Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Price Adjustment When Updating Requires Execution: Evidence from Automated Market Makers

Wen-Ting Wang

egpivo@gmail.com

National Chung Hsing University

Research Article

Keywords:

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11128970/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Price Adjustment When Updating Requires Execution: Evidence from Automated Market Makers

Wen-Ting Wang
Institute of Statistics, National Chung Hsing University, Taichung City 402202, Taiwan
egpivo@gmail.com / ORCID: 0000-0003-3051-7302

Abstract

When price updating requires an execution, aggregate correction toward a reference can conceal offsetting movement in the margins that generate it: which discrepancy states a market visits, how often execution occurs in those states, and how much correction follows conditional on execution. We formalize this as an exact decomposition – state composition, execution incidence, and conditional correction magnitude – estimated nonparametrically from sample analogs within fixed discrepancy bins. Passive Uniswap v3 pools make the decomposition unusually transparent: absent a focal-pool execution, native pool price does not change, so the non-execution correction branch is mechanically zero. We implement the decomposition in two matched 5 basis-point (bp) ETH–USDC pools on Ethereum and Base, drawing on a high-frequency analytic panel of 2.52 million block-level observations after reference and finiteness screening, a strict-prior two-venue external reference, and same-block measurement. A pre-specified fixed-hinge diagnostic shows that a simple fee-normalized discrepancy-response shape is not stable across the two non-overlapping windows. More importantly, raw aggregate correction is higher in Confirmation in both pools, while standardizing the two windows to a common discrepancy-state distribution reverses the comparison: positive state-composition and negative conditional-magnitude changes offset, with a smaller execution-incidence contribution, so that raw aggregate correction rises only modestly while these underlying margins move substantially. The offsetting pattern survives removal of the open-ended tail, while its magnitude is support-dependent and amplified on full support. The estimates are specific to the studied pools and windows; the transport analysis is descriptive and post-confirmatory.

Keywords: Price adjustment; Execution-mediated repricing; Automated market makers; Market microstructure; Execution incidence.
JEL classification: G14; G12; C14; C58; G15.

1 Introduction

Price adjustment is typically summarized by an aggregate statistic: an average correction toward some reference, computed over a window or a market state. That summary is rarely a fixed object. Hasbrouck (1991) measures the information content of trades through their persistent price response, while Dufour and Engle (2000) show that the estimated response varies systematically with trading intensity and market state: price impact, the speed of price adjustment, and the persistence of signed order flow all rise as the time between trades falls. Comparing an aggregate correction measure across samples without accounting for this state-dependence can mislead in either direction, whether the measure appears stable or appears to change. A further difficulty arises specifically when repricing is tied to a discrete event. Foucault et al. (2017), studying triangular arbitrage opportunities in foreign exchange, record two distinct ways a discrepancy can end: a trade executes against it, or

a dealer revises a quote with no trade at all, both occurring in their data. Budish et al. (2015) model the underlying race between liquidity providers cancelling stale quotes and traders picking them off. In such continuous, quote-driven venues, a non-execution price-update branch remains alongside the execution branch, so execution incidence and correction conditional on execution do not by themselves reconstruct the aggregate: an average correction conflates which discrepancy states a market visits, how often an execution occurs in those states, and how much correction follows conditional on that execution, and these three margins can move independently or offset one another. The measurement problem this creates – separating what changed in the states visited from what changed in the response to them – is not specific to this paper's setting: van der Beck et al. (2026) show that gross trading volume rose roughly fivefold while trading associated with persistent portfolio reallocation changed far less, and separately document a divergence between improvements in high-frequency liquidity measures and long-horizon price impact, and Barardehi and Bernhardt (2025) show that the textbook intraday U-shaped volatility and price-impact pattern changes materially under trade-time rather than calendar-time aggregation. Neither paper shares this paper's estimand or technique; both establish, in unrelated markets and periods, that a headline market-quality statistic can look stable or can move for reasons that are themselves decomposable into offsetting or reclassified underlying components – exactly the kind of concealment a raw aggregate correction measure risks in the execution-mediated setting studied here.

Passive Uniswap v3 concentrated-liquidity automated market maker (AMM) pools remove the non-execution branch by construction and make this decomposition exact rather than approximate. Because native-price revision in these pools requires a reserve-changing swap, a block without a focal-pool swap has, by construction, zero native-price correction: the non-execution adjustment branch is mechanically absent for native pool price, and pre- and post-state prices can be aligned exactly around the executing block. This measurement property belongs to the passive pools studied here rather than to onchain market making or trading venues in general – actively managed venues may revise pricing state separately from a user trade – and it does not make these pools representative of all financial markets, but it is why the exact incidence-times-conditional-correction decomposition is directly available here, without the additional structure a continuous, quote-driven venue would require. AMM arbitrage theory supplies the economic context for why a discrepancy can persist until execution: Milionis et al. (2023) show that with a proportional fee, an AMM price can drift within a no-arbitrage band around an external reference until an arbitrageur's trade corrects it, and that fees scale arbitrage profits by the fraction of time an arriving arbitrageur finds a profitable trade in the low-fee, fast-block asymptotic regime. That is a theoretical incidence analogue to the execution margin studied here, but their object remains model-based and structurally distinct from our empirical conditional-correction measure. Recent continuous-time CLMM theory models arbitrageur optimal control against an evolving liquidity profile rather than measuring realized incidence and conditional correction from data (Tung and Wang, 2026), and crypto execution and market-impact evidence on continuously quoted venues measures a trader's own optimal order-placement choice or signed order-flow's price impact rather than correction conditional on an execution-mediated update rule (Barucci et al., 2023; Bundi et al., 2024). Our object is instead a realized, bin-level decomposition of observed correction into state composition, execution incidence, and conditional correction, estimated directly from data without a structural model of arbitrage arrival or optimal execution.

We formalize the decomposition itself as an exact accounting identity, not as the paper's contribution. For window w and discrepancy state k, aggregate correction is $A_w = \sum_k \pi_{wk}\, p_{wk}\, m_{wk}$, where π is state composition, p is execution incidence, and m is conditional correction magnitude; Section 2 states this precisely. That an average factors this way once the non-execution branch is zero follows immediately from the law of total expectation, and we do not present it as new. What is not an accounting matter is that, in this setting, all three margins are separately, nonparametrically

observable from realized on-chain data – without modeling latent arbitrage arrivals or a trader's optimal execution problem – and can therefore be estimated, compared across samples, and reallocated to explain a change in the aggregate. The paper's contribution is this combination: an institutional setting in which the three margins are directly observable without a structural model; a realized high-frequency implementation of the decomposition in two matched 5 bp ETH–USDC Uniswap v3 pools on Ethereum and Base; a cross-window transport of the decomposition that separates changes in the states visited from changes in execution incidence and conditional correction; and the empirical finding that aggregate correction can move only modestly while these underlying margins move substantially, in offsetting directions. The clearest illustration of why this separation matters is a composition-induced sign reversal: raw aggregate correction rises in Confirmation relative to Discovery in both primary pools, whereas the common-state-weight Confirmation-minus-Discovery comparison is negative in both. The measurement logic is conceptually relevant to other settings in which repricing is tied to discrete, observable updating events, but the exact decomposition used here relies on the execution-mediated update rule of the studied pools; venues that can revise prices without execution require a different implementation, and the empirical finding below is specific to the studied pools and windows.

We implement the decomposition in two matched 5 bp ETH–USDC pools on Ethereum and Base, drawing on a 57-day high-frequency analytic panel comprising 2.52 million block-level observations after reference and finiteness screening across the two pools (354,506 for ETH5 and 2,164,574 for BASE5), a strict-prior reference from two external venues, and same-block measurement of execution incidence and conditional correction magnitude. Calendar span and block-level observation density are different dimensions: the panel is dense within a 57-day window, and the two non-overlapping sub-windows used below are correspondingly short in calendar time, a limitation we return to in the Discussion.

A deliberately simple fee-normalized diagnostic – a fixed hinge in the discrepancy coordinate, its directional prediction pre-specified after the discovery period but before holdout outcomes were inspected – is not stable across the two windows: the estimated slope change is negative rather than the predicted positive value in both primary pools. This provides supporting evidence against treating one global reduced-form response curve as stable across the studied windows.

Cross-window changes in aggregate correction conflate changes in the states a market visits with changes in the response conditional on those states. Response-map transport (RMT) applies the π, p, m decomposition to this question, transporting one window's response map onto the other's state frequencies; a symmetric Shapley allocation of the change in the aggregate functional $A(\pi, p, m)$ – an exact, order-independent accounting split, not a causal or structural decomposition – summarizes positive state-composition and negative conditional-magnitude contributions, with a smaller execution-incidence contribution. This pattern is concentrated in the support: it survives removing the open-ended tail bin in both primary pools, is not established on $x \leq 2$, and is amplified by a large compositional shift into that tail bin between the two windows, detailed in Section 4. RMT and its support-cap robustness were designed after the fixed-hinge diagnostic and are post-confirmatory exploratory rather than pre-specified. These allocations concern mean correction, not the hinge coefficient or its sign reversal.

Cryptocurrency microstructure evidence is otherwise close in setting but not in object (Makarov and Schoar, 2020; Brauneis et al., 2022a,b; Easley et al., 2026; Shen et al., 2025): the cited studies do not jointly separate state composition, execution incidence, and conditional correction or transport those margins across sample windows. Li et al. (2026) calibrate a related structural correction-strength parameter, Najnudel et al. (2024) provide adjacent theoretical context, and Lehar and Parlour (2025), Capponi and Jia (2025), and Aoyagi and Ito (2025) study passive AMMs as market institutions rather than as a laboratory for price-adjustment margins. Transporting one window's response

map onto another's state frequencies to explain a change in the aggregate follows the logic of direct standardization (Kitagawa, 1955); related semiparametric counterfactual reweighting applies the same composition-versus-conditional-structure logic to richer covariate distributions (DiNardo et al., 1996), and Section 4 makes the present, finite-cell correspondence precise. Observed swap incidence and conditional correction are not estimates of latent arbitrage arrivals or structural correction parameters.

2 Measurement Design: An Empirical Laboratory for Execution-Mediated Adjustment

For a focal pool and block t, let P_t^{pre} and P_t^{post} be the pool's native marginal price immediately before and after the block. Let X_t^* be the strict-prior external reference. The pre-state absolute discrepancy is

$$d_t^{pre} = |\log X_t^* - \log P_t^{pre}|, \qquad \gamma_f = -\log(1-\phi), \qquad x_t = d_t^{pre}/\gamma_f,$$

where ϕ is the nominal pool fee rate and γ_f is its log-fee equivalent. The same-block signed correction is

$$C_t = d_t^{pre} - |\log X_t^* - \log P_t^{post}|.$$

Positive values indicate movement toward the reference; negative values indicate movement away. Let $S_t = 1$ if the focal pool has at least one swap in the block.

Under the same-block native-price outcome construction,

$$C_t = 0 \quad \text{when } S_t = 0,$$

since a block without a focal-pool swap leaves the pool's native price unchanged. This is an accounting consequence of the outcome definition, not a claim about the Uniswap v3 protocol mechanics beyond that update rule. In a venue where prices can revise without execution, a non-execution response term would remain and this simplification would not hold. Thus, for each discrepancy value or fixed bin,

$$\mathbb{E}[C \mid x] = P(S = 1 \mid x)\mathbb{E}[C \mid S = 1, x].$$

The first factor is exactly the observed probability of a focal-pool swap, not $P(\text{corrective} \mid x)$. The second factor is exactly the observed conditional mean correction among swap blocks. Figure 1 records the timing, and Figure 2 (Section 3) reports both terms directly as response maps across the fixed discrepancy bins.

