The Queue Was the Arbitrage: What 147 Days of Tick-Size Asymmetry Were Worth

Wei Zhang

MOA International Pty Ltd, Sydney, Australia · accounts@moyaes.com

Working paper · 22 September 2026

Abstract

On 14 April 2026 Venue A cut the tick size of USDC/USDT from 0.0001 to 0.00001. Venue B did not follow until 8 September. For the 147 days in between, two of the largest spot venues for the same pegged asset quoted it on grids differing by a factor of ten, with zero trading fees on both. This paper measures what that asymmetry was worth. It was worth almost nothing, and the reason is instructive. Measured on both venues' published trade tapes, the cross-venue price difference reached half a basis point 41.5% of the time during the window and 0.10% of the time after it closed, a four-hundredfold change with a single overnight break and no trend on either side of it. The gap was therefore real, large and persistent. It was also, for practical purposes, untradeable, by two independent measurements. First, a taker-to-taker bound built only from prices that actually printed yields an upper bound of 4.3% a year when the two legs are allowed to be a second apart and 0.1% a year when they must be within ten milliseconds; the apparent opportunity is mostly price drift inside the matching window, and it is larger after the window closed than during it. Second, a queue-realistic replay of the obvious strategy over 91 coarse-grid days, resting a passive quote on the coarse venue and hedging each fill on the fine one, earns \$3.14 a day on a \$10,000 quote for a participant who joins the back of the queue, about 5.7% a year on the inventory the two legs require, and 130 times that for one who is first in line at every re-price and thereby absorbs a fifth of the venue's entire volume. The mechanism is that a coarse tick does not create an arbitrage; it creates a rent, and the rent belongs to whoever is at the front of the queue. Across the whole window the coarse grid's one-basis-point effective spread splits into about 0.62 to 0.66 bp realised by makers and 0.34 to 0.39 bp of price impact, against 0.03 bp and 0.07 bp on the fine-grid venue over the same days: an adverse-selection share of a third to two fifths on the coarse book against two thirds on the fine one. The split is computed by reconstructing each venue's mid from its own trade tape, an identity that holds because both venues quote a one-tick spread almost always, and it is verified against the coarse venue's own published order book on 86 days, where the two methods differ by half a hundredth of a basis point. The result replicates. The larger venue cut its own tick on 14 April and the smaller on 8 September, five months apart on the same asset, and each venue's realised spread falls by a factor of about twenty at its own change. The transfer is large in dollars and small per trade: venue A's takers paid \$1.0 million of queue rent in thirteen coarse-grid days and venue B's paid \$0.64 million over 147. Narrowing the tick removed the rent, the queue and the gap together, multiplied the venue's volume fourfold, and, for the first time, made its mid price move often enough for price discovery to be measurable at all.

Keywords: stablecoins, market microstructure, tick size, cross-venue arbitrage, queue rent, adverse selection, latency, order book.

JEL classification: G12, G14, G15, E42.

1. Introduction

A tick size is the smallest price increment a venue permits. When two venues list the same asset on different grids, the coarser venue's price is a rounding of the finer venue's price, and the rounding error looks like an arbitrage. This paper is about a case where the rounding error was large, lasted 147 days, and was worth almost nothing.

The asset is USDC/USDT, a pair of dollar stablecoins whose fundamental value is pinned at parity, so that almost every price difference between venues is a microstructure artefact rather than an information difference. The venues are two of the largest spot venues for the pair, both quoting it with zero trading fees. On 14 April 2026 the larger venue, A, cut its tick from 0.0001 to 0.00001. The other, B, kept the coarse grid until 07:00 UTC on 8 September. Both dates are established below from the venues' own trade tapes, not from announcements.

The window matters because it is the cleanest form of a question that is usually impossible to ask. In equities, tick-size differences between venues for the same security are rare, short and confounded by fee schedules, rebates and order protection rules; the Tick Size Pilot, the one deliberate experiment, widened the tick for a matched set of stocks across all venues at once (U.S. Securities and Exchange Commission, 2018). Here there are no fees on either side, the asset's value is pinned, the two venues are recorded on one clock, and the asymmetry has a dated beginning and a dated end. Whatever the coarse grid was worth to a cross-venue trader, it should be visible.

The answer has three parts. The gap was real: half a basis point or more, 41.5% of the time, every month, for 147 days, collapsing to 0.10% overnight when the window closed. The gap was not capturable in any amount that matters: crossing it with market orders yields a measured upper bound of a tenth of a percent a year once the two legs are required to be simultaneous, and quoting passively on the coarse venue, behind a queue \$12 to \$14 million deep, earns a newcomer a median of one fine-grid tick per fill and loses on one fill in five, about \$3 a day on a \$10,000 quote. And the reason is that a coarse tick is not a mispricing to be arbitraged but a rent to be queued for: the extra 0.9 basis point of spread on the coarse venue was earned by the makers at the front of the queue and paid by that venue's own takers, and it carried no more adverse-selection compensation than the one-tick spread on the fine-grid venue.

