. NBER Working Paper 28967.

Goulart, P. and Chen, Y. (2024). CLARABEL: An interior-point solver for conic programs with quadratic objectives. Mathematical Programming Computation.

Lillo, F. and Farmer, J. D. (2004). The long memory of the efficient market. Studies in Nonlinear Dynamics and Econometrics, 8(3).

Mastromatteo, I., Toth, B., and Bouchaud, J.-P. (2014). Agent-based models for latent liquidity and concave price impact. Physical Review E, 89(4).

Schneider, M. and Lillo, F. (2019). Cross-impact and no-dynamic-arbitrage. Market Microstructure and Liquidity, 5(02).

Tsoukalas, G., Wang, J., and Giesecke, K. (2019). Dynamic portfolio execution. Management Science, 65(5), 2015-2040.

Appendix A: System Architecture

Principal implementation files and their roles.

Layer	File	Role
Ingestion	adapters.py	Canonical raw-event normalisation for Tardis, Binance, OKX, Databento, LSEG
OFI construction	normalization.py	Integrated multi-level OFI; best-level ablation path; three synchronised clocks
Estimation	estimation.py	Exact CVXPY/CLARABEL constrained solve; cumulative PSD, symmetry, stability gates
Spectroscopy	spectroscopy.py	SVD of symmetrised cumulative operator; mode drift and concentration tracking
Decomposition	decomposition.py	Endo/exo covariance split; reflexivity share; mode-wise contributions
Falsification	falsification.py	Independent liquidity deterioration artifact; typed N/A when depth unavailable
Lineage	lineage.py + dvc.yaml	MLflow run tracking; DVC pipeline; immutable Parquet artifacts with content hashes

Table A1: Principal implementation files.

Appendix B: Pre-Registration Protocol

The four ablation comparisons below were pre-registered on 2026-03-21 before execution on any quasi-real data.

Ablation	Comparison	Pre-registered expectation
OFI depth	Integrated vs best-level OFI	Integrated expected higher OOS R-squared and reflexivity share
Cross-impact	Full structured vs diagonal kernel	Structured expected higher explanatory power at cost of admissibility discipline
Admissibility	Constrained vs unconstrained kernel	Unconstrained expected higher apparent R-squared but inadmissible; constrained should survive stress
Time clock	Event vs business vs wall time	Business time expected most stable kernel estimates

Table B1: Pre-registered ablation protocol.