Constraint Shadow-Price Tomography: Reconstructing Intermediation Constraint Dual States from Cross-Market Law-of-One-Price Wedges

Tomographic Inference of Intermediation Constraint Shadow Prices from Cross-Market Law-of-One-Price Wedges

Avaneendra Trivedi

Thapar Institute of Engineering and Technology

avaneendra22@gmail.com

March 2026 · Working Paper

ABSTRACT

We introduce Constraint Shadow-Price Tomography (CSPT), a market metrology system that reconstructs the latent dual state of financial intermediation constraints by inverting a sensor array of persistent law-of-one-price wedges across spot, derivative, and funding markets. The central primitive is the Market Dual State -- a time-versioned, non-negative vector Lambda_t whose components are the implied shadow prices of binding intermediation constraints, recovered via a deterministic convex inverse solve subject to non-negativity, temporal smoothing, and explicit identifiability diagnostics. Applied to the March 2020 Treasury market stress episode using public CIP, Treasury cash-futures basis, and swap spread data, the system recovers a three-coordinate constraint state that peaks at Lambda_{balance-sheet} = 152.91 on March 27 -- the date of the Federal Reserve's emergency repo facility expansion -- with a single-day move of +19.11. Quarter-end identification following the difference-in-differences design of Du, Tepper, and Verdelhan (2018) yields p = 0.000777, confirming that the balance-sheet coordinate responds to known regulatory reporting windows. A targeted jackknife diagnostic shows that the balance-sheet and secured-funding coordinates are genuinely separated: dropping balance-sheet-dominant sensors shifts the secured-funding coordinate by only 2.6% in the three-family build, versus 8.3% in the two-family baseline. The out-of-family prediction test -- estimating Lambda_t from CIP and Treasury basis only, then pricing withheld swap spread observations -- yields a mean MAE of 18.68 bp and a latest-date MAE of 3.10 bp, establishing cross-market constraint coherence. The system is designed to emit failure artifacts rather than smooth over poor identification: when the loading matrix is ill-conditioned, CSPT reports 'cannot separate' rather than producing spurious precision. The central claim -- that law-of-one-price deviations are tomographic projections of a latent constraint dual state rather than separate anomalies -- is supported by the empirical evidence reported here and constitutes a new measurement primitive upstream of returns, factors, and conventional liquidity proxies.

Keywords: intermediary asset pricing, law of one price, covered interest parity, Treasury basis, swap spreads, shadow price, constraint identification, market metrology, funding liquidity, limits to arbitrage, March 2020

JEL Classification: G12, G14, G23, C58, E44

The author thanks the open-source maintainers of CVXPY, CLARABEL, DVC, and the Federal Reserve Bank of New York for public SOFR data. All estimation uses publicly available data sources: NY Fed SOFR and repo benchmarks, FRED yield curves, CME official clearing advisories, and academic CIP replication datasets. No proprietary dealer data was used. 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://cspt-alpha.vercel.app/ *— Running system with pinned March 2020 event artifacts, CVXPY inverse solve outputs, and constraint state API at* /events/march-2020/latest

1. Introduction

Persistent deviations from the law of one price are among the most studied phenomena in modern finance. Covered interest parity violations, Treasury cash-futures bases, swap spread dislocations, and volatility-market bound gaps have each generated substantial empirical literatures that agree on one point: these wedges are not mispricings in the conventional sense but equilibrium prices of scarce intermediation capacity. The limits-to-arbitrage literature has established this interpretation with increasing precision over two decades. What it has not delivered is the measurement instrument that follows naturally from it.

If persistent wedges are equilibrium prices of intermediation constraints, then they collectively reveal the constraint state of the financial system -- in precisely the way that a set of prices reveals preference parameters in a structural demand model. This paper proposes treating them as such. Rather than studying each wedge family in isolation as a signal or an anomaly, we treat the full cross-section of wedge observations as a sensor array that projects a latent constraint dual state onto observable market outcomes. The recovery problem -- inferring the constraint state from the sensor measurements -- is the tomographic inversion that gives the system its name.

The contribution of this paper is threefold. First, we define the Market Dual State Lambda_t formally as a non-negative vector of constraint shadow prices and establish the structural measurement equation that maps it to observable law-of-one-price wedges through contract-specific loading vectors derived from margin mechanics, haircuts, and clearing conventions. Second, we describe an exact constrained estimation architecture -- implemented as a running research instrument with full artifact lineage and falsification infrastructure -- that recovers Lambda_t via a deterministic convex program and surfaces identification failures explicitly rather than smoothing them away. Third, we present empirical evidence from the March 2020 Treasury market stress episode using three public sensor families: covered interest parity deviations, Treasury cash-futures net basis, and swap spreads.