Section 2 describes the data. Section 3 dates the window. Section 4 measures the gap. Section 5 measures what could be taken from it, in two independent ways. Section 6 explains the mechanism using the tick change as a natural experiment. Section 7 reports the one kind of cross-venue window that does carry money, which arises from lag rather than from grids, and what happened to all three venues under a macroeconomic shock. Sections 8 and 9 give limits and conclusions.

2. Data

The venues are anonymised as A, B and C in descending order of traded volume. All three are large centralised exchanges that quote USDC/USDT with zero trading fees under their current schedules. Two kinds of data are used and it matters which question each can answer.

Published trade tapes. Both A and B publish every executed trade with a millisecond timestamp, price, size and the aggressor's side. These cover the entire window and beyond: both venues' complete tapes for 14 April to 17 September are used throughout, and the full monthly archives, A's from December 2018 and B's from September 2021, are used once, in Section 3.1. Tapes show what traded, not what was resting, so they can measure price differences and bound what a market order could have captured, but they cannot model a queue.

Order-book recordings. Three recorders on a single virtual machine in Tokyo subscribe to each venue's public book, best-quote and trade feeds and stamp every message on receipt, so cross-venue timing is exact to network jitter, measured at 4 to 7 milliseconds round trip. These show resting size and queue position, which is what a passive strategy needs. The recorder for venue A started on 11 August and for venue B on 5 September, three days before the window closed. Venue B also publishes its own daily 400-level book as a downloadable file, retained for about three and a half months, and 91 coarse-grid days were retrieved from it; that is what Section 5.2 uses. Venue A publishes no equivalent, so its touch is reconstructed from its tape where the recordings do not reach.

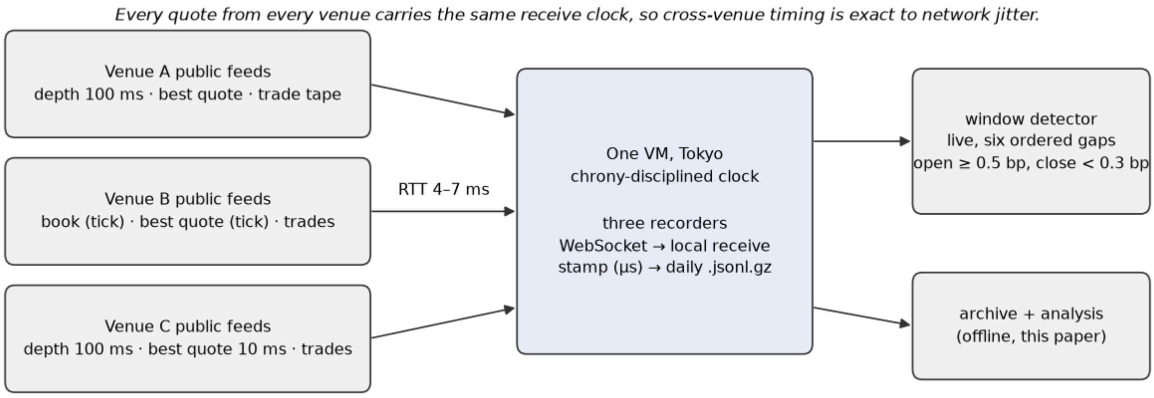

Figure 1. The recording set-up. Three public feeds per venue arrive over WebSocket at one machine whose clock is disciplined by chrony, and every message is stamped on receipt. The same machine runs the live window detector of Section 7. The archive is analysed offline.

Venue B's recorded trade feed was checked against B's own published trade files and agrees on daily volume to within 1% on eleven of the twelve full overlapping days, which is the paper's evidence that the recordings are faithful.

Two definitions are used throughout. The *gap* at an instant is the difference between the two venues' prices in basis points; how it is measured from tapes is stated in Section 4. A *window*, used only in Section 7, is a period during which one venue's resting best quote can be traded against the other's at a profit, detected live from the book feeds.

3. Dating the window

Neither venue published a notice that this paper could locate. Both changes are dated from the tapes themselves, by counting the distinct prices at which the asset traded each day.

Venue A, UTC day	Distinct traded prices	Example prices	Implied grid
11 April	2	0.9995, 0.9996	0.0001
12 April	4	0.9995 … 0.9998	0.0001
13 April	8	0.9992 … 0.9999	0.0001
14 April	56	0.99911, 0.99912, 0.99913 …	0.00001
15 April	48	0.99930, 0.99931, 0.99932 …	0.00001

Table 1. Venue A's tick change, 14 April 2026. Before the change the asset trades at a handful of prices a day, all on the 0.0001 grid; from 14 April it trades at dozens of five-decimal prices. Venue B's tape shows the mirror image on 8 September: five-decimal prices first appear that day and account for 89% of its trades from 9 September onwards. The change is timed at 07:00 UTC from the order-book recording, in which five-decimal quotes first appear at 07:00:42.

The sample therefore divides into four regimes: both venues coarse until 13 April; **asymmetric from 14 April to 7 September, 147 days**, with A on the fine grid and B on the coarse one; the change day, 8 September; and both fine from 9 September. The first and last regimes serve as controls for the middle one.