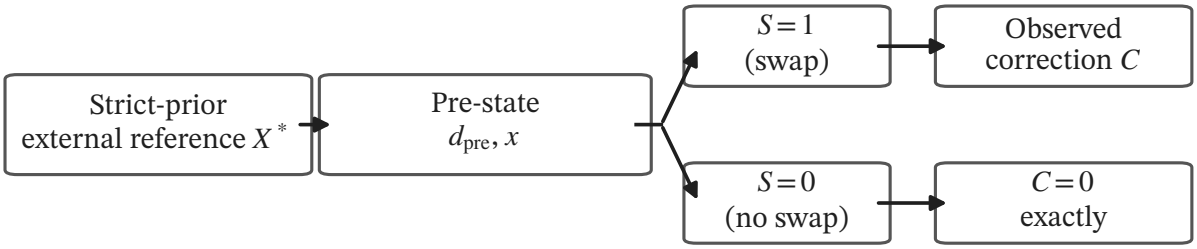
This identity organizes the paper's measurement representation. Partition the discrepancy coordinate x into fixed bins indexed by k, and let window w denote a sample period (Discovery or Confirmation) or a market. Define

$$\pi_{wk} = P_w(x \in B_k), \qquad p_{wk} = P_w(S = 1 \mid x \in B_k), \qquad m_{wk} = \mathbb{E}_w[C \mid S = 1, x \in B_k],$$

so that π is state composition, p is execution incidence, and m is conditional correction magnitude. Because $C = 0$ whenever $S = 0$, $\mathbb{E}_w[C \mid x \in B_k] = p_{wk} m_{wk}$, and aggregate mean correction over window w is

$$A_w = \sum_k \pi_{wk}\, p_{wk}\, m_{wk}.$$

The three objects are estimated by their sample analogs within each fixed bin: $\hat{\pi}_{wk} = N_{wk}/N_w$, $\hat{p}_{wk} = N_{wk}^S/N_{wk}$, and $\hat{m}_{wk}$ is the sample mean of C among swap blocks in bin k, giving the plug-in estimate $\hat{A}_w = \sum_k \hat{\pi}_{wk}\hat{p}_{wk}\hat{m}_{wk}$. This is a nonparametric, sample-analog decomposition rather than a parametric or structural model: it estimates observed conditional moments within fixed bins and



$$E[C \mid x] = P(S=1 \mid x)\, E[C \mid S=1, x]$$

Figure 1: Measurement timing and exact decomposition. The non-swap branch has $C = 0$ by the native-price update rule; the displayed equality is therefore an accounting identity for this outcome.

combines them by the accounting identity above, and it is not presented as a new method-of-moments or GMM estimator. Section 4 applies A_w across windows; Supplementary Appendix G gives the full correspondence between π, p, m, their sample analogs, and the response-map transport and Shapley calculations reported below.

2.1 Sample and Reference Construction

The high-frequency sample runs from 9 June through 5 August 2026. It includes the 5 bp Uniswap v3 ETH–USDC pool on Ethereum (ETH5), and 5 bp (BASE5) and 1 bp (BASE1) Uniswap v3 ETH–USDC pools on Base. All three pools are confirmed Uniswap v3 by direct on-chain verification of their factory and protocol fee/tick-spacing settings. The primary empirical comparison is ETH5 versus BASE5; BASE1 is reported as a secondary market because its later contextual classification is mixed. Table 1 separates the correction sample from the longer historical context panel.

The primary strict-prior reference uses Bybit ETHUSDC and OKX ETH–USDC direct-USDC best bid and offer (BBO) quotes. For each block it takes the last quote strictly before the block timestamp from each venue, requires both quotes to be no more than three seconds old, and uses the equal-weight average of their log midpoints. There is no one-venue fallback. This makes the reference a direct, exchange-time construction rather than a proxy stabilized by a USDT basis estimate. The timing sequence is: external quotes are recorded by each venue, the last quote strictly before the focal block timestamp fixes the strict-prior reference, the focal pool's pre-state price is read at that same block, any within-block swap(s) execute, and the pool's post-state price is read at block close (Figure 1). Because the reference is fixed before the block timestamp rather than resampled at the moment of the swap, C_t measures change in distance to this fixed strict-prior benchmark, not correction toward the external price contemporaneous with the swap; the external market can move within the block. Reference age – the gap between the strict-prior quote and the block timestamp – has median 0.11–0.16 seconds and 95th percentile 1.2–2.4 seconds across the two venues and two chains, comfortably inside the three-second freshness gate; the full age distribution is reported in Supplementary Appendix A. Transient microstructure noise in either venue's quote can still perturb X_t^*, although the both-venues-required construction, the equal-weight averaging of the two direct-USDC log midpoints, and the three-second freshness requirement reduce reliance on any

Table 1: Sample construction: candidate, reference-flag, and analytic populations

Pool	Candidate N	Reference-flag N (%)	Analytic N (%)
ETH5	393,305	377,250 (95.9%)	354,506 (90.1%)
BASE5	2,393,726	2,302,099 (96.2%)	2,164,574 (90.4%)
BASE1	2,393,726	2,302,099 (96.2%)	—

Notes: ETH5 (Ethereum, 5 bp), BASE5 (Base, 5 bp), and BASE1 (Base, 1 bp) are confirmed Uniswap v3 ETH–USDC pools (Section 2.1). All three share the same 09 Jun–05 Aug 2026 candidate window. $x = d^{pre}/\gamma_f$, where $\gamma_f = -\log(1-\phi)$ and ϕ is the nominal pool fee rate. Same-block correction C is defined only in the 57-day reference window (57 of the 58 calendar dates spanned; see Section 2.1), using the two-venue direct-USDC reference. **Reference-flag** N is the population with a real on-chain timestamp, finite pre/post pool prices, and a valid external-reference flag; this flag does not itself guarantee the underlying reference value is a finite number. **Analytic** N is the strictly smaller population (Discovery + Confirmation) additionally requiring the derived discrepancy x and correction C to be finite – the population used for the primary-pool π, p, m, raw mean correction, RMT, Shapley, and their calendar-day bootstrap (Supplementary Appendix G); the holdout regression uses the corresponding finite-x/C population restricted to swap blocks. Not reported for BASE1, which is secondary and for which no RMT is run. The two populations differ by about 6% at both primary pools, traced to blocks where the reference-validity flag is set but the underlying reference value is not itself finite. Each pool also has a longer historical daily-context series, distinct from this reference-qualified sample and containing neither x nor C, ending within one calendar day of the candidate window above; its start dates and its position within each pool's own history are reported in Supplementary Appendix F and Table S7. Pool addresses are reported in Appendix B.

single stale quote; X_t^* remains an empirical strict-prior benchmark, not a latent efficient price.

Alternative USD and stablecoin-basis proxies are evaluated against this primary reference in Supplementary Appendix A; neither is used to extend the correction sample. The upstream reference-eligibility flag – a real on-chain timestamp, finite pre/post pool prices, and a valid external-reference flag – marks 377,250 ETH5 blocks and 2,302,099 Base blocks. Requiring the derived discrepancy x and correction C to also be finite leaves strictly smaller analytic populations of 354,506 ETH5 and 2,164,574 BASE5 observations, used in the primary decomposition and transport analysis (Table 1). Prices are reconstructed deterministically from on-chain events, with pre- and post-state prices aligned to the same block. The execution indicator $S_t = 1$ records whether the focal pool has at least one swap in the block; it is not a count of swaps and not an estimate of an arbitrageur's arrival probability, and C_t remains the net block-level outcome across however many swaps occur.

The reference-qualified panel contains 57 represented UTC days: 30 in Discovery and 27 in Confirmation, although the endpoint span covers 58 named calendar dates (9 June–5 August 2026). A separate historical panel places this window in each pool's own distribution of AMM-native activity, liquidity, and volatility, including an AMM-native realized-volatility measure (not a CEX-based volatility estimate). ETH5 and BASE5 are broadly typical on all six substantive metrics, including median liquidity (26.2nd and 28.4th percentile of each pool's own history, respectively) and realized volatility (37.6th and 32.2nd percentile, respectively): on these AMM-native measures, the study window is not an extreme-liquidity or extreme-volatility episode for either primary pool. BASE1 is mixed, with Mint/Burn activity at its 98.3rd percentile. Unusually intense liquidity-provision churn of this kind changes active-liquidity and tick-depth conditions during the window – not price directly, since Mint/Burn events are not swaps and do not enter S – which makes BASE1's historical window less comparable as a stable regime and is why it is treated as secondary. Supplementary Appendix F retains the full definitions, table, and figure. These histories contain no external reference and cannot extend the price-correction or transport evidence.

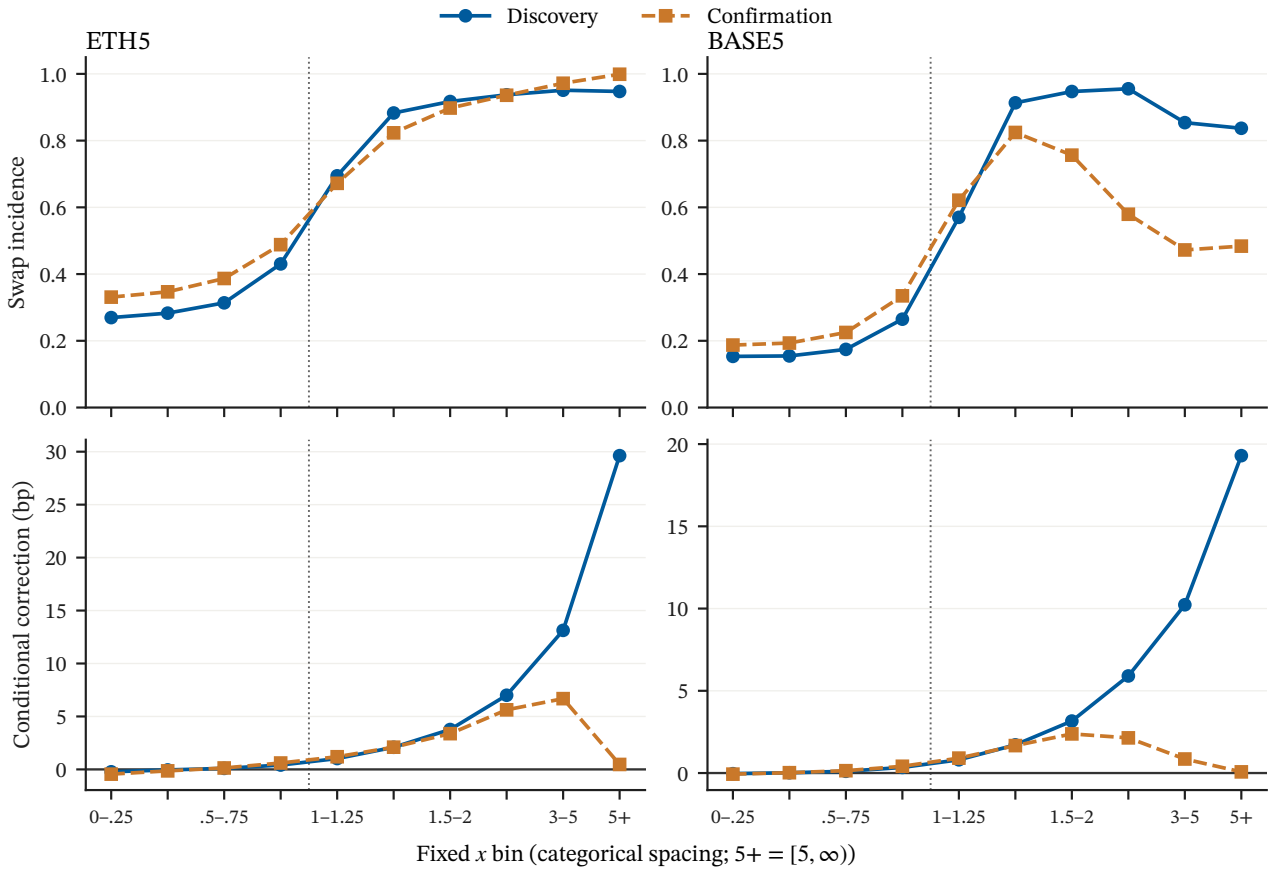


Figure 2: Observed incidence and conditional-correction response maps (post-confirmatory transport is discussed in Section 4). Discovery and Confirmation incidence (top) and conditional correction magnitude (bottom) use the existing fixed bins for ETH5 and BASE5. All ten fixed bins are retained; some tick labels are omitted for readability. The 5+ label denotes the open-ended $[5, \infty)$ bin. Horizontal spacing is categorical, and connecting lines are visual guides only. The dashed line marks the existing $x = 1$ coordinate, not an estimated breakpoint.