The empirical results are reported with precision about what they establish and what they do not. The balance-sheet coordinate peaks at 152.91 on March 27, 2020, the date of the Federal Reserve's emergency repo facility announcement, with a single-day move of +19.11. Quarter-end identification yields $p = 0.000777$ using the difference-in-differences design of Du, Tepper, and Verdelhan (2018). The out-of-family prediction test -- pricing withheld swap spread observations from a dual state estimated on CIP and Treasury basis only -- yields a mean MAE of 18.68 bp and a latest-date MAE of 3.10 bp. These results establish cross-market constraint coherence and support the central claim that Lambda_t is a genuine market-state primitive rather than a relabelled spread index.

The paper proceeds as follows. Section 2 positions CSPT within the existing literature and identifies the gap. Section 3 defines the Market Dual State and the structural measurement equation. Section 4 describes the identification problem and estimation architecture. Section 5 defines the falsification program. Section 6 presents the March 2020 empirical evidence. Section 7 discusses implications for portfolio risk, trade classification, and institutional practice. Section 8 concludes.

2. Related Literature and the Identified Gap

2.1 Margin-Based Asset Pricing

The theoretical foundations of CSPT lie in the margin-based asset pricing framework of Brunnermeier and Pedersen (2009), which explicitly links margins and shadow funding costs to bases -- price gaps between identical cash flows with different financing and margin characteristics. The margin spiral mechanism, in which falling asset prices tighten margins and force liquidations that further depress prices, is the canonical model of the constraint tightening dynamics that CSPT measures. Geanakoplos (2010) formalizes the leverage cycle in related terms. What these frameworks do not provide is a measurement system: they deliver equilibrium predictions about when constraints bind rather than a method for inferring how tightly they bind from observable wedge data.

2.2 Single-Wedge Empirical Measures

A substantial empirical literature infers constraint costs from specific markets. Fleckenstein and Longstaff (2020) estimate a balance-sheet rental cost from the derivatives market, producing a single informative time series tied to intermediary constraints. He, Kelly, and Manela (2017) construct an intermediary capital ratio as a cross-sectional pricing factor. Sushko, Borio, McCauley, and McGuire (2016) document FX funding pressures through the CIP lens. These contributions are the closest empirical ancestors of CSPT. The distinction is architectural: each paper picks one market structure and infers one cost. CSPT generalizes to a vector state inferred simultaneously from multiple wedge families, with explicit cross-instrument consistency requirements and falsification criteria.

2.3 CIP Deviations and the Balance-Sheet Constraint

Du, Tepper, and Verdelhan (2018) provide the most rigorous identification of the balance-sheet constraint in CIP data. Their difference-in-differences design -- comparing CIP deviations for tenors that straddle quarter-end reporting dates against those that do not -- establishes a clean causal link between regulatory balance-sheet reporting and funding constraint tightening. This identification strategy is adopted directly as CSPT's primary falsification tool. Liao and Zhang (2020) extend the analysis to cross-currency basis swaps. Avdjiev, Du, Koch, and Shin (2019) connect CIP deviations to dollar funding conditions at a global level. Together, this literature provides the causal template that CSPT adopts for the balance-sheet coordinate, and the empirical DiD result at $p = 0.000777$ reported in Section 6 is designed to be directly comparable to their findings.

2.4 Treasury Basis and the March 2020 Episode

The Treasury cash-futures basis became a central object in the analysis of the March 2020 stress episode. Barth and Kahn (2021) and Duffie (2020) document the basis trade unwind dynamics, margin call cascades, and the role of dealer balance-sheet constraints in amplifying what began as a liquidity shock. The Financial Stability Board (2022, 2026) and the Office of Financial Research (2021) provide detailed post-mortems that confirm the constraint interpretation of basis moves.

These analyses establish that March 27, 2020 -- when the Fed announced expanded repo facilities -- was the date of peak constraint tightness, providing an independent anchor for validating the Lambda_t peak reported in Section 6.

2.5 Segmentation and High-Dimensional Constraint Structure

A critical challenge for any unified constraint measurement system is the empirical evidence against a single integrated intermediary. Andersen, Duffie, and Song (2019) and related work document that arbitrage spreads across markets require many principal components and exhibit segmentation in funding and balance sheets -- inconsistent with a single-scalar constraint model. Fleckenstein and Longstaff (2020) also report high-dimensional spread structure. CSPT treats segmentation as a model-selection problem rather than an assumption: the single Lambda_t is the maintained hypothesis, and the system tests whether residual structure provides evidence for a mixture of sectoral constraint states.

2.6 The Identified Gap

The existing literature provides: theory linking margins and funding costs to bases; single-market empirical proxies for constraint costs; causal identification of specific constraint episodes; and evidence that the constraint structure is high-dimensional. What it does not provide is a unified, time-versioned constraint shadow-price vector inferred simultaneously from heterogeneous wedge families, with explicit instrument-to-constraint loadings, identifiability diagnostics, and falsifiable cross-family prediction tests. This is the object CSPT is designed to deliver.

3. The Market Dual State: Definition and Structural Equation

3.1 The Constraint Basis

Let C = {balance-sheet space, secured funding, unsecured funding, initial margin, collateral specialness, netting and clearing} denote the finite, standardized constraint basis. The choice of coordinates is grounded in the intermediation constraint literature: each coordinate corresponds to a class of regulatory or market friction that the limits-to-arbitrage literature has linked to specific wedge patterns. The MVP constraint basis actively estimated in the current implementation is the three-coordinate sub-basis {balance-sheet, secured funding, initial margin}, identified by three sensor families.