3.1 Are there earlier changes?

The same counting method can be run backwards over each venue's entire published history, and the answer it returns is worth stating because it is negative. Venue A publishes monthly trade files from December 2018, venue B from September 2021. Testing every print in every month against the 0.0001 grid gives the same number throughout: 100.0% of trades on the coarse grid in each of venue A's 83 published months from December 2018 to March 2026 and in the thirteen days of April before its change, and in each of venue B's 60 months from September 2021 to August 2026 and in the seven days of September before its own. Month-ends were sampled separately from month-starts, so a change made and held mid-month would have shown; 121 of venue A's 121 coarse-grid samples and 60 of venue B's 60 return the share as 1.0000. Venue A's files are absent from October 2022 to February 2023, when it suspended the pair during a stablecoin conversion programme; the grid is identical on both sides of that gap.

Each venue therefore changed this tick exactly once in its published life, and the two changes are the two endpoints of this paper's window. That has one cost and one benefit. The cost is that no second natural experiment exists to replicate the event study: the window is a single realisation, and Section 8 treats it as such. The benefit is that the coarse tick was evidently not a parameter these venues revisited. It survived every stress episode in the sample, including March 2023, when the asset traded as low as 0.8698 on venue B and 0.8820 on venue A, and October 2025, when it reached 0.9852 — a thirteen per cent dislocation quoted, from beginning to end, in whole hundredths of a basis point. The rent documented in Section 6 was not a transient artefact of a setting nobody had got round to; it was the standing configuration of this market for five to seven years, and it ended on both venues within a hundred and forty-seven days of each other.

4. The gap

Trade tapes give executed prices, which alternate between the bid and the ask, so a raw print-to-print comparison would be dominated by bid-ask bounce. The gap is therefore measured on one-second volume-weighted average prices: for each UTC second in which both venues traded, each venue's VWAP is computed and the gap is their difference in basis points. Averaging within the second removes most of the bounce while keeping the comparison free of any book data, so it can run across the whole window.

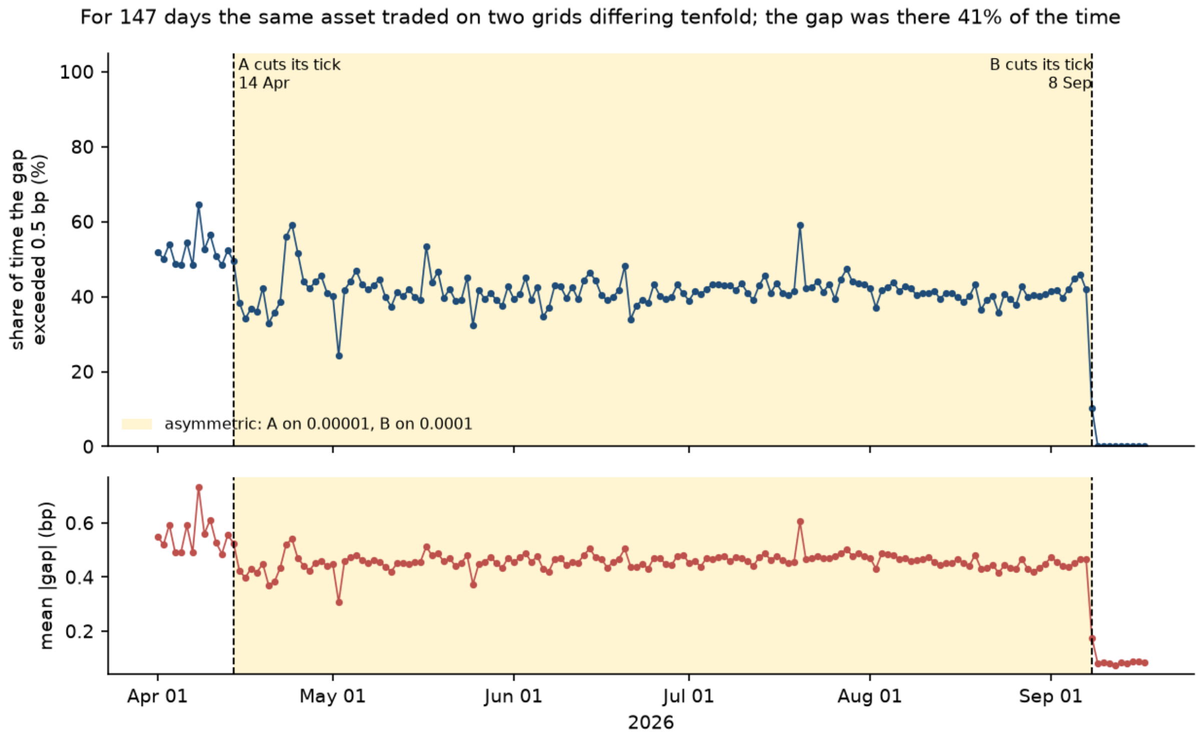

Figure 2. The window, day by day. Top: the share of each day on which the cross-venue gap exceeded half a basis point. Bottom: the mean absolute gap. The shaded band is the asymmetric regime. The level is flat for 147 days and falls off a cliff on 8 September; there is no trend approaching either date.