3 Observed Response Maps and Cross-Window Stability

The Discovery window is 6 July–5 August 2026. It is exploratory, and its role is to describe the estimated p_{wk} and m_{wk} maps rather than to establish a transportable fee law. Figure 2 displays these maps directly for the primary 5 bp pools in both the non-overlapping Discovery and Confirmation windows; Supplementary Figures S2–S3 and Supplementary Table S6 retain every fixed bin in per-margin and full tabular form.

For ETH5, observed execution incidence $\hat{p}_k$ rises sharply across the central bins and remains high in the upper bins. Conditional correction $\hat{m}_k$ among swap blocks also rises across the Discovery support. For BASE5, the two margins are again distinct: incidence is lower at small discrepancies and conditional correction increases through the middle of the support. These are descriptions of realized blocks. They do not say that an arbitrary opportunity must produce a swap, nor that the observed conditional mean is a structural correction coefficient.

ETH5 and BASE5, both confirmed Uniswap v3 pools (Section 2) at the matched 5 bp fee tier, show increasing observed swap incidence through the central discrepancy range and increasing conditional correction in the Discovery bins, but their levels and large-support behavior differ: a

matched-protocol, matched-fee cross-chain comparison that improves institutional comparability without identifying a chain effect or other causal contrast. BASE1 is reported alongside them but is not used to assert matched-fee transport, since its historical high-frequency window is a mixed regime because of unusually intense liquidity-provision churn (Section 2.1); Supplementary Appendix C retains a distinct 12-second tracking object for readers who require a fixed-time comparison, but its decomposition is not exact.

3.1 Fixed-Hinge Stability Diagnostic

The response maps above already suggest that a single global discrepancy-response shape may not summarize the data well. We evaluate this directly with a parsimonious fixed-hinge diagnostic.

The fee-normalized coordinate $x = d^{pre}/\gamma_f$ is economically motivated rather than an arbitrary rescaling: with a proportional trading fee, an AMM price can drift within a no-arbitrage band around an external reference until an arbitrageur's trade corrects it (Milionis et al., 2023), so discrepancy relative to the fee is an economically meaningful coordinate. This theoretical context motivates the *scale* on which correction is studied here; it does not predict an exact breakpoint, a positive β_2, or the specific hinge specification below – $x = 1$ is a simple fixed empirical coordinate, not a theoretically established structural breakpoint. The Discovery response map above motivates the *shape*: conditional correction rises across most of the Discovery support in both primary pools, suggesting that a single global linear slope may be too restrictive, so a fixed hinge around the one-fee coordinate is a parsimonious reduced-form diagnostic worth testing out of sample. We accordingly specify a same-block, swap-only conditional-correction regression,

$$C = \beta_0 + \beta_1 x + \beta_2 (x - 1)_+ + \varepsilon,$$

where β_1 is the slope below the fixed hinge, β_2 is the change in slope beyond $x = 1$, and $\beta_1 + \beta_2$ is the slope above the hinge. This reduced-form regression is a parsimonious diagnostic of the Discovery shape, not a structural arbitrage model, a breakpoint estimator, or the paper's main π, p, m measurement framework. The breakpoint is fixed at $x = 1$; there is no breakpoint search. Because the Discovery shape suggested steeper conditional correction beyond the one-fee region, the tested directional prediction was $\beta_2 > 0$. After the exploratory Discovery analysis and before inspecting outcomes in the held-out confirmation window, the breakpoint, this expected sign, and the calendar-day bootstrap protocol were specified. The Confirmation window is 9 June–5 July 2026, contains 27 calendar days, excludes the Discovery window, and uses 2,000 calendar-day bootstrap replications with seed 20260829. With 27 day-level clusters, the reported intervals are descriptive uncertainty summaries under calendar-day resampling rather than evidence of independent day-level replication; ETH5 and BASE5 share overlapping calendar dates and market conditions, so their confirmation-window results are not statistically independent of one another.

The directional result is unambiguous. For ETH5, $\hat{\beta}_2 = -3.653$ bp per unit x, with 95% interval $[-4.146, -3.026]$; for BASE5, $\hat{\beta}_2 = -1.478$ bp per unit x, with interval $[-1.775, -1.107]$. In both primary pools, the estimated slope change has the opposite sign from the Discovery-motivated positive directional prediction. The secondary BASE1 estimate is also negative. The positive slope-change pattern suggested by Discovery does not carry over globally to the holdout window: this fixed-hinge diagnostic is not stable across the two windows. This concerns the parsimonious fixed-hinge diagnostic specifically and does not extend to the paper's main π, p, m measurement framework, which Section 4 unpacks directly.

The instability is itself support-dependent: local windows around $x = 1$ and caps through $x \leq 2$ have positive hinge estimates in both primary pools and both windows, whereas full-support estimates reverse. Supplementary Appendix D (Table S4, Figure S1) reports the complete fixed ladder. These

Table 2: Pre-specified holdout diagnostic for the global fixed-$x = 1$ hinge

Market	N	Days	$\hat{\beta}_0$ (bp)	$\hat{\beta}_1$ (bp per x)	$\hat{\beta}_2$ (bp per x) [95% CI]
ETH5	89,577	27	-1.428	3.608	-3.653 [-4.146, -3.026]
BASE5	309,109	27	-0.474	1.444	-1.478 [-1.775, -1.107]
BASE1	490,893	27	-0.797	1.674	-1.680 [-1.896, -1.407]

Notes: Same-block conditional correction among swap blocks, 09 June–05 July 2026. Model: $C = \beta_0 + \beta_1 x + \beta_2(x-1)_+ + \varepsilon$, where x is dimensionless. After exploratory Discovery analysis and before holdout outcomes were examined, the breakpoint, expected positive β_2 sign, and calendar-day bootstrap protocol were specified. The confirmation analysis uses 2,000 calendar-day bootstrap replications (seed 20260829). $\hat{\beta}_2$ is negative at every pool, opposite the pre-specified positive expectation. BASE1 is secondary.

diagnostics characterize the support sensitivity of the chosen summary; they do not identify a threshold.

The support sensitivity of the hinge estimate points to an economic distinction between arbitrage incentive and realized price adjustment. Trading fees help determine when a discrepancy creates a profitable opportunity, but the correction actually observed within a block also depends on execution conditions and available liquidity (Milionis et al., 2023; Li et al., 2026). Table 2's $\hat{\beta}_1 + \hat{\beta}_2$ – the fitted slope above the hinge – is close to zero at full support in both pools (-0.045 for ETH5, -0.034 for BASE5), which is not the same as observing negative correction at large discrepancy; it reflects how a single linear-plus-hinge curve fits across the full support, not that price adjustment reverses direction. Execution incidence and realized correction can also diverge sharply at the sparsest states: in the Confirmation $[5, \infty)$ bin, ETH5 has a swap in nearly every block (12,871 of 12,885) yet a conditional correction mean under 0.5 bp, while BASE5's Confirmation incidence in the same bin is about half that (26,763 of 55,322, roughly 48%) with an even smaller conditional correction mean. Neither pattern is evidence about arbitrage participation, MEV, or delay specifically – the composition of discrepancy values within this open-ended bin can itself differ across windows, and the strict-prior reference is fixed before the block – but together they show that a block containing a swap does not by itself indicate that the discrepancy was corrected.

4 State Composition and Response-Map Transport

Cross-window changes in aggregate correction conflate changes in state exposure with changes in the response conditional on state. The π, p, m representation separates these margins directly: when aggregate adjustment $A_w = \sum_k \pi_{wk} p_{wk} m_{wk}$ differs across samples, how much reflects a change in π (the states visited), how much reflects a change in p (how often execution occurs in those states), and how much reflects a change in m (how strongly prices move conditional on execution)? Response-map transport (RMT) applies the π, p, m decomposition to this question. This analysis was designed after the holdout evaluation in Section 3.1 and is therefore post-confirmatory. Applying window D's incidence and magnitude maps to window H's state frequencies gives the transported aggregate $A_{H \leftarrow D} = \sum_k \pi_{Hk}\, p_{Dk}\, m_{Dk}$, compared against realized $A_H = \sum_k \pi_{Hk} p_{Hk} m_{Hk}$; the difference isolates how much of the change in aggregate correction is attributable to D's response map rather than to H's own state composition. This finite-cell transport follows the logic of direct standardization and composition-rate decomposition associated with Kitagawa (1955); related semiparametric counterfactual reweighting extends the same composition-versus-conditional-structure logic to richer covariate distributions (DiNardo et al., 1996). Discovery is later in calendar time than Confirmation: transporting its map backward is a retrospective comparison, not a prospective forecast. For each window and discrepancy bin, we measure its frequency, swap incidence, and

Table 3: Post-confirmatory exploratory response-map transport (bp)

Quantity	ETH5	BASE5
Raw Discovery mean $\bar{C}_D$	0.484	0.107
Raw Confirmation mean $\bar{C}_H$	0.580	0.143
Discovery map on H support $C_{H\leftarrow D}$	2.792	1.190
Transport error $TE_{D\to H}$	-2.212 $[-4.559, -0.519]$	-1.047 $[-2.600, -0.261]$
Standardized Discovery $C_{D,\text{std}}$	1.638	0.648
Standardized Confirmation $C_{H,\text{std}}$	0.497	0.124
Standardized difference STD_DIFF	-1.141 $[-2.328, -0.287]$	-0.525 $[-1.308, -0.126]$
State composition Δ_π	$+1.256$ $[0.457, 2.406]$	$+0.489$ $[0.158, 1.115]$
Execution incidence Δ_p	$+0.011$ $[-0.009, 0.043]$	-0.083 $[-0.224, -0.019]$
Conditional magnitude Δ_m	-1.171 $[-2.391, -0.280]$	-0.370 $[-0.877, -0.088]$

Notes: Designed after the pre-specified holdout evaluation showed the opposite directional result; RMT is post-confirmatory exploratory. D: 6 July–5 August 2026; H: 9 June–5 July 2026. Common weights are $(\pi_D + \pi_H)/2$. Brackets give 95% intervals from 2,000 within-window calendar-day bootstrap replicates. State composition, execution incidence, and conditional magnitude allocate $\bar{C}_H - \bar{C}_D$, the raw window-to-window change in aggregate mean correction (not STD_DIFF above, and not β_2). Adding-up holds before rounding.

conditional correction magnitude. Their frequency-weighted product gives aggregate mean correction. Supplementary Appendix G gives the exact formulas and bootstrap design.

Applying Discovery's incidence and magnitude maps to Confirmation's state frequencies overpredicts realized correction. The transport error, defined as realized minus transported mean correction, is -2.212 bp for ETH5 and -1.047 bp for BASE5, with both 95% intervals below zero (Table 3). Raw mean correction nevertheless rises only modestly, from 0.484 to 0.580 bp for ETH5 ($+0.096$ bp) and from 0.107 to 0.143 bp for BASE5 ($+0.036$ bp); the transport error is several times larger in magnitude than this raw change in both pools, which is exactly what it means for a modestly higher aggregate to conceal substantially offsetting internal movement. Giving both maps the same bin weights—the equal-weight average of the two windows' state frequencies—yields Confirmation-minus-Discovery differences of -1.141 bp and -0.525 bp, also with intervals below zero. The sign itself reverses after state standardization: raw aggregate correction is higher in Confirmation than Discovery in both pools, whereas the common-weight Confirmation-minus-Discovery comparison is negative in both. The raw cross-window comparison therefore points in the opposite direction from the comparison obtained after holding the state distribution fixed. Common bins do not equalize within-bin discrepancy distributions, so the standardized comparison remains subject to within-bin composition differences. Discovery represents 30 and Confirmation 27 calendar days (Section 2.1); as elsewhere, these are descriptive-uncertainty clusters, not independent replications.

The bin-level observation mass behind this result (Supplementary Table S5) shows the great majority of mass in both pools and both windows lies below $x = 1$; the open-ended $[5, \infty)$ bin is the single largest compositional shift in the support, growing from 0.06%/0.01% of Discovery mass to 7.6%/5.3% of Confirmation mass for ETH5/BASE5 respectively. This sharp compositional shift into the sparsest, most extreme bin helps explain why the full-support result is materially larger than the $x \leq 2$ result reported below. The Discovery $[5, \infty)$ cells contain 108 of 114 ETH5 blocks and 77 of 92 BASE5 blocks with at least one focal-pool swap; Supplementary Table S6 reports the swap-block count for every bin, window, and pool.