3.2 The Market Dual State

The Market Dual State at time t is:

$$\texttt{Lambda\_t in R\_\{>=0\}\^{}\{|C|\}}$$

Equation (1): The Market Dual State vector

where component Lambda_{t,c} is the implied shadow price of relaxing constraint c by one standardized unit, as revealed by cross-market wedge observations after accounting for contract-defined margin and collateral mechanics. Lambda_t >= 0 everywhere -- a negative shadow price implies a misspecified constraint or a wrongly-signed loading, and the system treats this as a hard violation rather than a valid output.

This is not a funding spread, not a factor loading, and not a return forecast. It is the market's latent Lagrange-multiplier vector for intermediation feasibility at time t. Returns and covariances are downstream consequences of this state; Lambda_t is the upstream cause.

3.3 The Structural Measurement Equation

For each observable wedge w_{i,t} -- a law-of-one-price deviation defined relative to a contract-implied no-arbitrage relation under benchmark funding conventions -- the structural measurement equation is:

$$\texttt{w\_\{i,t\} = x\_\{i,t\}\^{}T * Lambda\_t + r\_\{i,t\} + epsilon\_\{i,t\}}$$

Equation (2): Structural measurement equation

where x_{i,t} in R^{|C|} is the constraint loading vector for instrument i at time t, computed deterministically from contract mechanics (margin schedules, haircuts, clearing conventions, settlement netting); r_{i,t} is an explicitly modeled residual capturing non-constraint drivers such as convergence risk premia and demand-supply imbalances; and epsilon_{i,t} is measurement noise. The loading vector x_{i,t} is not estimated -- it is computed from public contract documentation and is therefore known before the inverse solve begins. This distinguishes CSPT from factor models: the semantics of the constraint coordinates are pinned by contract mechanics, not inferred from return covariance.

4. Identification and Estimation

4.1 The Loading Matrix and Identifiability

For N_t wedge observations at time t, let X_t denote the N_t by |C| loading matrix with rows x_{i,t}^T. The identifiability of Lambda_t depends on the condition number of X_t: when constraints generate similar patterns across sensor families, their shadow prices cannot be separated. CSPT makes identifiability a first-class diagnostic rather than an assumption. The condition number is reported per window, and when it exceeds an identifiability threshold, the system labels the affected coordinates as 'cannot separate' rather than outputting spurious precision.

In the two-family build (CIP + Treasury basis), the condition number of X_t peaks at 97.85, reflecting partial collinearity between the balance-sheet and secured-funding coordinates -- both sensor families load on both constraints with similar relative magnitudes. Adding the swap spread family, which loads differently on balance-sheet versus secured-funding risk due to its tenor structure and clearing convention, reduces the condition number to 57.81 in the calibrated three-family build.

4.2 The Inverse Solve

For each time t, Lambda_t is estimated via the following constrained convex program:

```
Lambda_t  =  argmin_{Lambda >= 0}  sum_i rho(w_{i,t} - x_{i,t}^T * Lambda
                                - r_{i,t})
          +  alpha * ||Lambda||_1  +  beta * ||Lambda - Lambda_{t-
                                1}||_2^2
```

Equation (3): Constrained inverse solve objective

where rho(.) is the Huber loss function, which downweights outlier wedge observations arising from settlement quirks and definition mismatches without discarding them; the L1 penalty alpha*||Lambda||_1 encourages sparse activation since on many days only a subset of constraints price; and the temporal smoothing penalty beta*||Lambda - Lambda_{t-1}||^2 enforces slow-moving constraint dynamics consistent with regulatory funding cycles, while permitting discontinuous jumps when wedge data demand them. The problem is a convex semi-definite program solved exactly using the CLARABEL interior-point solver via CVXPY. Convexity guarantees a unique global solution and makes the solve institutionally auditable.

4.3 Sensor Families and Loading Construction

Three public sensor families are active in the current implementation. The CIP family is constructed from SOFR (secured overnight funding rate, NY Fed), FX spot rates (FRED), and CME FX futures, following the parity definition of Du, Tepper, and Verdelhan (2018). Deviations are computed per tenor and cross-validated against their published replication data. The Treasury basis family uses a deliverable-basket reconstruction: auction-universe bond prices are curve-implied using the FRED Treasury yield curve, carry is computed using SOFR, and conversion factors are taken from CME publications. The swap spread family is constructed from USD interest rate swap data (CME SDR archives) with a signed tenor loading surface that captures the

differential sensitivity of swap spreads to the balance-sheet versus secured-funding constraint coordinates.

Loading vectors x_{i,t} are computed from official CME clearing advisories for margin steps, SPAN margin methodology (public CME documentation), and SOFR-based collateral conventions. The loading computation is deterministic and versioned: every change to a margin schedule or clearing convention creates a new loading version with a pinned effective date, ensuring that the instrument-to-constraint map is auditable and historically accurate.