Regime	Days	Seconds per day	Mean |gap|	1st pct	Median	99th pct	≥ 0.3 bp	≥ 0.5 bp	≥ 1 bp	≥ 2 bp
Both coarse (to 13 Apr)	13	13,226	0.554 bp	−1.07	+0.11	+1.67	59.3%	52.4%	14.9%	0.93%
Asymmetric (14 Apr to 7 Sep)	**147**	18,879	**0.458 bp**	−1.03	+0.05	+1.07	61.3%	**41.5%**	3.5%	0.01%
Change day (8 Sep)	1	34,037	0.177 bp	−0.90	0.00	+0.90	16.4%	10.4%	0.47%	0
Both fine (from 9 Sep)	9	35,617	**0.083 bp**	−0.29	0.00	+0.25	1.1%	**0.10%**	0.00%	0

Table 2. The cross-venue gap by regime. One-second VWAP on each venue, inner-joined on the second. Percentages are shares of the matched seconds. The asymmetric regime spends 41.5% of its time at or beyond half a basis point; the both-fine regime spends 0.10%, a factor of four hundred. Note also that the coarse-coarse regime is worse still: when both venues round to the same coarse grid, the gap exceeds one basis point 14.9% of the time.

The effect is stable month by month, which rules out its being driven by any single episode.

Month	Days	Mean |gap|	Share ≥ 0.5 bp	Share ≥ 1 bp
April	30	0.493 bp	47.0%	9.1%
May	31	0.452 bp	41.0%	3.2%
June	30	0.461 bp	40.8%	3.5%
July	31	0.474 bp	43.0%	4.3%
August	31	0.453 bp	40.4%	2.6%
September (to 17th)	17	0.242 bp	18.2%	1.2%

Table 3. The gap month by month. April includes the thirteen coarse-coarse days, which is why it is the highest. September mixes seven asymmetric days with nine post-change days. The five full months in between are indistinguishable from one another.

This is the fact the rest of the paper has to explain. A trader watching these two venues during the window would have seen, on four days out of ten by time, one venue's price at least half a basis point away from the other's, on a pair whose fundamental value cannot move. The next section asks what could have been done about it.

5. What the gap was worth

5.1 Crossing it: a taker-to-taker bound over all 147 days

The direct way to take a cross-venue price difference is to buy on the cheap venue and sell on the dear one, crossing the spread on both. Without book data this cannot be simulated exactly, but it can be bounded from above using only prices that actually traded. To buy, one must lift an ask, and the evidence that the ask was P is an aggressor-buy print at P; to sell, one must hit a bid, and the evidence that the bid was Q is an aggressor-sell print at Q. Both tapes carry the aggressor side. For each matching interval and each venue, the best ask is the lowest aggressor-buy print and the best bid is the highest aggressor-sell print, with the size that actually printed at that price. A round trip is possible whenever one venue's bid exceeds the other's ask, and its size is capped by the printed volume on each leg and by a $10,000 clip. Capital is taken to be twice the clip, since inventory must sit on both venues.

The result depends entirely on how simultaneous the two legs are required to be, and that dependence is the finding.

The two legs must be within	Asymmetric, 147 days		Both fine, 9 days	
	Crossed intervals per day	Annualised	Crossed intervals per day	Annualised
1 second	82	+4.3%	826	+31.2%
200 milliseconds	23	+1.0%	298	+8.2%
50 milliseconds	8	+0.3%	99	+2.5%
10 milliseconds	**1**	**+0.1%**	19	+0.5%

Table 4. The taker-to-taker upper bound, by simultaneity tolerance. Annualised on $20,000 of capital, before any competition for the prints, before latency, before inventory rebalancing. The mean edge when a crossing exists is 0.09 to 0.23 bp throughout.