Because $A(\pi, p, m) = \sum_k \pi_k p_k m_k$ is a nonlinear function of the three margins, the change $A_H - A_D$ does not split additively into a composition term, an incidence term, and a magnitude term without a choice of allocation rule. We use the Shapley value to obtain a symmetric, order-independent

allocation of the change across the three margins, following the decomposition logic developed by Shorrocks (2013), rather than a new identification method: for factors π, p, m and all $3! = 6$ orderings σ in which they are swapped from their Discovery to Confirmation values, the contribution of factor j is

$$\Delta_j = \frac{1}{6} \sum_{\sigma} \Big[A(\cdot \mid \sigma \text{ up to and including } j) - A(\cdot \mid \sigma \text{ up to but excluding } j) \Big],$$

averaged over orderings, so that $\Delta_\pi + \Delta_p + \Delta_m = A_H - A_D$ exactly. This is an exact accounting allocation of a realized change, not a causal or structural decomposition. The three components allocate $A_H - A_D$, the raw window-to-window change in aggregate mean correction; they do not allocate the standardized difference reported above, and, as already noted, they do not allocate β_2. The magnitude component Δ_m is a change in the fixed-bin conditional-magnitude map – the average correction among swap blocks within each pre-defined discrepancy cell – and should be read as a coarse-cell descriptive allocation rather than a behavioral change holding exact discrepancy fixed, since discrepancy values can differ within a cell. For ETH5, composition contributes +1.256 bp and magnitude −1.171 bp, with incidence close to zero. BASE5 has composition +0.489 bp, magnitude −0.370 bp and incidence −0.083 bp. Relative to the 5 bp fee tier, ETH5's composition and magnitude contributions are each roughly one-quarter of one fee tier in size (+25% and −23% of 5 bp); BASE5's are smaller, at roughly one-tenth (+10% and −7%). Correction basis points and fee basis points are not the same economic object – one measures movement toward a reference, the other a pool's per-swap charge – so this comparison is for orientation only. These exact allocations concern the change in aggregate mean correction, not the hinge coefficient or its sign reversal. Positive composition and negative magnitude contributions offset in both matched 5 bp pools. Figure 2 (Section 3) displays the underlying incidence and magnitude maps.

4.1 Support Dependence of the Transport Result

The support-cap robustness exercise was designed after the holdout evaluation showed instability. Before estimating the support-cap results, we fixed the ladder $x \leq 1.5, 2, 3, 5$ and Full, retaining complete fixed bins and renormalizing frequencies within each window. The qualitative composition-up, magnitude-down pattern persists after excluding the open-ended $[5, \infty)$ bin in both pools, but is not established on $x \leq 2$. Full support substantially amplifies the differences. The complete ladder and pointwise intervals are in Supplementary Appendix H (Table S8).

At $x \leq 5$, ETH5 has transport error −0.137 bp, standardized difference −0.095 bp, composition +0.217 bp and magnitude −0.085 bp. BASE5 has corresponding estimates of −0.210, −0.106, +0.136 and −0.069 bp. All these intervals exclude zero in the stated direction. The raw-versus-standardized sign reversal also remains at $x \leq 5$: capped mean correction is higher in Confirmation in both pools (+0.123 bp for ETH5, +0.044 bp for BASE5), while the common-weight standardized difference remains negative in both. At $x \leq 3$, ETH5 retains the full pattern; BASE5's standardized-difference and magnitude intervals include zero. The pattern is not established on $x \leq 2$; this does not imply central-map equivalence. No cutoff was searched or threshold estimated. Because caps renormalize weights, Full-minus-cap differences do not allocate an additive contribution to the excluded tail.

This support dependence bears directly on the most natural objection to the headline result: whether the full-support composition-up, magnitude-down pattern is manufactured entirely by the sparsest, most extreme bin. Figure 3 shows it is not manufactured entirely by that bin – the qualitative pattern survives through $x \leq 5$ in both primary pools – but it is support-dependent: it is not established on $x \leq 2$, and full support substantially amplifies the magnitude of the estimates relative to any capped subset.

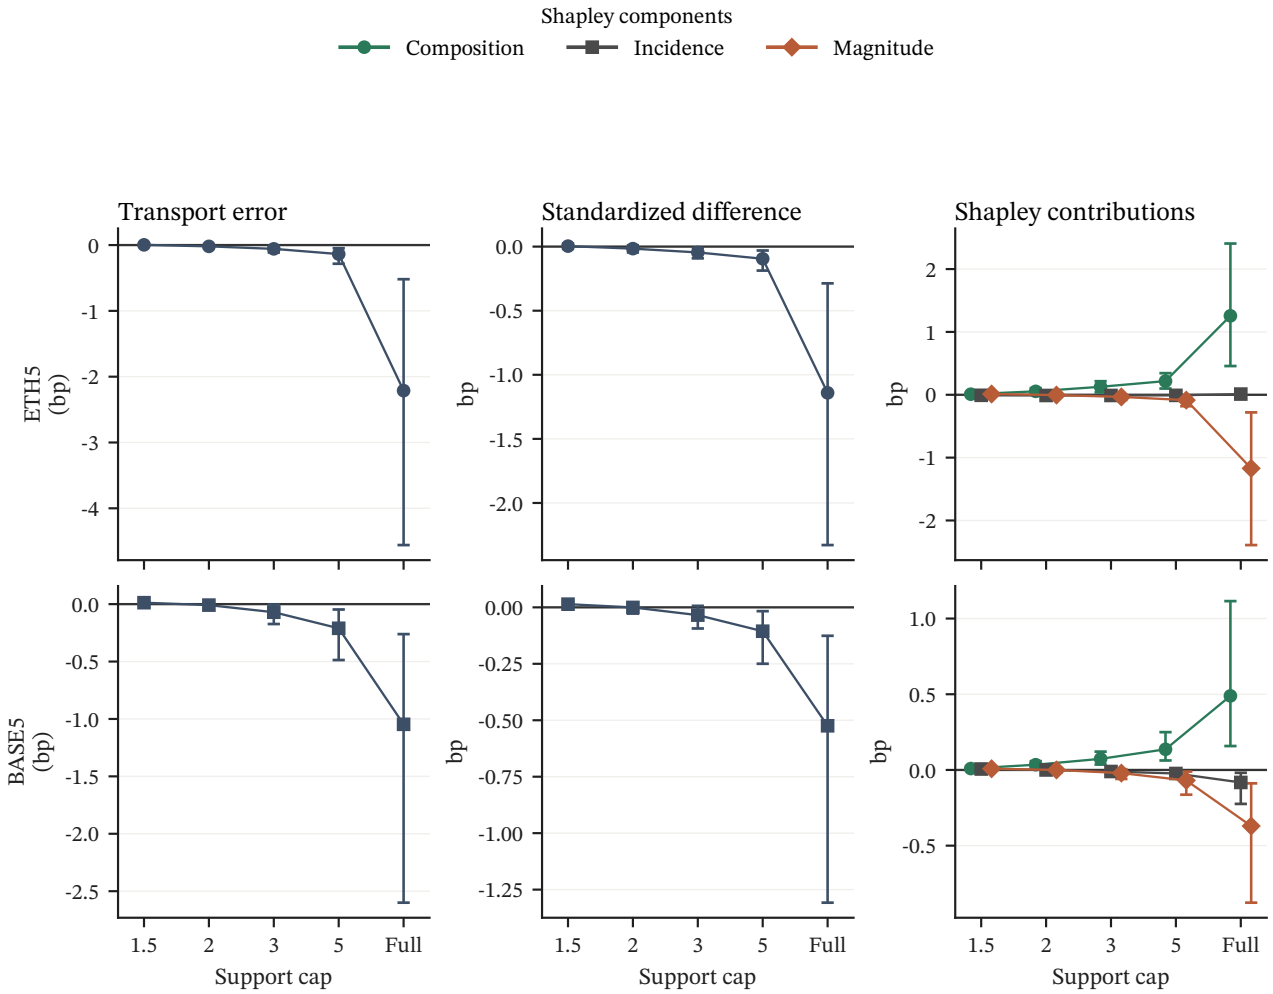


Figure 3: Post-confirmatory exploratory robustness across support caps fixed before this diagnostic was run. Left panels report transport error and standardized difference; right panels report mean-change contributions (Composition, Incidence, Magnitude) from the symmetric Shapley allocation of the change in aggregate mean correction. Lines only connect these fixed cap points; there is no cutoff search, smoothing or threshold estimation. Intervals are pointwise 95% calendar-day bootstrap intervals. Full-support magnitude is strongly amplified with inclusion of the open-ended tail. The qualitative composition-up, magnitude-down pattern persists before that bin in both primary pools, though it is not established on $x \leq 2$. Each cap renormalizes its weights. Horizontal spacing is categorical; small component offsets separate intervals. Panels use separate labeled linear scales.

5 Discussion and Conclusion

The post-confirmatory transport analysis reveals a support-dependent descriptive pattern across the two non-overlapping sample windows, and the π, p, m representation is what makes it visible. Raw mean correction rises only modestly in Confirmation (Section 4), yet its full-support response map yields lower correction under the common bin weights. It records offsetting positive composition and negative conditional-magnitude contributions in both primary 5 bp pools. Incidence contributes little to the ETH5 aggregate change and a smaller negative amount on BASE5. These terms account for the change in aggregate mean correction in the exact allocation; they do not explain or decompose

β_2. The holdout evaluation of the fixed-hinge summary and its support-cap diagnostic (Section 3.1) retain a separate role from the RMT support-dependence result (Section 4.1): the holdout evaluation documents instability in a parsimonious reduced-form summary, while the subsequent decomposition examines cross-window changes in aggregate correction through separately observed composition, execution-incidence, and conditional-magnitude margins.

This reversal is itself the key statistical fact. Aggregate mean correction is $A_w = \sum_k \pi_{wk} p_{wk} m_{wk}$, so a raw comparison across windows can change either because the response map (p and/or m) changes or because the distribution of visited states (π) changes: raw aggregate correction is not composition-invariant, and a modest or positive change in it does not by itself imply that the response map strengthened. Here raw aggregate correction rises in both pools, yet under common state weights the Confirmation-minus-Discovery comparison is negative in both – the composition shift toward the sparsest, most extreme discrepancy states (Section 4; Table S5) coincides with a lower state-standardized response. Statistically, this composition-induced sign reversal is analogous to the logic behind Simpson-type aggregation effects: an unconditional comparison can move in one direction while a comparison under common state weights moves in the other. We use the analogy only for the aggregation logic; the analysis does not assert the stronger within-cell reversal conditions associated with the classical Simpson paradox, and the role of standardization here is descriptive – it makes the two response maps comparable under common state exposure, not a causal or treatment-effect adjustment. The finance implication is concrete: a higher headline correction statistic can arise because the market visits different discrepancy states, even when its state-standardized response is lower, so a cross-window, cross-regime, or cross-venue comparison of an aggregate price-adjustment statistic can be misleading unless one first asks whether the state distribution itself changed. The decomposition also separates an extensive execution margin (p, whether a focal-pool execution occurs) from an intensive conditional-correction margin (m, correction conditional on execution); in this sample the incidence contribution is quantitatively small relative to composition and magnitude, but conditioning only on executed blocks would still omit a distinct component of aggregate adjustment.

Two prior lines of work supply the knowledge this paper's evidence adds to. AMM fee/arbitrage theory establishes the economic incentive and the fee-mediated price-tracking machinery: Milionis et al. (2023) show why a discrepancy can persist within a no-arbitrage band until an arbitrageur's trade corrects it, Li et al. calibrate a fixed-fraction correction-strength object in a stylized price-tracking model (Li et al., 2026), Najnudel et al. provide adjacent theoretical context for fee-mediated AMM price dynamics (Najnudel et al., 2024), and Fukasawa et al. characterize arbitrage-free reserve processes and loss-versus-rebalancing within a continuous-time no-arbitrage construction (Fukasawa et al., 2025). That literature provides theoretical mechanisms linking fees, arbitrage incentives, liquidity conditions, and AMM price tracking; our evidence adds that realized execution occurrence and realized same-block correction are separate empirical margins – our conditional mean correction m and observed execution incidence p are observed conditional moments, not estimates of a structural correction-strength parameter or an arbitrage-arrival probability – and that a single global discrepancy-response shape need not transport across non-overlapping windows of the same execution-mediated setting. AMM oracle and state-construction questions have a distinct lineage (Angeris and Chitra, 2020), and our strict-prior external reference plays a different role, as an empirical measurement input for a same-block correction statistic rather than a term in a theoretical no-arbitrage construction.