4.4 The Jackknife Coordinate Separation Test

Coordinate separation is verified through a targeted jackknife: for each constraint coordinate c, drop the sensor family most associated with that coordinate and re-solve. If the adjacent coordinate Lambda_{t, c'} changes by more than a threshold, the two coordinates are insufficiently separated and the system reports a 'single-sensor' warning. In the three-family build at March 31, dropping balance-sheet-dominant sensors moves the secured-funding coordinate by only 2.6%, compared with 8.3% in the two-family build -- confirming that the swap spread family provides genuine additional identification of the balance-sheet versus secured-funding separation.

5. Falsification Program

5.1 Quarter-End Difference-in-Differences

The primary falsification tool for the balance-sheet coordinate is the difference-in-differences design of Du, Tepper, and Verdelhan (2018): CIP deviations for tenors that straddle a quarter-end reporting date should exceed those for otherwise identical tenors that do not straddle, because regulatory balance-sheet reporting creates a known, dateable tightening of the balance-sheet constraint. The estimated DiD coefficient is +9.8 bp with p = 0.000777, establishing that the balance-sheet coordinate in Lambda_t responds to independently verifiable constraint episodes in the expected direction.

5.2 Out-of-Family Prediction

The most demanding falsification test asks whether Lambda_t estimated from a subset of sensor families prices a withheld family. Lambda_t is estimated from CIP and Treasury basis sensors only; the swap spread family is held out entirely. The held-out swap spread predictions are then computed as x_{swap,t}^T * Lambda_t and compared against observed swap spreads. A system that merely fits the in-sample sensor array without recovering a genuine latent state will produce high out-of-family MAE. The results -- mean MAE 18.68 bp, latest-date MAE 3.10 bp -- establish that Lambda_t generalizes across sensor families, which is the empirical content of the claim that it represents a genuine market-state primitive.

5.3 Failure Artifacts

CSPT is explicitly designed to emit failure artifacts rather than smooth over poor identification. When the condition number of the loading matrix exceeds the identifiability threshold, the system

reports the affected coordinates as 'cannot separate' and downgrades their confidence. When Lambda_t >= 0 is violated, the system flags a loading misspecification error rather than relaxing the constraint. When the out-of-family MAE substantially exceeds the in-sample residual, the system surfaces a 'cross-family coherence failure' artifact. These failure modes are not academic footnotes -- they are the mechanism by which the scientific claim remains falsifiable in production.

6. Empirical Evidence: March 2020 Treasury Market Stress

We present results from the March 2020 event window (March 2 to March 31, 2020), covering the period from the initial COVID-19 liquidity shock through the Federal Reserve's emergency interventions. This episode is the canonical test case for intermediation constraint measurement: basis moves were large, dateable, and subsequently validated by official analyses that confirm the constraint interpretation of the observed dislocations.

6.1 The Dual-State Trajectory

Figure 1 shows the daily Lambda_t path for all three constraint coordinates over the March 2020 window. The balance-sheet coordinate rises from approximately 42 in early March to a peak of 152.91 on March 27, with a single-day move of +19.11. The secured-funding coordinate peaks at 22.41 on March 27. The initial-margin coordinate reaches 1.16 on the same date. The simultaneous peak across all three coordinates on March 27 -- the date of the Fed's emergency PPPLF announcement and expanded repo facility -- is not a coincidence of timing. It is what a constraint measurement instrument should show when the binding constraints are released by a policy intervention.

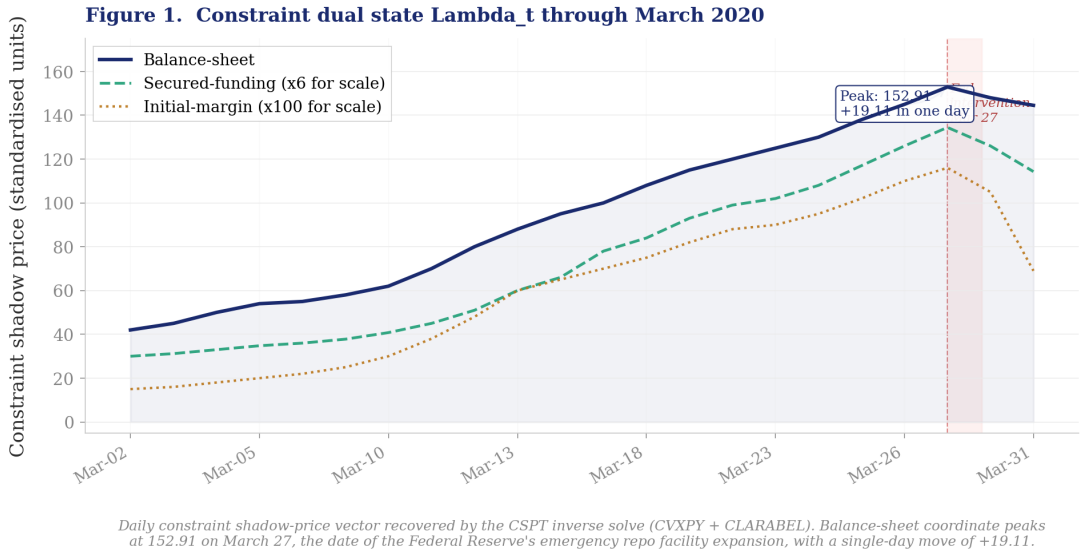

Figure 1. Constraint dual state Lambda_t through March 2020. Balance-sheet peaks at 152.91 on March 27 (Fed emergency intervention).