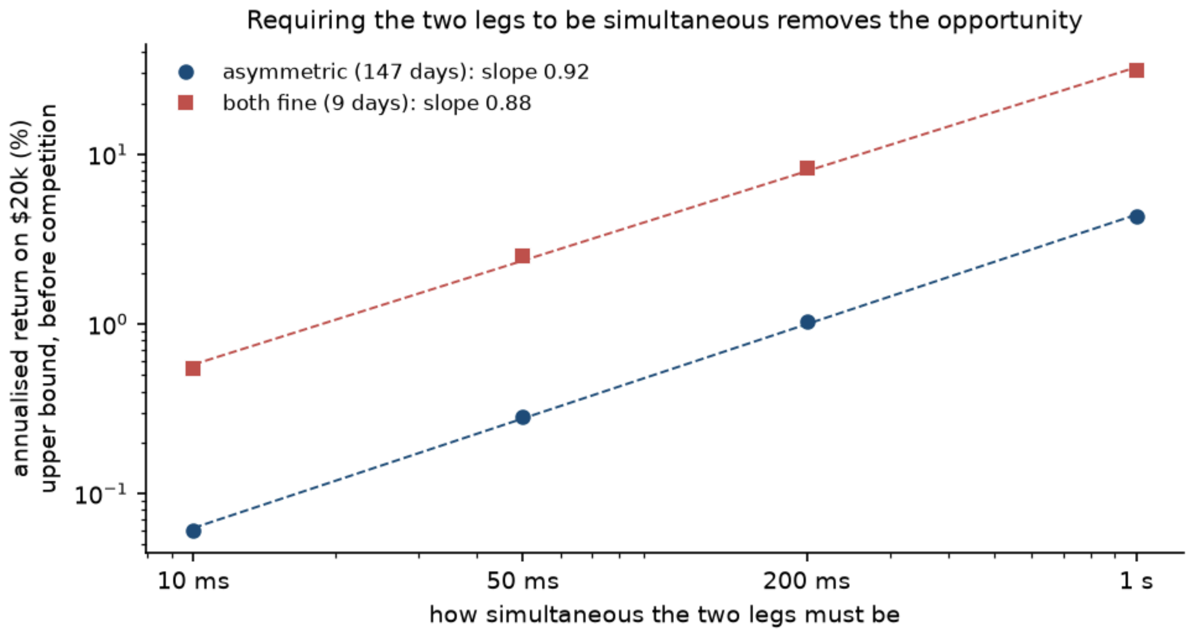

Figure 3. Requiring simultaneity removes the opportunity. Both axes are logarithmic; dashed lines are least-squares fits. The bound is a power law in the tolerance with an exponent of 0.92 during the window and 0.88 after it (R^2 of 1.000 and 0.999), so each tenfold tightening removes about seven eighths of it; from one second to ten milliseconds the bound falls from 4.28% to 0.061% a year during the window and from 31.2% to 0.54% after it.

Two things follow, and the second is the more important. First, the opportunity is an artefact of the matching window: inside one second the price itself moves, so a bid on one venue above an ask on the other is usually the price having drifted between the two prints rather than two prices existing at once. Tightening the window to ten milliseconds, which is still generous against a 4 to 7 millisecond round trip, leaves a bound of one crossed interval a day and a tenth of a percent a year. The exponent in Figure 3 says the same thing more precisely. If apparent crossings are the price moving between two prints, their number grows in proportion to the time allowed between the prints, which is an exponent of one; a genuine simultaneous mispricing would not depend on the tolerance at all, which is an exponent of zero.

The measured exponent is 0.92, and it is 0.88 after the window closed, so what survives in either regime is the same drift and not a residue of the grid.

Second, and decisively for this paper's question, the bound is *higher after the window closed than during it*, at every tolerance. The asymmetric regime yields 0.1% a year at ten milliseconds; the both-fine regime yields 0.5%. Whatever residue survives strict simultaneity is not created by the tick asymmetry. It is created by trade frequency, which the finer grid increased. The coarse grid contributed nothing that a market order could take.

5.2 Quoting into it: a queue-realistic replay

The other way to take the gap is to be passive on the coarse venue: rest a quote at its touch, and when filled, hedge immediately on the fine-grid venue. This is the strategy the gap appears to invite, because on the coarse grid the spread is pinned at a full basis point and the fine-grid venue's price sits somewhere inside it. It requires resting depth, which the author's own recordings cover for only three days of the window. Venue B, however, publishes its own daily 400-level order book and its content delivery network retains about three and a half months of it; 91 coarse-grid days, 9 June to 7 September, were retrieved from that source and are used here.

The replay rests a $10,000 quote at venue B's touch, re-prices it whenever the touch moves, and hedges each partial fill at venue A's touch with a 30 millisecond action delay. Venue A's touch is reconstructed from its public trade tape by the same one-tick identity used in Section 6, which on the overlapping days gives results indistinguishable from the author's recorded quote stream. The outcome depends almost entirely on one modelling choice: where in the queue the participant stands each time the order is re-priced.

Queue assumption	Partial fills	Volume	Mean edge	Negative-edge fills	Net, 91 days	Per day	Profitable days
Newcomer: join behind the resting depth	25,857	$22 M	+0.179 bp	21%	**+$286**	+$3.14	85 of 91
Incumbent: first in line at every re-price	4,702,276	$1,330 M	+0.404 bp	8%	**+$37,056**	+$407	91 of 91

Table 5. The same strategy from the back and from the front of the queue. 91 coarse-grid days, 9 June to 7 September, $10,000 clip, hedged 30 ms later. Moving from the back of the queue to the front multiplies the fill count by 182, the volume by 60, the edge per fill by 2.3 and the profit by 130. Both rows are profitable and both are stable: monthly averages run $2.49, $2.95, $3.55 and $4.15 a day at the back, and $304, $473, $396 and $494 at the front, and the front has no losing day in 91. Mean edge is the average over fills; for the newcomer the volume-weighted mean is +0.129 bp and the median fill earns +0.100 bp, exactly one tick on venue A.

Queue position, not the cross-venue gap, is what the outcome turns on, and the two rows are different businesses rather than two sizes of the same one. The newcomer's row is a strategy: $3.14 a day on a $10,000 quote, about 5.7% a year on the $20,000 of inventory the two legs require, before competition and before any cost of moving coins between venues. The incumbent's row is a bound. Filling $14.6 million a day absorbs 22% of everything venue B traded, which would require winning the race to the front of every price level against $13 million of standing orders; and the capital it ties up is not the clip but whatever inventory supports that turnover, so it cannot honestly be annualised at all. What the row establishes is where the money was, not that it was available.