State-dependent microstructure evidence and the standardization literature together establish that aggregate comparisons across samples depend on the distribution of states, not only on the conditional response within them: Hasbrouck (1991) and Dufour and Engle (2000) (Introduction) show price response depends on market state, direct standardization (Kitagawa, 1955) and semiparametric counterfactual reweighting (DiNardo et al., 1996) formalize how to separate composition from conditional structure, and the invariance literature asks whether transaction-cost and bet-size relationships stay

stable once rescaled to business time (Kyle and Obizhaeva, 2016). Our evidence adds a concrete execution-mediated case in which the raw cross-window comparison and the common-state-weight comparison actually have opposite signs, rather than merely differing in magnitude: the present exercise decomposes, without rescaling, why an aggregate response differs or appears stable across two windows of the same setting.

Several limits are consequential and setting-specific rather than limits on the general measurement problem. The primary strict-prior reference uses only the specified two direct-USDC venues and requires both to be timely; it is not a latent efficient price. Swap occurrence is endogenous, and same-block correction measures a native-price change rather than realized trading surplus. Discovery and Confirmation are non-overlapping and Confirmation is earlier in calendar time; RMT transports one window's response map onto the other's state frequencies retrospectively rather than forecasting either window, and it does not assume the response map is stationary across them – whether that map is stable across the two windows is itself the empirical question RMT and the holdout evaluation both address, and both answer it in the negative in this sample. The longer AMM-native historical panel improves context but contains no accepted high-frequency external reference and cannot define x, C, RMT, or the holdout regression; BASE1 remains secondary throughout. That panel places the full high-frequency study window in broadly typical portions of each primary pool's own liquidity and native-volatility history (Section 2.1; Supplementary Appendix F), but because those diagnostics are not constructed separately for Discovery and Confirmation, they do not identify the source of the cross-window shift into the sparsest discrepancy states documented in Section 4: the data establish that the distribution of discrepancy states changed sharply across the two windows, not which underlying market mechanism produced that shift. The sharp shift in state composition is economically compatible with several channels emphasized in AMM microstructure theory: fee-mediated arbitrage incentives (Milionis et al., 2023), execution and price-tracking conditions (Li et al., 2026), and liquidity provision on blockchain-based decentralized exchanges (Capponi and Jia, 2025) jointly shape whether and how a discrepancy is corrected. The present design observes the resulting state, execution-incidence, and conditional-correction margins, but does not separately identify which of these channels generated the cross-window shift. The RMT support-dependence result is bounded by the bin-mass shift that drives full-support amplification (Table S5), the range over which the pattern is and is not established (Section 4.1), and the uncontrolled within-bin composition that coarse fixed bins cannot rule out. The value of the two studied pools is that every margin is directly observable without a structural model, not that ETH5 and BASE5 are a representative sample of AMMs or of execution-gated markets generally: this is a demonstration that offsetting compositional change and changes in the fixed-bin conditional-magnitude map can occur and are directly measurable in such a market, not a claim about how often or how strongly this occurs elsewhere. Agreement between ETH5 and BASE5 is itself descriptive rather than confirmatory in a statistical sense – the two pools share overlapping calendar dates and broader market conditions, and the calendar-day bootstrap used throughout does not by itself establish independence across pools any more than across days.

When price updating requires execution, aggregate correction that rises only modestly can conceal offsetting changes in the states a market visits (π), how often execution occurs in those states (p), and the correction that follows conditional on execution (m); a higher raw aggregate correction can coexist with a lower state-standardized response because the mix of visited states changed. The two primary 5 bp Uniswap v3 pools studied make this decomposition exact rather than approximate, because their execution-mediated native-price update rule eliminates the non-execution native-price adjustment branch entirely. In the observed windows, the composition-up, magnitude-down pattern persists through $x \leq 5$ in both primary pools, is not established on $x \leq 2$, and is substantially amplified on full support by a large compositional shift into the sparsest, open-ended bin. The empirical pattern is

specific to the studied pools and windows, but the measurement exercise shows how an execution-gated setting can make otherwise hidden state-composition, execution-incidence, and conditional-magnitude changes directly observable, without requiring a structural model of arbitrage arrival.

Statements and Declarations

Funding. Wen-Ting Wang received funding from the National Science and Technology Council, Taiwan (Grant ID NSTC 113-2118-M005-005-MY2). The funding source had no involvement in the study design; in the collection, analysis and interpretation of data; in the writing of the report; or in the decision to submit the article for publication.

Competing Interests. The author declares no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author Contributions. Wen-Ting Wang: Conceptualization, Methodology, Formal analysis, Software, Validation, Investigation, Data curation, Visualization, Writing—original draft, Writing—review and editing.

Data Availability. The data used in this study were derived from publicly available Ethereum and Base blockchain records. The external reference is built from Bybit and OKX best bid and offer archives, obtained directly from the venues; these exchange-sourced quote archives may be subject to third-party redistribution restrictions, and the manuscript reports their construction rule in full so that the measurement can be reconstructed from the same public sources.

Code Availability. The analysis code is publicly available at https://github.com/egpivo/execution-mediated-price-adjustment.

Ethics Approval. This study uses publicly observable blockchain records and historical exchange quote archives, and did not involve human participants or animal subjects. Ethics approval and informed consent were not required.

Declaration of Generative AI and AI-Assisted Technologies in the Writing Process. During the preparation of this work, the author used Claude Code to assist with manuscript organization, language editing, LaTeX preparation, code-level formatting of data visualizations, and submission-compliance checks. Empirical results and specifications were produced by author-controlled code, and all figures were rendered deterministically from fixed numerical inputs using author-reviewed Matplotlib scripts; no generative image model was used. After using this service, the author reviewed and edited the content as needed and takes full responsibility for the content of the article.

References

Guillermo Angeris and Tarun Chitra. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91, 2020. doi: 10.1145/3419614.3423251. Also arXiv:2003.10001.

Jun Aoyagi and Yuki Ito. Coexisting exchange platforms: Limit order books and automated market makers. *Journal of Political Economy Microeconomics*, 3(3):611–648, 2025. doi: 10.1086/732831.

Yashar H. Barardehi and Dan Bernhardt. Revisiting the U-shaped patterns in volatility and price impacts: Novel results using trade-time estimates. *Journal of Financial Markets*, 74:100971, 2025. doi: 10.1016/j.finmar.2025.100971.

Emilio Barucci, Giuseppe Giuffra Moncayo, and Daniele Marazzina. Market impact and efficiency in cryptoassets markets. *Digital Finance*, 5:519–562, 2023. doi: 10.1007/s42521-023-00095-9.

Alexander Brauneis, Roland Mestel, Ryan Riordan, and Erik Theissen. The anatomy of a fee change—evidence from cryptocurrency markets. *Journal of Empirical Finance*, 67:152–167, 2022a. doi: 10.1016/j.jempfin.2022.03.003.

Alexander Brauneis, Roland Mestel, Ryan Riordan, and Erik Theissen. Bitcoin unchained: Determinants of cryptocurrency exchange liquidity. *Journal of Empirical Finance*, 69:106–122, 2022b. doi: 10.1016/j.jempfin.2022.08.004.

Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4):1547–1621, 2015. doi: 10.1093/qje/qjv027.

Nils Bundi, Ching-Lin Wei, and Khaldoun Khashanah. Optimal trade execution in cryptocurrency markets. *Digital Finance*, 6(2):283–318, 2024. doi: 10.1007/s42521-023-00103-y.

Agostino Capponi and Ruizhe Jia. Liquidity provision on blockchain-based decentralized exchanges. *The Review of Financial Studies*, 38(10):3040–3085, 2025. doi: 10.1093/rfs/hhaf046.

John DiNardo, Nicole M. Fortin, and Thomas Lemieux. Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach. *Econometrica*, 64(5):1001–1044, 1996. doi: 10.2307/2171954.

Alfonso Dufour and Robert F. Engle. Time and the price impact of a trade. *The Journal of Finance*, 55(6):2467–2498, 2000. doi: 10.1111/0022-1082.00297.

David Easley, Maureen O'Hara, Songshan Yang, and Zhibai Zhang. Microstructure and market dynamics in crypto markets. *Journal of Financial Markets*, page 101071, 2026. doi: 10.1016/j.finmar.2026.101071.

Thierry Foucault, Roman Kozhan, and Wing Wah Tham. Toxic arbitrage. *The Review of Financial Studies*, 30(4):1053–1094, 2017. doi: 10.1093/rfs/hhw103.

Masaaki Fukasawa, Basile Maire, and Marcus Wunsch. Liquidity provision of utility indifference type in decentralized exchanges. *Digital Finance*, 7:255–273, 2025. doi: 10.1007/s42521-025-00128-5.

Joel Hasbrouck. Measuring the information content of stock trades. *The Journal of Finance*, 46(1):179–207, 1991. doi: 10.1111/j.1540-6261.1991.tb03749.x.

Evelyn M. Kitagawa. Components of a difference between two rates. *Journal of the American Statistical Association*, 50(272):1168–1194, 1955. doi: 10.1080/01621459.1955.10501299.

Albert S. Kyle and Anna A. Obizhaeva. Market microstructure invariance: Empirical hypotheses. *Econometrica*, 84(4):1345–1404, 2016. doi: 10.3982/ECTA10486.

Alfred Lehar and Christine A. Parlour. Decentralized exchange: The uniswap automated market maker. *The Journal of Finance*, 80(1):321–374, 2025. doi: 10.1111/jofi.13405.

Peihao Li, Nadia Dahmani, and Wenqi Cai. Arbitrage and the stability of AMM price tracking. arXiv:2605.06060 [cs.CE], 2026.

Igor Makarov and Antoinette Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319, 2020. doi: 10.1016/j.jfineco.2019.07.001.

Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. Automated market making and arbitrage profits in the presence of fees. arXiv:2305.14604 [q-fin.MF], 2023. Current version July 23, 2025.

Joseph Najnudel, Shen-Ning Tung, Kazutoshi Yamazaki, and Ju-Yi Yen. An arbitrage driven price dynamics of automated market makers in the presence of fees. *Frontiers of Mathematical Finance*, 3(4):560–571, 2024. doi: 10.3934/fmf.2024018.

Shulin Shen, Yixuan Zhang, and Eric Zivot. Improving information leadership share for measuring price discovery. *Journal of Empirical Finance*, 83:101638, 2025. doi: 10.1016/j.jempfin.2025.101638.

Anthony F. Shorrocks. Decomposition procedures for distributional analysis: A unified framework based on the shapley value. *The Journal of Economic Inequality*, 11(1):99–126, 2013. doi: 10.1007/s10888-011-9214-z.

Shen-Ning Tung and Tai-Ho Wang. A mathematical framework for modelling CLMM dynamics in continuous time. *Digital Finance*, 8:39, 2026. doi: 10.1007/s42521-026-00198-z.

Philippe van der Beck, Lorenzo Bretscher, and Julie Zhiyu Fu. Rethinking volume: Has the market become more liquid? Technical Report 26-003, Harvard Business School, 5 2026. URL https://priceimpactbound.github.io/rethinkingvolume/rethinking-volume.pdf. Revised May 2026; originally circulated as "Rethinking Volume: The Illusion of Liquidity".

Supplementary Material

Price Adjustment When Updating Requires Execution:
Evidence from Automated Market Makers

This supplement retains measurement details, historical context, full bin-level results and exploratory diagnostics. The pre-specified holdout test and the exploratory transport analyses have distinct inferential status.

A Reference Construction and Quality Assurance

The primary reference (Tier A in the calibration records) uses direct-USDC BBOs from Bybit ETHUSDC and OKX ETH–USDC. For each block, the implementation selects a quote strictly prior to the block timestamp from each venue, requires both selected quotes to be at most three seconds stale, and computes the equal-weight log midpoint. It does not substitute a single venue, ETHUSDT, or a USDC–USDT basis proxy when a direct quote is absent.