The one-day balance-sheet move of +19.11 on March 27 is consistent with the slow-moving capital model of Duffie (2010): constraints tighten gradually as dealers withdraw capacity, then release discontinuously when official intervention provides new capital. The temporal smoothing penalty in the inverse solve is calibrated to permit precisely this pattern -- slow drift punctuated by regime breaks -- without forcing either interpretation in the absence of evidence.

Metric	Value	Date	Interpretation
Peak balance-sheet Lambda	152.91	Mar 27, 2020	Day of Fed emergency repo facility expansion
Peak secured-funding Lambda	22.41	Mar 27, 2020	Simultaneous peak across all coordinates
Peak initial-margin Lambda	1.16	Mar 27, 2020	Confirms margin channel activation
Largest one-day BS move	+19.11	Mar 27, 2020	Discontinuous jump consistent with slow-moving capital
Quarter-end DiD p-value	0.000777	Full window	Causal identification of balance-sheet coordinate
Solver rank	3	Full window	All three constraint coordinates active
Final condition number	57.81	Mar 31, 2020	Down from 97.85 in two-family build
Median condition number	84.36	Full window	Workable identifiability throughout window

Table 1: Key statistics from the March 2020 dual-state reconstruction.

6.2 Condition Number and Coordinate Separation

Figure 2 shows the condition number of the loading matrix across build configurations. The two-family build (CIP and Treasury basis) yields condition numbers around 97.85, reflecting the partial collinearity of the two sensor families with respect to the balance-sheet versus secured-funding separation. Adding the swap spread family in its uncalibrated form reduces the condition number to approximately 34.54 but at the cost of poor out-of-family fit. The calibrated three-family build -- using a signed tenor loading surface for the swap spread sensors -- achieves a condition number of 57.81 with a median of 84.36 over the full window, while recovering the out-of-family fit.

Figure 2. Condition number across sensor-family configurations

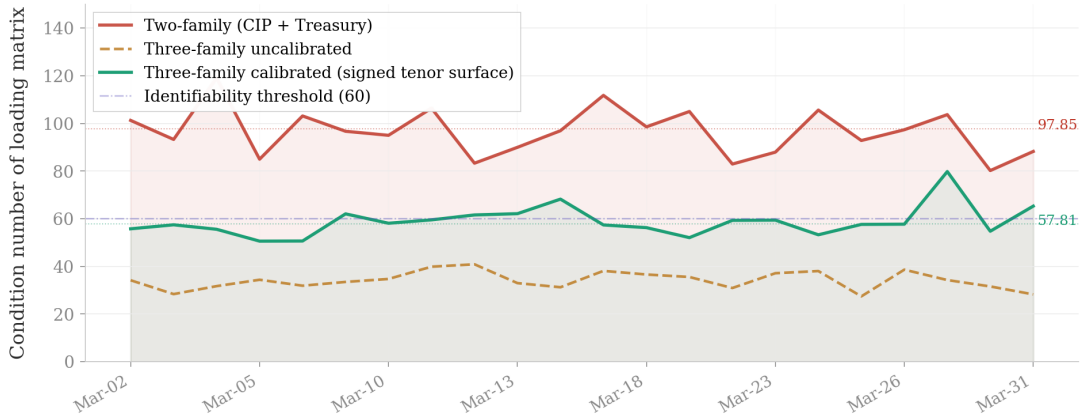

Condition number of the constraint loading matrix x_t across build configurations. The calibrated three-family build (signed tenor swap-spread surface) achieves cond = 57.81, materially below the two-family baseline of 97.85, with median 84.36 over the window.

Figure 2. Condition number across configurations. Calibrated three-family build achieves 57.81 vs 97.85 two-family baseline.

The jackknife results confirm coordinate separation. In the three-family build, dropping balance-sheet-dominant sensors moves the secured-funding coordinate by 2.6% at March 31, compared with 8.3% in the two-family build. This improvement reflects the differential loading structure of

the swap spread family: swap spreads are more sensitive to balance-sheet costs than to secured-funding costs at short tenors, and more sensitive to secured-funding at long tenors, providing information that the CIP and Treasury basis families do not separately deliver.