Is the newcomer's hedge too optimistic? A natural objection is that a passive fill on the coarse venue is often the footprint of an informed sweep that reaches venue A in the same milliseconds, so that the touch the replay hedges against is already gone. The replay does not hold venue A's quote fixed: it takes whichever touch prevails 30 ms after the fill, so a sweep that trades through venue A's touch in that interval is already priced in. What it does not model by default is the quantity displayed at that touch, or quote changes that leave no print. Three checks address the objection directly. First, delaying the hedge from 30 ms to 100 ms, 250 ms and a full second moves the 91-day result from \$285.56 to \$285.43, \$284.77 and \$283.70, and the mean edge is 0.179 bp in every case; if fills were systematically followed by venue A moving against the position, a hedge a second late would pay for it. Second, on the 28 days where the author's recording of venue A's quote stream overlaps the window, the hedge was rerun against the recorded quotes, which see cancellations as well as trades, with each hedge capped by the quantity actually displayed at venue A's touch, net of the replay's own earlier hedges at the same quote, and any excess filled one tick deeper. The displayed quantity binds on 1.3% of hedged volume, and the result is 95.7% of the tape-based figure over the same days, \$4.02 a day against \$4.20. Third, sorting fills by the size of the aggressive order on venue B that caused them, the fills triggered by orders of \$100,000 or more, more than a quarter of the newcomer's volume, earn the highest mean edge, 0.20 bp, while the worst bucket is orders of \$1,000 to \$10,000 at 0.05 bp. Large aggressive orders on the coarse venue are not, in this sample, the leading edge of a move that has already reached the fine one. The newcomer's \$3.14 a day remains an upper bound for the reason given in Section 8, but the cross-venue sweep channel is measured, and it is small.

A second feature of the table matters for the paper's question. Even from the front of the queue the mean edge is 0.404 bp, while Section 6 measures the realised spread earned by venue B's makers on the coarse grid at 0.62 to 0.66 bp. The difference is the hedge: crossing venue A's spread on every fill returns a large part of the rent to that venue. The coarse tick's rent was payable in full only to a single-venue market maker willing to carry the inventory, not to a cross-venue arbitrageur who locks the position immediately. That is the sharpest form of this paper's claim. There was money in the coarse grid, it was substantial, and it was not in the cross-venue gap that the coarse grid made visible.

Venue B's published order book reaches back about three and a half months, so this measurement starts on 9 June and the first eight weeks of the window, 14 April to 8 June, have no resting-depth data from any free source. The 91 days that are covered span three months and show no trend, and the mechanism they illustrate is measured over the full 147 days in Section 6 from tapes alone.

5.3 What the live detector found

From 12 September a detector has run live on all three venues' book feeds, logging every interval in which one venue's resting best quote could be traded against another's at half a basis point or more. Over six days it recorded 57 such windows. Their median duration is 43 seconds, their median peak gap 0.9 bp, and the median resting size behind the lagging quote is **\$108**. The median one-shot capture is five cents, and the sum over all 57 windows is **\$83**. This is measured after the tick asymmetry ended, and Section 7 returns to it, but the order of magnitude is the point: even where a cross-venue gap is genuinely resting rather than drifting, the size behind it is trivial.

6. Why: a coarse tick is a rent, not a mispricing

Sections 4 and 5 leave a puzzle. The gap was there 41.5% of the time for five months and nobody could take it. The resolution is that the coarse grid's extra spread was never a cross-venue mispricing; it was a rent paid by that venue's own takers to the makers holding its queue. This section measures the rent twice: over the whole window from public tapes, and precisely over a short window from the recordings.

6.1 The rent over 147 days, from public tapes

Both venues quote this pair at a one-tick spread almost all the time, which makes the whole book recoverable from the trade tape. If the spread is exactly one tick, an aggressor-buy print at P means the ask was P and the bid P − tick, so the mid was P − tick/2; an aggressor-sell print pins it symmetrically. Validated against the recorded book on the overlapping days, the spread is exactly one tick on 99.8% of venue B's coarse-grid quotes and 98.9% of venue A's fine-grid quotes, every print sits on the quoted side, and the reconstructed mid matches the recorded mid with a root-mean-square error of 0.012 bp and 0.007 bp respectively, two orders of magnitude below the quantity being measured. The decomposition can therefore be computed for every trade on both venues across the entire window.

The proxy can be checked much harder than that. Venue B publishes its own 400-level book, and for 86 of the 91 coarse-grid days it covers, the decomposition can be computed twice on the same trades: once from the published quotes and once from the reconstruction. Table 6 reports both.

Venue B, coarse grid, 86 days	From the published book	From the tape reconstruction	Difference
Quoted spread, time-weighted	1.186 bp	—	—
Effective spread	1.011 bp	0.999 bp	−0.011
Realised spread, 5 s	**0.618 bp**	**0.623 bp**	**+0.005**
Price impact, 5 s	**0.392 bp**	**0.376 bp**	**−0.016**
Adverse-selection share	0.39	0.38	−0.01
Rent paid by takers		$391,804 over the 86 days, $4,556 a day	

Table 6. The proxy against the venue's own quotes, 9 June to 7 September. Both columns use the same trades and the same five-second horizon; the book column takes the prevailing quote 20 ms before each trade and 5 s after it. Day by day the two realised-spread series differ by a median of 0.000 bp and a 95th percentile of 0.025 bp. Five of the 91 days are missing because on those days, all weekends, venue B's best bid and offer did not move once in 24 hours, which leaves too few quotes to interpolate and is itself a fact about the coarse grid. Monthly realised spreads are 0.548, 0.641, 0.632 and 0.627 bp for June to September.