The Tier-B and Tier-C comparisons in Table S1 are calibration diagnostics only. Their purpose is to record why the lower-cost constructions are not used to expand the correction panel. The Tier-A reference-flag row counts and shares are reported in Table 1.

Table S1: Reference-tier decision

Tier	Construction	Use in this paper
A	Bybit ETHUSDC + OKX ETH–USDC; strict prior; both $\leq$ 3s	Benchmark correction sample
B	USD proxy plus basis adjustment	Fails calibration; not used
C	ETHUSDT plus USDC–USDT basis	Calibration only; not used

Notes: Full stored calibration metrics are reported in Appendix Table S2. Tier A has no one-venue fallback.

Table S2: Full stored reference-tier calibration metrics and agreement semantics

Tier	Both retained N	MAE (bp)	p95 (bp)	$x < 1$ side full overlap agreement	$x < 1$ band-side conditional agreement	Within-$x < 1$ sub-bin aggregate agreement	$[2, 3)$ bin-side agreement	$[3, 5)$ bin-side agreement	$x \geq 5$ bin-side agreement
A	–	0.00	0.00	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
B	120,225	20.25	73.13	35.6%	–	5.8%	1.5%	1.4%	0.7%
C	358,376	0.64	1.63	92.7%	90.5%	55.9%	67.8%	81.2%	99.3%

Notes: Side agreement compares $\mathbf{1}\{x_{candidate} < 1\}$ and $\mathbf{1}\{x_{TierA} < 1\}$ over all both-retained rows. The conditional band-side value is the earlier Tier-C stored calculation restricted to rows for which at least one series is below one; it is not the full-overlap side agreement and was not stored for Tier B. The within-$x < 1$ sub-bin value is the count-weighted aggregate of the three stored union-based sub-bin records ($[0, .5)$, $[.5, .75)$, $[.75, 1)$); it is not an exact mutually exclusive multi-class agreement. Other bin-side columns use the same stored union-based convention. The calibration artifact records both-retained counts but does not separately store a Tier-B/Tier-C freshness rate.

Table S3 reports the distribution of reference quote age – the gap between the last-prior quote from each venue and the focal block timestamp – separately by venue and chain, over the full candidate-block population before the freshness gate is applied. Because the reference is fixed strictly prior to the block timestamp rather than resampled at the moment of the focal swap, and because the external market can move within a block, same-block correction C is a change in distance to this fixed strict-prior benchmark, not a change relative to the external price contemporaneous with the swap itself; no swap-level or transaction-time reference is constructed or claimed.

Table S3: Reference quote age (seconds between the last-prior quote and the focal block timestamp), 9 Jun–5 Aug 2026

Venue	Chain	p50	p75	p90	p95	p99	Share ≤1s	Share ≤3s
Bybit	Ethereum	0.107	0.301	0.711	1.171	2.490	93.8%	99.5%
OKX	Ethereum	0.155	0.507	1.371	2.405	5.984	86.3%	96.3%
Bybit	Base	0.105	0.297	0.714	1.178	2.460	93.7%	99.5%
OKX	Base	0.151	0.502	1.335	2.301	5.915	86.6%	96.5%

Notes: computed over the full 9 Jun–5 Aug candidate-block population (393,305 Ethereum blocks; 2,393,726 Base blocks), before the three-second freshness gate and both-venues-required rule that determine reference-flag status (Table 1). All four venue-chain age distributions are comfortably inside the three-second gate at the median and 90th percentile; OKX ages run systematically higher than Bybit's, consistent with the gate's role in determining which candidate blocks receive the reference flag.

B On-chain Price Reconstruction

The price panel is reconstructed deterministically from on-chain events. Pre- and post-state native prices are associated with the same block; the focal-pool swap indicator is computed from that pool's swap events. A block without a focal-pool swap is retained in the denominator of the extensive margin and has zero correction under the native-price outcome definition. The reconstruction fixes the price and execution measurements used in the analysis.

The focal pools are ETH5, `0x88e6a0c2ddd26feeb64f039a2c41296fcb3f5640` on Ethereum (chain ID 1); BASE5, `0xd0b53d9277642d899df5c87a3966a349a798f224` on Base (chain ID 8453); and BASE1, `0xb4cb800910b228ed3d0834cf79d697127bbb00e5` on Base. Each pool's identity was verified by direct on-chain calls to the live contract: ETH5 reports the Uniswap v3 Ethereum factory `0x1f98431c8ad98523631ae4a59f267346ea31f984`, and BASE5 and BASE1 report the Uniswap v3 Base factory `0x33128a8fc17869897dce68ed026d694621f6fdfd`. The fee and tick-spacing settings returned by the contracts are 500 pips with tick spacing 10 for ETH5 and BASE5, and 100 pips with tick spacing 1 for BASE1, matching the 5 bp and 1 bp tiers used throughout. Discriminating calls for concentrated-liquidity forks revert on all three pools. Event ordering and the reconstruction quality checks follow the procedure described above and in Appendix A.

C Twelve-Second Outcome as a Secondary Appendix

The primary manuscript uses same-block correction. A 12-second outcome appeared in the previous theory-note illustration and remains a distinct fixed-time tracking object. It is not an alias for same-block correction, is not pooled with the primary estimand, and is not used in the confirmation result reported in Table 2.

The distinction matters because a fixed-time horizon may span later state changes not attributable to the focal block. The primary analysis therefore keeps the pre-state, focal swap, and post-state in a single block. This makes the swap indicator the exact on/off switch in the decomposition; it would not be one for the 12-second object.

D Diagnostic Status

The support caps and symmetric windows in Table S4 were analyzed only after the holdout evaluation returned the opposite sign. They use fixed support definitions and are presented to document sensitivity. They are not discovery-window re-searches and not a new holdout evaluation.

After that holdout evaluation, we evaluated a fixed diagnostic ladder of support caps and local

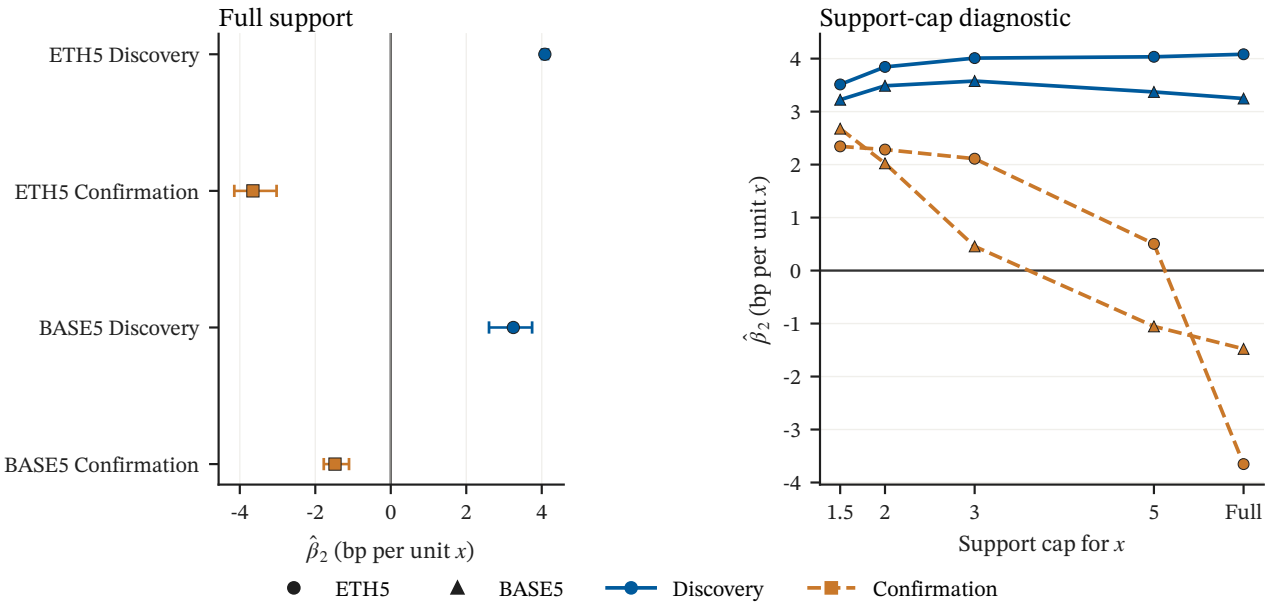


Figure S1: Global hinge estimates and post-confirmatory support diagnostic. Both panels report $\hat{\beta}_2$ in basis points per unit x. The left panel uses the full support; the right panel reports the cap ladder fixed before this diagnostic was run. Marker shape encodes market and color/line style encode window. The diagnostic does not alter the holdout-window directional result: in both primary pools the estimated slope change has the opposite sign from the Discovery-motivated positive prediction.

windows to characterize the result's *support dependence*. The diagnostic characterizes support dependence and does not alter the primary holdout result or identify a new breakpoint. The fixed-$x = 1$ hinge is positive in both markets and both periods in local windows around $x = 1$; capped estimates through $x \leq 2$ are also positive. ETH5's capped estimates remain positive through $x \leq 5$; its full-support estimate differs sharply from them, and the two sets of estimates differ only in whether the open-ended $[5, \infty)$ bin is retained. BASE5's estimate declines as broader support is retained and is negative over the broader capped and full-support sets. No breakpoint or threshold is estimated in either market. Local/capped behavior and full-support behavior are different empirical objects: the diagnostic shows where the global summary does not transport, not a local law that supersedes it. These statements describe the support diagnostic, not a newly estimated boundary or a universal $x = 1$ rule.

Table S4: Post-confirmatory support and local-window diagnostic

Market	Period	$x \leq 1.5$	$x \leq 2$	$x \leq 3$	$x \leq 5$	Full	$.5 \leq x \leq 1.5$	$.75 \leq x \leq 1.25$
ETH5	Discovery	3.513	3.843	4.010	4.034	4.082	3.020	1.719
ETH5	Confirmation	2.344	2.283	2.111	0.504	-3.653	1.801	1.228
BASE5	Discovery	3.224	3.487	3.577	3.371	3.247	2.625	1.668
BASE5	Confirmation	2.678	2.024	0.457	-1.053	-1.478	1.981	1.370

Notes: Entries are fixed-$x = 1$ hinge slope changes in bp per unit x. All support caps and local windows fixed before this diagnostic was run are reported. This diagnostic was conducted after the holdout evaluation showed the opposite directional result; it characterizes support dependence and does not alter that result.

Table S5: Bin-level observation mass (compact; full per-bin p, m detail in Table S6)

	ETH5				BASE5			
Bin	N_D	π_D	N_H	π_H	N_D	π_D	N_H	π_H
[0,.25)	38,432	20.9%	31,207	18.3%	264,361	23.8%	238,171	22.6%
[.25,.5)	37,720	20.5%	30,937	18.1%	268,626	24.1%	234,484	22.3%
[.5,.75)	36,157	19.7%	30,221	17.7%	265,851	23.9%	227,243	21.6%
[.75,1)	32,343	17.6%	26,306	15.4%	211,310	19.0%	163,950	15.6%
[1,1.25)	19,602	10.7%	15,771	9.2%	76,855	6.9%	62,979	6.0%
[1.25,1.5)	9,098	4.9%	8,393	4.9%	16,265	1.5%	22,479	2.1%
[1.5,2)	6,919	3.8%	8,069	4.7%	7,040	0.6%	16,890	1.6%
[2,3)	3,016	1.6%	4,621	2.7%	1,613	0.1%	13,809	1.3%
[3,5)	592	0.3%	2,103	1.2%	349	0.0%	16,885	1.6%
[5,∞)	114	0.1%	12,885	7.6%	92	0.0%	55,322	5.3%

Notes: N is analytic block count (finite x and C; Table 1); π is bin share of window mass. D=Discovery (30 represented days), H=Confirmation (27 days). The open-ended $[5, \infty)$ bin grows from 0.06% (ETH5) / 0.01% (BASE5) of Discovery mass to 7.6% / 5.3% of Confirmation mass – the single largest compositional shift in the support, and it helps explain the full-support amplification.

E Full Bin-Level Decomposition

Table S5 reports the bin-level observation mass behind the response-map transport result discussed in the main text (Section 4).