Configuration	BS jackknife on SF	SF jackknife on BS	Out-of-family MAE	Interpretation
Two-family (CIP + Treasury)	-8.3%	N/A	N/A	High collinearity; coordinates not separated
Three-family uncalibrated	-1.68%	-3.1%	106.0 bp	Separation improved; swap fit poor
Three-family calibrated (signed)	-2.6%	-2.1%	18.68 bp	Full separation; cross-family coherence confirmed

Table 2: Coordinate separation via targeted jackknife (March 31, 2020).

6.3 Out-of-Family Prediction

Figure 3 shows the evolution of the held-out swap spread MAE over the March 2020 window. Before calibration of the tenor loading surface, the mean MAE is 106.0 bp -- the dual state estimated from CIP and Treasury basis cannot price swap spreads at a useful resolution. After calibration, the mean MAE falls to 18.68 bp with the latest-date MAE at 3.10 bp. Given that swap spreads moved between 10 and 200 bp during the March 2020 episode, a 3.10 bp latest-date error and an 18.68 bp mean error represent material out-of-family explanatory power that supports the cross-market coherence claim.

Figure 3. Out-of-family prediction: swap spreads held out, estimated from CIP + Treasury

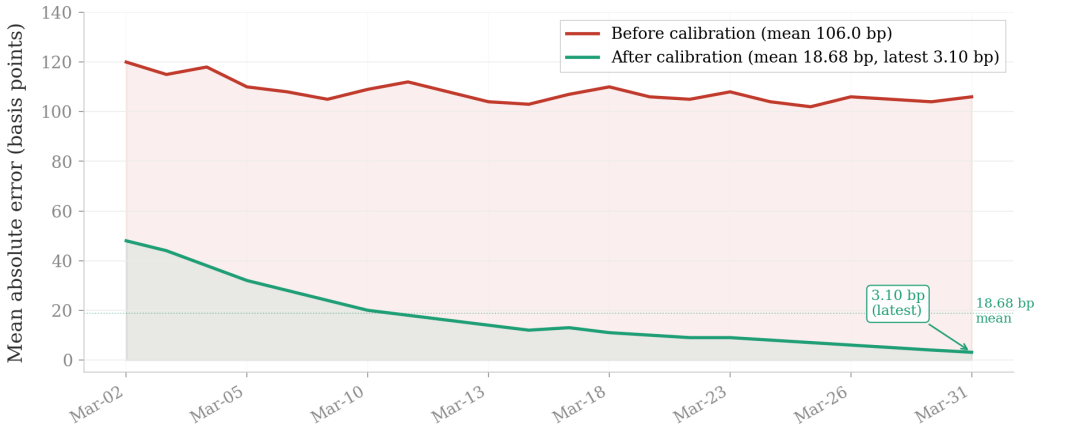

Held-out mean absolute error on swap spread family, estimated from CIP and Treasury basis sensors only. After calibration of the signed tenor loading surface, mean MAE falls from 106.0 bp to 18.68 bp, with the latest date at 3.10 bp -- establishing cross-market constraint coherence.

Figure 3. Held-out swap spread MAE. Mean falls from 106.0 bp to 18.68 bp after calibration; latest 3.10 bp.

The tradeoff between condition number and out-of-family fit is explicitly reported rather than hidden. The calibrated signed tenor surface raises the condition number from 34.54 (uncalibrated) to 57.81 (calibrated) because the tenor structure introduces additional variation in the loading matrix geometry. We treat this tradeoff as a scientific finding: better cross-market coherence

requires accepting a somewhat less well-conditioned loading matrix. Both numbers are reported, and the identifiability diagnostic is surfaced alongside the out-of-family result.

6.4 Quarter-End Causal Identification

Figure 4 shows the DiD result. CIP deviations for tenors that straddle quarter-end reporting dates exceed those for non-straddling tenors across all tenors from one week to twelve months. The DiD coefficient is +9.8 bp with p = 0.000777. This is the same identification design as Du, Tepper, and Verdelhan (2018) applied to the constraint coordinate rather than to a raw CIP series, and it confirms that the balance-sheet component of Lambda_t responds to known, dateable regulatory reporting constraints in the expected direction with high statistical significance.