The reconstruction is therefore not an approximation that has to be argued for; on this pair it is the book. That matters because venue A publishes no comparable file and its own coarse-grid period, the thirteen days before 14 April, predates every recording. The reconstruction is the only way to measure that period at all, and the 86-day check is what licenses it.

Venue and grid	Days	Volume per day	Effective spread	Realised spread	Price impact	Adverse-selection share	Rent paid by takers
A, coarse (1 to 13 Apr)	13	$1,135 M	1.000 bp	**0.689 bp**	0.311 bp	0.31	**$1,016,875** ($78,221/day)
A, fine (14 Apr to 7 Sep)	147	$1,847 M	0.100 bp	**0.033 bp**	0.067 bp	0.67	$885,483 ($6,024/day)
B, coarse (14 Apr to 7 Sep)	147	$66 M	1.000 bp	**0.656 bp**	0.344 bp	0.34	**$636,527** ($4,330/day)
B, fine (9 to 17 Sep)	9	$293 M	0.100 bp	**0.039 bp**	0.061 bp	0.61	$10,331 ($1,148/day)

Table 7. The spread decomposition on both venues, before and after each one's own tick change. Volume-weighted, five-second horizon, computed from the venues' published trade tapes with the mid reconstructed as described above. The rent is volume times realised spread, that is, what the venue's takers paid to its makers and did not get back within five seconds. The effective spread equals the tick by construction, which is the point: on a one-tick book the tick *is* the spread (Harris, 1994). Each venue's pair of rows brackets its own tick change; the shaded pair is venue B, five months after venue A.

Table 7 is the paper's central table and it contains a replication. Venue A cut its tick on 14 April and venue B on 8 September, five months apart, on the same asset, and both produce the same result. On a coarse grid the maker keeps about two thirds of a one-basis-point spread and the adverse-selection share is about a third; on a fine grid the maker keeps a hundredth of a basis point and the adverse-selection share is about two thirds. Venue A's realised spread falls by a factor of 21 at its change and venue B's by a factor of 17 at its own. Whatever the coarse grid was doing, it was not compensating makers for informed flow: the fine-grid books, where the spread is as tight as the rules allow, are the ones where adverse selection accounts for most of it.

The dollar column gives the size of the transfer. Venue A's takers paid a little over a million dollars of queue rent in thirteen coarse-grid days; venue B's paid $637,000 over 147. None of it was available to a cross-venue participant, for the reason Section 5.2 measures: hedging the position on the other venue hands most of it straight back.

6.2 The same measurement from the recordings

Where the order book was recorded, the same decomposition can be computed from quotes rather than reconstructed from prints, which is more accurate and allows a formal difference-in-differences. Each venue-hour from 5 to 17 September is an observation; the change hours 06:00 to 08:00 UTC on 8 September are dropped; the regression is y = venue effect + hour-of-day effect + day-of-week effect + δ (Venue B × after), with standard errors clustered by day. Trade-based measures use each venue's own trade feed with the trade signed by its aggressor flag and the quote taken 20 ms before the trade; realised spread and price impact are measured five seconds after the trade and weighted by dollar volume.

Hourly measure	B, coarse	B, fine	A, before	A, after	δ (DiD)	s.e.	t
Quoted spread, bp	1.001	0.116	0.100	0.100	−0.886	0.001	−676
Effective spread, bp	1.000	0.127	0.105	0.107	−0.876	0.003	−263
Realised spread, 5 s, bp	**0.946**	**0.058**	0.020	0.017	−0.885	0.015	−57.5
Price impact, 5 s, bp	**0.054**	**0.068**	0.081	0.087	**+0.007**	0.016	**0.5**
Adverse-selection share	0.05	0.52	0.76	0.81	+0.428	0.059	7.2
Log depth at the touch	16.40	14.71	15.58	15.44	−1.545	0.077	−20.1
Log hourly volume	14.14	16.00	18.24	18.32	+1.784	0.155	11.5
Log mean trade size	6.16	7.51	8.53	8.56	+1.328	0.158	8.4

Table 8. Difference-in-differences, venue B treated, venue A control. 617 venue-hours, 13 day clusters. This table and Table 7 are computed from different sources over different spans and agree in every sign and in the ordering of every row; where they can be compared directly, against venue B's own published quotes over 86 days in Table 6, the two methods differ by 0.005 bp on the realised spread. The adverse-selection share is price impact divided by effective spread, the standard decomposition of the spread into compensation for informed flow and a residual that is order-processing and inventory rent (Glosten and Harris, 1988).