Table S6: Exact decomposition by fixed discrepancy bin

Market/window	x bin	N	N_S	$P(S \mid x)$	$E[C \mid S, x]$ (bp)	$E[C \mid x]$ (bp)	Product (bp)	Residual
ETH5 Discovery	[0,0.25)	38,432	10,359	0.270	-0.226	-0.061	-0.061	0.00e+00
	[0.25,0.5)	37,720	10,672	0.283	-0.049	-0.014	-0.014	0.00e+00
	[0.5,0.75)	36,157	11,345	0.314	0.097	0.030	0.030	0.00e+00
	[0.75,1.0)	32,343	13,920	0.430	0.411	0.177	0.177	0.00e+00
	[1.0,1.25)	19,602	13,613	0.694	1.024	0.711	0.711	0.00e+00
	[1.25,1.5)	9,098	8,032	0.883	2.088	1.844	1.844	0.00e+00
	[1.5,2.0)	6,919	6,346	0.917	3.766	3.454	3.454	0.00e+00
	[2.0,3.0)	3,016	2,827	0.937	7.002	6.564	6.564	0.00e+00
	[3.0,5.0)	592	563	0.951	13.128	12.485	12.485	0.00e+00
	[5.0,inf)	114	108	0.947	29.626	28.067	28.067	-4.34e-19
ETH5 Confirmation	[0,0.25)	31,207	10,327	0.331	-0.447	-0.148	-0.148	0.00e+00
	[0.25,0.5)	30,937	10,729	0.347	-0.138	-0.048	-0.048	8.47e-22
	[0.5,0.75)	30,221	11,689	0.387	0.135	0.052	0.052	0.00e+00
	[0.75,1)	26,306	12,841	0.488	0.604	0.295	0.295	0.00e+00
	[1,1.25)	15,771	10,598	0.672	1.195	0.803	0.803	0.00e+00
	[1.25,1.5)	8,393	6,910	0.823	2.099	1.728	1.728	0.00e+00
	[1.5,2)	8,069	7,243	0.898	3.377	3.031	3.031	5.42e-20
	[2,3)	4,621	4,325	0.936	5.625	5.265	5.265	1.08e-19
	[3,5)	2,103	2,044	0.972	6.682	6.494	6.494	1.08e-19
	[5,inf)	12,885	12,871	0.999	0.468	0.467	0.467	-6.78e-21
BASE5 Discovery	[0,0.25)	264,361	40,522	0.153	-0.040	-0.006	-0.006	0.00e+00
	[0.25,0.5)	268,626	41,534	0.155	0.018	0.003	0.003	0.00e+00
	[0.5,0.75)	265,851	46,342	0.174	0.106	0.018	0.018	-2.12e-22
	[0.75,1.0)	211,310	55,956	0.265	0.342	0.091	0.091	1.69e-21
	[1.0,1.25)	76,855	43,810	0.570	0.805	0.459	0.459	6.78e-21
	[1.25,1.5)	16,265	14,854	0.913	1.714	1.565	1.565	2.71e-20
	[1.5,2.0)	7,040	6,668	0.947	3.161	2.994	2.994	0.00e+00
	[2.0,3.0)	1,613	1,541	0.955	5.902	5.638	5.638	0.00e+00
	[3.0,5.0)	349	298	0.854	10.230	8.735	8.735	1.08e-19
	[5.0,inf)	92	77	0.837	19.297	16.151	16.151	-2.17e-19
BASE5 Confirmation	[0,0.25)	238,171	44,596	0.187	-0.067	-0.013	-0.013	2.12e-22
	[0.25,0.5)	234,484	45,323	0.193	0.021	0.004	0.004	-5.29e-23
	[0.5,0.75)	227,243	51,126	0.225	0.142	0.032	0.032	0.00e+00

Table S6 (continued)

Market/window	x bin	N	N_S	$P(S \mid x)$	$E[C \mid S, x]$ (bp)	$E[C \mid x]$ (bp)	Product (bp)	Residual
	[0.75,1)	163,950	54,882	0.335	0.411	0.138	0.138	1.69e-21
	[1,1.25)	62,979	39,136	0.621	0.907	0.564	0.564	0.00e+00
	[1.25,1.5)	22,479	18,532	0.824	1.669	1.376	1.376	0.00e+00
	[1.5,2)	16,890	12,776	0.756	2.379	1.800	1.800	0.00e+00
	[2,3)	13,809	8,000	0.579	2.138	1.239	1.239	0.00e+00
	[3,5)	16,885	7,975	0.472	0.853	0.403	0.403	0.00e+00
	[5,inf)	55,322	26,763	0.484	0.071	0.034	0.034	-4.24e-22

Notes: $C = |\log X^* - \log P^{pre}| - |\log X^* - \log P^{post}|$ is same-block signed correction. S indicates at least one focal-pool swap. N_S is the swap-block count; N_S/N reconstructs $P(S \mid x)$ exactly before rounding for display. The product equals the mean by construction because no focal-pool swap implies $C = 0$ exactly. Discovery is exploratory; confirmation rows are a non-overlapping pre-discovery sample.

The separate margin plots in Figures S2 and S3 retain the bin profiles on a logarithmic horizontal display; the combined response maps appear in Figure 2 in the main text (Section 3).

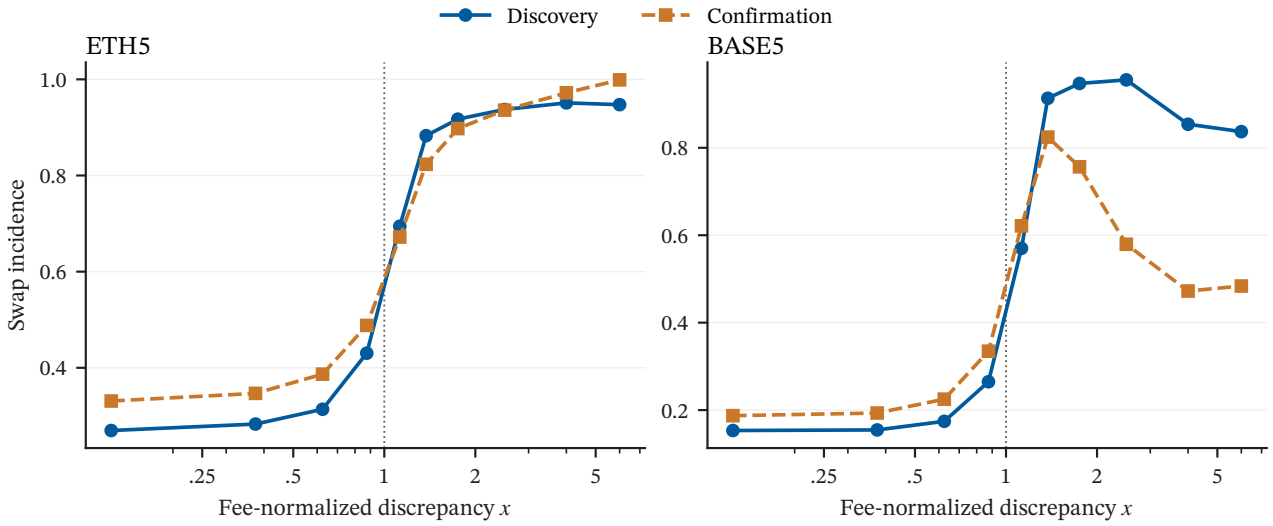


Figure S2: Observed focal-pool swap incidence by fixed discrepancy bin. Discovery is exploratory; confirmation is an earlier, non-overlapping window. Lines connect fixed-bin summaries only.

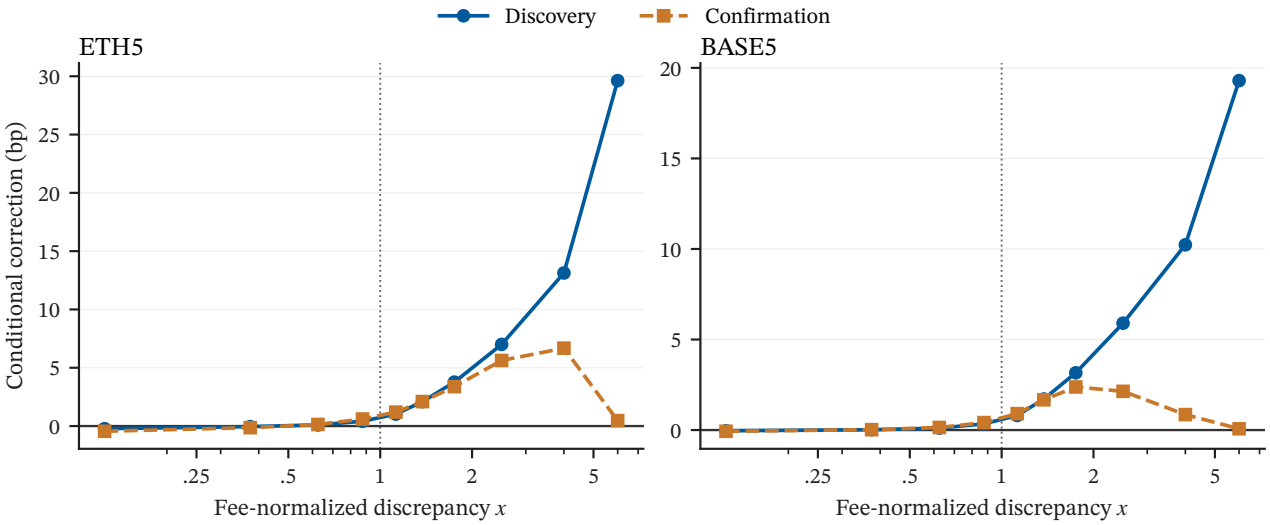


Figure S3: Observed same-block correction conditional on a focal-pool swap, by fixed discrepancy bin. This is the intensive margin of the exact decomposition.

F Historical Context and Scope

The high-frequency correction variables require the primary strict-prior BBO reference and are therefore limited to 57 days. A separate AMM-native panel provides context over the available histories: event reconstruction begins on 5 May 2021 for ETH5, 6 September 2023 for BASE5, and 16 September 2023 for BASE1; the daily series used in Table S7 starts on the following full UTC day. This panel contains no external reference, x, or C.

For each pool, the table places the high-frequency window's mean daily AMM-native observable in its own full-history distribution. The labels are descriptive, not estimated regimes: broadly typical means that all six substantive metrics lie between the 10th and 90th percentiles, while mixed regime records at least one substantive tail value. ETH5 and BASE5 meet the former description. BASE1 is mixed because its Mint/Burn count is at the 98.3rd percentile. We exclude median seconds between swaps and pool age because they are structural timestamp or end-of-sample artifacts, not comparable behavioral context metrics.

Table S7: Historical position of the high-frequency window

Pool	Swap-block percentile	Swap count percentile	Volume proxy percentile	Liquidity percentile	Mint/Burn percentile	1h RV percentile	Classification
ETH5	27.5	41.2	21.3	26.2	38.3	37.6	Broadly typical
BASE5	31.1	42.1	34.2	28.4	27.1	32.2	Broadly typical
BASE1	45.7	58.1	62.0	63.7	98.3	26.8	Mixed regime

Notes: Percentiles place the 57-day window's mean daily AMM-native observable in the pool's own full-history daily distribution. No external reference is used; x and C are not defined here. Median seconds-between-swaps and pool age are excluded because they are structural timestamp/end-of-sample artifacts.