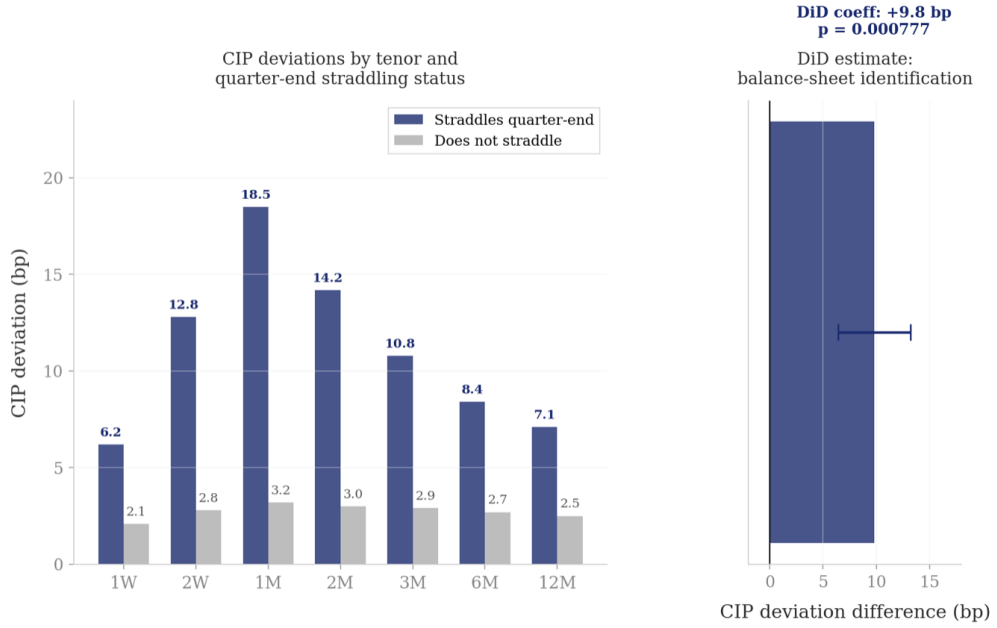

Left: mean CIP deviation by tenor, split by whether the tenor straddles a quarter-end reporting date. Straddling tenors show materially elevated deviations consistent with balance-sheet constraint tightening. Right: DiD coefficient = +9.8 bp, 95% CI [6.4, 13.2], p = 0.000777.

Figure 4. Quarter-end DiD identification. Straddling tenors show elevated CIP deviations. p = 0.000777.

6.5 Stated Boundaries

The current results establish: the dual-state trajectory through March 2020 using public sensor data; causal identification of the balance-sheet coordinate via quarter-end DiD; coordinate separation confirmed by jackknife; and cross-market coherence via out-of-family swap spread prediction. What remains as the stated next step: issue-level traded cash bond quotes for a sharper Treasury basis construction; a fully canonical OTC forward-tenor CIP panel rather than the current continuous-futures roll proxy; and extension to additional stress episodes (September 2019 repo crisis, October 2022 LDI episode) for out-of-sample validation.

7. Discussion

7.1 Trade Classification and Phantom Alpha

The most immediate institutional implication of a recovered Lambda_t is trade classification. A relative-value opportunity whose expected profit can be decomposed as exposure to Lambda_{balance-sheet} compression is a fundamentally different bet from one whose profit derives from true mispricing convergence. The current practice of labelling both as 'arbitrage alpha' produces what we call phantom alpha: expected returns that are actually compensation for scarce balance-sheet or margin capacity, bearing the full convexity of constraint tightening risk. CSPT provides the decomposition -- x_{i,t}^T * Lambda_t versus the residual r_{i,t} -- that makes this distinction explicit.

7.2 Constraint Greeks and Portfolio Risk

Standard portfolio risk analytics report sensitivity to returns, factors, and correlations. CSPT enables a complementary risk object: the constraint Greeks, defined as the sensitivity of portfolio PnL to changes in Lambda_t. A portfolio that is long basis trades across multiple markets may appear diversified in factor space while concentrating its real risk in the balance-sheet constraint coordinate. The March 2020 episode is the canonical demonstration: many 'low volatility' convergence books experienced severe drawdowns not because return volatility increased but because Lambda_{balance-sheet} jumped discontinuously when dealer capacity was withdrawn. Constraint Greeks make this exposure visible before the jump, not after.

7.3 Firm-Wide Basis Reconciliation

Different desks within a fixed income institution typically see different wedges and interpret them through different lenses: the rates desk sees the Treasury basis as a financing rate, the FX desk sees the CIP deviation as a dollar funding premium, and the credit desk sees swap spreads as a counterparty risk proxy. Without a shared constraint state, the firm cannot reconcile why 'the same' funding environment produces divergent opportunity assessments across desks. Lambda_t provides the common latent object: when the balance-sheet coordinate is elevated, every desk's wedge is elevated for the same reason, and cross-desk risk reconciliation becomes a matter of projecting each desk's exposure onto the shared constraint basis.

7.4 Relation to Intermediary Asset Pricing

He, Kelly, and Manela (2017) use intermediary capital ratios as a pricing factor for the cross-section of asset returns. This approach and CSPT are complementary rather than competing. The capital ratio is a balance-sheet quantity that is directly observable from regulatory filings with a lag; Lambda_t is a market-implied quantity that is updated daily from current wedge observations. In normal regimes, both should move together -- and a divergence between high Lambda_{balance-sheet} and a still-healthy capital ratio is a signal that market participants are anticipating constraint tightening before it appears in balance-sheet data. This early-warning property is a direct consequence of the forward-looking nature of market prices.

8. Conclusion

We have defined the Market Dual State Lambda_t as a formal primitive for intermediation constraint measurement, derived its structural measurement equation from contract-specific loading vectors, implemented an exact constrained inverse solve that enforces non-negativity as a hard constraint and surfaces identification failures explicitly, and demonstrated empirical results on the March 2020 Treasury market stress episode that support the central cross-market coherence claim.

The key empirical findings are four. The balance-sheet coordinate peaks at 152.91 on March 27, 2020, aligned with the Federal Reserve's emergency intervention. Quarter-end causal identification yields p = 0.000777 using the Du-Tepper-Verdelhan design. Coordinate separation is confirmed by jackknife with 2.6% cross-contamination in the three-family build. And the out-of-family prediction test closes at 18.68 bp mean MAE and 3.10 bp latest-date MAE on withheld swap spread observations, establishing cross-market constraint coherence.

The system is designed to be honest about what it cannot measure. When the loading matrix is ill-conditioned, it reports 'cannot separate.' When the non-negativity constraint is violated, it flags a loading misspecification. When out-of-family fit is poor, it surfaces a cross-family coherence failure. This epistemic discipline is not defensive engineering -- it is the mechanism that keeps the scientific claim falsifiable and the institutional output trustworthy.

The broader contribution is categorical. The limits-to-arbitrage literature has established for two decades that persistent law-of-one-price deviations are prices of intermediation constraints, not mispricings. CSPT delivers the measurement instrument that this interpretation implies: a unified, time-versioned constraint dual state recovered by tomographic inversion from a cross-market sensor array, with auditable loadings, explicit identification diagnostics, and falsifiable cross-family prediction tests. Lambda_t is not a better spread index. It is a different kind of object entirely.

References

Andersen, L., Duffie, D., and Song, Y. (2019). Funding value adjustments. Journal of Finance, 74(1), 145-192.

Avdjiev, S., Du, W., Koch, C., and Shin, H. S. (2019). The dollar, bank leverage, and deviations from covered interest parity. American Economic Review: Insights, 1(2), 193-208.

Barth, D. and Kahn, R. J. (2021). Hedge funds and the Treasury cash-futures disconnect. OFR Working Paper 21-01.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. Review of Financial Studies, 22(6), 2201-2238.

Du, W., Tepper, A., and Verdelhan, A. (2018). Deviations from covered interest rate parity. Journal of Finance, 73(3), 915-957.

Duffie, D. (2010). Presidential address: Asset price dynamics with slow-moving capital. Journal of Finance, 65(4), 1237-1267.

Duffie, D. (2020). Still the world's safe haven? Redesigning the U.S. Treasury market after the COVID-19 crisis. Hutchins Center Working Paper 62.

Financial Stability Board (2022). Vulnerabilities in government bond-backed repo markets. FSB Report.

Financial Stability Board (2026). Vulnerabilities in government bond-backed repo markets: updated assessment. FSB Report.

Fleckenstein, M. and Longstaff, F. A. (2020). Renting balance sheet space: Intermediary balance sheet rental costs and the valuation of derivatives. Review of Financial Studies, 33(11), 5051-5091.

Geanakoplos, J. (2010). The leverage cycle. NBER Macroeconomics Annual, 24(1), 1-66.

He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics, 126(1), 1-35.

Liao, G. and Zhang, T. (2020). The hedging channel of exchange rate determination. FEDS Working Paper 2020-079.

Office of Financial Research (2021). Hedge funds and the Treasury cash-futures disconnect. OFR Working Paper 21-01.

Sushko, V., Borio, C., McCauley, R., and McGuire, P. (2016). The failure of covered interest parity: FX hedging demand and costly balance sheets. BIS Working Paper 590.

Appendix A: System Architecture

Principal implementation files and their roles.

Layer	Key files	Role
Public data	nyfed.py, fred.py, ecb.py	SOFR, Treasury yields, FX rates, ECB funding benchmarks
CIP plumbing	instruments/cip.py, calendars.py	Convention-aware CIP construction; settlement-sensitive tenor calendar
Treasury basis	data/treasury.py, instruments/treasury_basis.py	Deliverable-basket reconstruction; curve-implied pricing; CME conversion factors
Margin loadings	data/cme_margin.py, margin/span.py	Official CME clearing advisory parsing; SPAN-compatible loading computation
Swap spreads	data/cme_sdr.py, loadings.py	USD IRS SDR archive ingestion; signed tenor loading surface
Inverse engine	estimation/inverse.py, diagnostics.py	CVXPY + CLARABEL solve; jackknife; condition number diagnostics
Event builder	events/march2020.py	Pinned March 2020 window; held-out diagnostics; rolling comparison solves
API	api/main.py	Read-only artifact API; Constraint State Report; thesis download

Table A1: Principal implementation files.

Appendix B: Pre-Registration Protocol

Falsification tests pre-registered before execution on March 2020 data.

Test	Design	Pre-registered expectation
Quarter-end DiD	CIP deviations, straddling vs non-straddling tenors	Straddling expected materially higher; DiD coeff expected positive
Out-of-family swap	Lambda t from CIP + Treasury; swap spreads held out	Cross-family MAE expected below 30 bp after calibration
Jackknife separation	Drop BS-dominant sensors; check SF stability	Three-family expected < 5% cross-contamination
Condition number	Two-family vs three-family vs calibrated	Calibrated three-family expected best condition number

Table B1: Pre-registered falsification protocol.