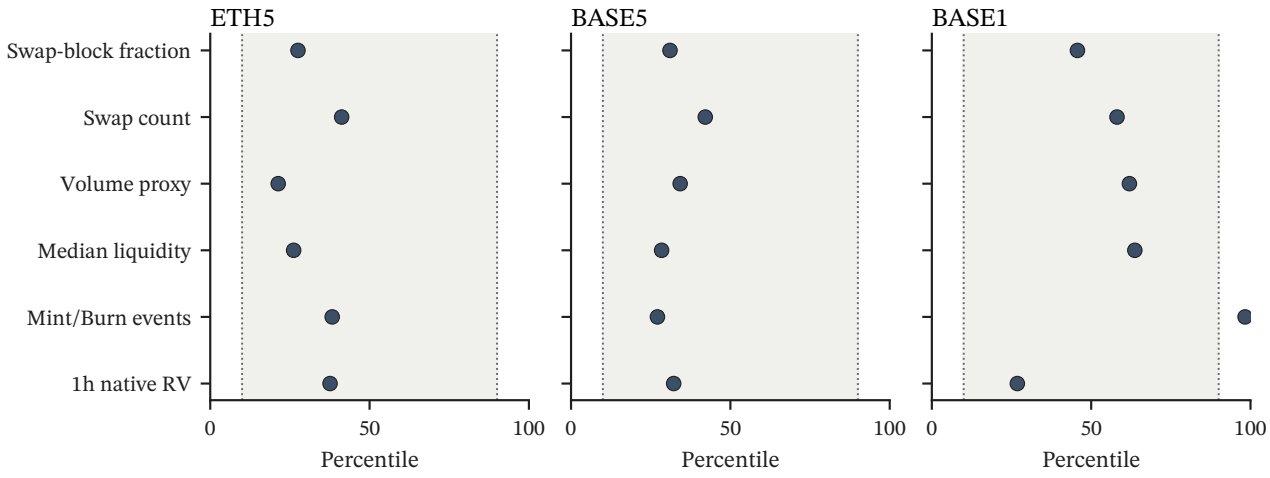


Figure S4: AMM-native historical context. Points locate the high-frequency window's mean daily AMM-native observables within each pool's own history; the shaded band is the descriptive 10th–90th-percentile range. These series do not use an external reference and do not define correction variables.

G Response-Map Transport Methods

RMT was designed after the pre-specified holdout test's result and is post-confirmatory exploratory. Discovery D covers 6 July–5 August 2026; Confirmation H covers 9 June–5 July 2026. The labels describe analysis roles rather than calendar order. Retain the existing high-frequency eligibility filter and finite x, C, including non-swap blocks with $C = 0$. The primary-market counts and bin summaries are those in Table S6: Discovery has 30 represented UTC days and Confirmation 27. No BASE1 RMT is included.

For each primary market, let D denote Discovery and H Confirmation. The labels refer to analytical roles: Confirmation is the earlier calendar window, so $D \to H$ is a retrospective transport exercise, not a prospective forecast. Using the existing fixed discrepancy bins b, define

$$\pi_w(b) = P_w(x \in b), \qquad p_w(b) = P_w(S = 1 \mid x \in b), \qquad m_w(b) = \mathbb{E}_w[C \mid S = 1, x \in b].$$

Writing $\bar{C}_w = \mathbb{E}_w[C]$, the accounting identity gives $\bar{C}_w = \sum_b \pi_w(b) p_w(b) m_w(b)$. The Discovery response map applied to Confirmation's state composition is

$$C_{H \leftarrow D} = \sum_b \pi_H(b) p_D(b) m_D(b), \qquad TE_{D \to H} = \bar{C}_H - C_{H \leftarrow D}. \tag{S1}$$

Use the same bin weights for both response maps, $\bar{\pi}(b) = \frac{1}{2}[\pi_D(b) + \pi_H(b)]$, and define

$$\begin{aligned} C_{D,\mathrm{std}} &= \sum_b \bar{\pi}(b) p_D(b) m_D(b), \quad C_{H,\mathrm{std}} = \sum_b \bar{\pi}(b) p_H(b) m_H(b), \\ \mathrm{STD_DIFF} &= C_{H,\mathrm{std}} - C_{D,\mathrm{std}}. \end{aligned} \tag{S2}$$

We use a symmetric Shapley allocation to summarize the contributions of state composition, execution incidence, and conditional magnitude. Let $F(\pi, p, m) = \sum_b \pi(b) p(b) m(b)$. Averaging marginal changes over all six orders of replacing Discovery factors by Confirmation factors gives

$$\Delta = F(\pi_H, p_H, m_H) - F(\pi_D, p_D, m_D) = \Delta_\pi + \Delta_p + \Delta_m. \tag{S3}$$

The components allocate changes in mean correction to state composition, the execution-incidence map, and the conditional-magnitude map, respectively. The allocation includes interactions through the six-order average.

The fixed bins are

$$[0,.25),\ [.25,.5),\ [.5,.75),\ [.75,1),\ [1,1.25),$$
$$[1.25,1.5),\ [1.5,2),\ [2,3),\ [3,5),\ [5,\infty).$$

For retained counts $n_w(b)$, swap counts $n_{S,w}(b)$, and summed swap-block correction $c_w(b)$, estimate

$$\widehat{\pi}_w(b) = \frac{n_w(b)}{\sum_j n_w(j)}, \qquad \widehat{p}_w(b) = \frac{n_{S,w}(b)}{n_w(b)}, \qquad \widehat{m}_w(b) = \frac{c_w(b)}{n_{S,w}(b)}.$$

Equations (S1)–(S3) use these estimates and all bins, without trimming. The common weights $\bar{\pi} = (\pi_D + \pi_H)/2$ give equal weight to windows, not to pooled sample sizes.

For completeness, let $v(A)$ be F with factors in $A \subseteq \{\pi, p, m\}$ taken from H and the others from D. For each factor j, the exact six-permutation Shapley allocation is

$$\Delta_j = \frac{1}{6} \sum_{\sigma \in \mathcal{S}_3} \left[v\big(A_\sigma(j) \cup \{j\}\big) - v\big(A_\sigma(j)\big) \right],$$

where $A_\sigma(j)$ contains the factors preceding j in order σ. The components sum to $\bar{C}_H - \bar{C}_D$ before rounding; they do not allocate the hinge slope change β_2.

The reported 95% intervals are percentile intervals from 2,000 bootstrap replicates (seed 20260907), resampling UTC calendar days with replacement separately within each window and retaining all observations in each sampled day. Bin frequencies and both maps are recomputed jointly in every replicate. A replicate with an undefined required bin mean is invalid for the corresponding statistic; there is no zero-filling, borrowing or merging. All reported headline statistics have 2,000 valid replicates in both markets. These intervals describe the exploratory analysis and are not adjusted for its post-confirmatory selection.

Secondary omnibus diagnostics compare the full incidence or magnitude vectors across bins using a Moore–Penrose Wald statistic and centered bootstrap draws. They do not drive the transport conclusion. The intensive-margin covariance matrices have full rank (10) but condition numbers of approximately 2.0×10^4 for ETH5 and 2.3×10^6 for BASE5. This numerical conditioning limits the interpretation of the omnibus tests; the main evidence is the transport, standardization and component intervals in Table 3.

H Support-Cap Robustness

This post-confirmatory exploratory robustness uses the existing cap ladder $x \le 1.5, 2, 3, 5$ and Full, retaining respectively the first 6, 7, 8, 9 and all 10 fixed bins in Appendix G. Complete half-open bins are used: the finite retained sets are $[0, c)$. No retained observation is exactly at a finite cap, so these sets coincide with the stated $x \le c$ labels for this sample and its resamples.

For each window and cap, set

$$\pi_{w,c}(b) = \frac{n_w(b)}{\sum_{j \text{ retained}} n_w(j)}.$$

The within-bin incidence and magnitude maps are unchanged by restriction. Equations (S1)–(S3) then apply over the retained bins with these normalized frequencies and common weights $(\pi_{D,c} + \pi_{H,c})/2$. Thus the capped mean is $\bar{C}_{w,c} = \sum_b \pi_{w,c}(b) p_w(b) m_w(b)$, and its exact Shapley allocation remains a

decomposition of a mean change, not the hinge result. Different caps condition on different samples; Full-minus-cap differences are not an additive tail decomposition.

The bootstrap retains the original window-specific UTC-day universes and uses the same 2,000 resamples (seed 20260907) across caps. All maps and cap frequencies are recomputed within each replicate. A retained bin with undefined incidence or magnitude invalidates that cap's required statistics, without filling or merging; excluded bins do not invalidate a cap. All headline market/cap/statistic combinations have 2,000 valid replicates and zero invalid share. Table S8 reports the full ladder; the corresponding figure appears in the main text as Figure 3 (Section 4.1). The intervals are pointwise and do not adjust for selecting the analysis after the holdout evaluation showed instability. They do not provide simultaneous coverage across caps, and crossing zero does not establish equivalence.

Table S8: Post-confirmatory exploratory support-cap robustness

Panel A: retained counts and capped means (bp)

Market	Cap	N_D	N_H	$\bar{C}_{D,c}$	$\bar{C}_{H,c}$
ETH5	$x \leq 1.5$	173,352	142,835	0.20003	0.21285
ETH5	$x \leq 2$	180,271	150,904	0.32493	0.36354
ETH5	$x \leq 3$	183,287	155,525	0.42759	0.50917
ETH5	$x \leq 5$	183,879	157,628	0.46641	0.58902
ETH5	Full	183,993	170,513	0.48351	0.57982
BASE5	$x \leq 1.5$	1,103,268	949,306	0.07603	0.09926
BASE5	$x \leq 2$	1,110,308	966,196	0.09453	0.12899
BASE5	$x \leq 3$	1,111,921	980,005	0.10257	0.14463
BASE5	$x \leq 5$	1,112,270	996,890	0.10528	0.14900
BASE5	Full	1,112,362	1,052,212	0.10661	0.14297

Panel B: transport and mean-change allocation (bp)

Cap	TE_c	STD_DIFF$_c$	$\Delta_{\pi,c}$	$\Delta_{p,c}$	$\Delta_{m,c}$
ETH5					
$x \leq 1.5$	+0.00327 [−0.01765, 0.02351]	+0.00370 [−0.01728, 0.02402]	+0.00911 [−0.01306, 0.03177]	−0.00834 [−0.01405, −0.00179]	+0.01205 [−0.00581, 0.02851]
$x \leq 2$	−0.01955 [−0.05060, 0.00803]	−0.01591 [−0.04552, 0.01015]	+0.05448 [0.00266, 0.10595]	−0.01115 [−0.01796, −0.00334]	−0.00473 [−0.03276, 0.01751]
$x \leq 3$	−0.05755 [−0.11321, −0.01671]	−0.04549 [−0.09156, −0.00997]	+0.12703 [0.03736, 0.21564]	−0.01109 [−0.01871, −0.00263]	−0.03437 [−0.08208, −0.00102]
$x \leq 5$	−0.13671 [−0.28377, −0.04822]	−0.09466 [−0.18699, −0.03042]	+0.21746 [0.09923, 0.34369]	−0.00939 [−0.01776, −0.00023]	−0.08547 [−0.18161, −0.02093]
Full	−2.21198 [−4.55869, −0.51885]	−1.14083 [−2.32839, −0.28676]	+1.25606 [0.45739, 2.40648]	+0.01113 [−0.00936, 0.04342]	−1.17089 [−2.39130, −0.28048]
BASE5					
$x \leq 1.5$	+0.01256 [−0.00390, 0.02622]	+0.01407 [−0.00155, 0.02759]	+0.00918 [−0.00179, 0.02121]	+0.00614 [0.00020, 0.01166]	+0.00791 [−0.00320, 0.01751]
$x \leq 2$	−0.00854 [−0.04549, 0.01655]	−0.00031 [−0.02746, 0.02005]	+0.03451 [0.01349, 0.05810]	−0.00008 [−0.00931, 0.00812]	+0.00003 [−0.01827, 0.01305]
$x \leq 3$	−0.07041 [−0.17318, −0.00311]	−0.03445 [−0.09355, 0.00660]	+0.07327 [0.03574, 0.12100]	−0.01032 [−0.02545, 0.00260]	−0.02090 [−0.05960, 0.00494]
$x \leq 5$	−0.21035 [−0.48694, −0.04679]	−0.10572 [−0.24976, −0.01690]	+0.13637 [0.06257, 0.24964]	−0.02342 [−0.05215, −0.00306]	−0.06923 [−0.16336, −0.01065]
Full	−1.04663 [−2.59992, −0.26118]	−0.52453 [−1.30778, −0.12596]	+0.48914 [0.15795, 1.11506]	−0.08253 [−0.22426, −0.01874]	−0.37025 [−0.87652, −0.08841]

Notes: Brackets are pointwise 95% percentile intervals from 2,000 within-window calendar-day bootstrap replicates, seed 20260907. Every market/cap/statistic has 2,000 valid replicates and 0% invalid share. Complete fixed bins are retained; no observations equal the finite cap boundaries. State frequencies are renormalized within each window and cap. Full reproduces Table 3 before rounding. Shapley components allocate capped aggregate mean changes, not β_2. Neither simultaneous nor post-selection coverage is asserted.