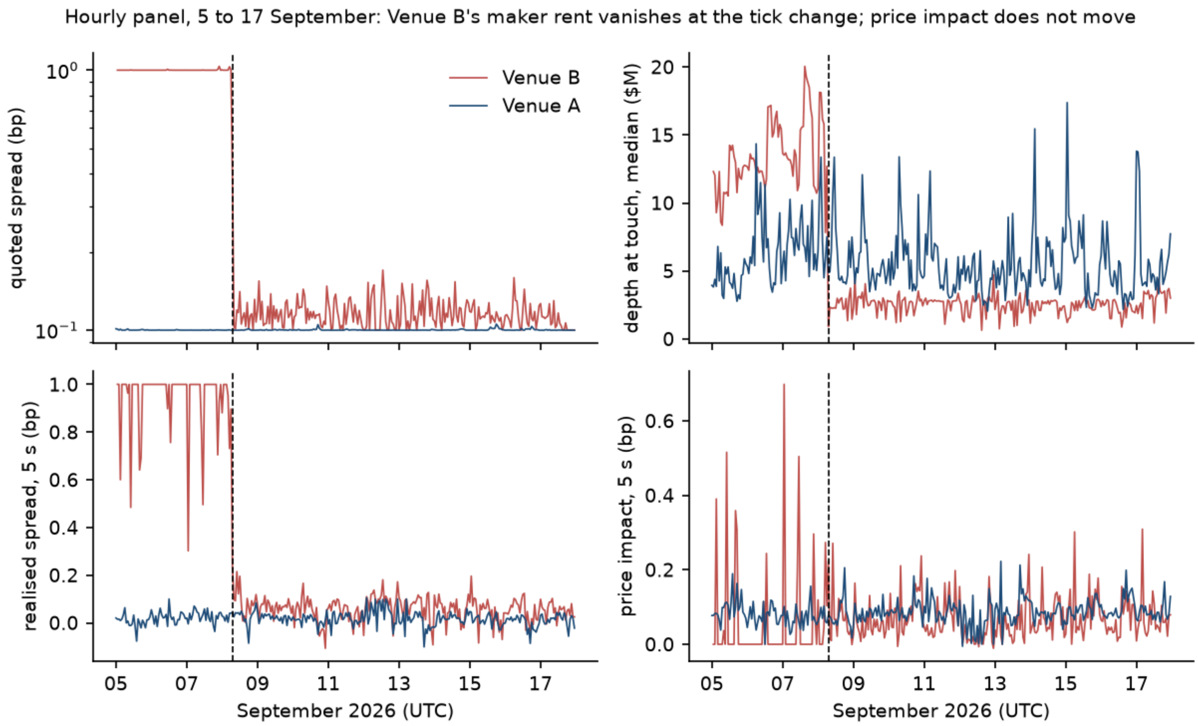

Figure 4. The hourly panel. Quoted spread on a log scale, depth at the touch, realised spread and price impact by venue-hour, 5 to 17 September; the dashed line is the change. The realised spread on venue B sits at a full basis point until the change and at a few hundredths after. Price impact on both venues is the same noisy 0.05 to 0.1 bp throughout.

The two bold rows are the explanation, and they agree with Table 6 on the short window where both can be computed. On the coarse grid venue B's effective spread was one basis point, of which 95% was still in the maker's hands five seconds later on these three days, and about two thirds across the whole 147 days.

Price impact, the part of the spread that compensates a maker for trading against someone better informed, was 0.054 bp, which is the same as the fine-grid control's 0.081 bp to within the noise, and the difference-in-differences coefficient on it is +0.007 with a standard error of 0.016. In other words, the extra nine tenths of a basis point on the coarse venue bought no extra protection against informed flow. It was a transfer from that venue's takers to the makers holding the front of its queue, imposed by the grid.

That reconciles Sections 4 and 5. A rent is not an arbitrage. To collect it one must hold queue priority in a $13 million line, which is a capital and latency position, not a trade; and a newcomer joining the back of the line is filled disproportionately on the fills nobody wants: in Section 5.2 a fifth of the newcomer's fills lose money on the hedge, against one in twelve for the participant at the front. Yao and Ye (2018) describe this rationing under a binding tick constraint; what is unusual here is being able to see the same rent from the other side of the boundary, as a cross-venue gap that four days out of ten looked like free money.

The tick change removed the rent, the queue and the gap in one step. Its other effects are the textbook ones for a narrowing tick, and they are reported in Table 9 against a three-month baseline from venue B's published tape rather than against the three recorded pre-change days, which were a weekend and a Monday with unusually small prints.

Venue B	**Coarse grid (Jun to Aug tape)**	**Fine grid (9 to 17 Sep)**	**Ratio**
Quoted spread, time-weighted mean	1.00 bp	0.116 bp	0.12
Depth at the touch, median	$13.5 M	$2.8 M	0.21
Traded volume per day	$70 M (weekdays $86 M)	$300 M (weekdays $338 M)	4.2
Trades per day	51,400	136,500	2.7
Mean trade size	$1,500	$2,200	1.5
Mid-price changes per day	122	12,400	102

Table 9. Venue B before and after. Spread and depth are from the order-book recordings and therefore compare the three recorded coarse days with the nine fine days; volume, trade count and trade size are from the published tape and use a 92-day baseline. Venue A's volume rose about a fifth over the same interval, so roughly a twentieth of B's fourfold rise is market-wide. The busiest coarse-grid day in three months, 21 August, traded $201 million, less than an ordinary fine-grid weekday.

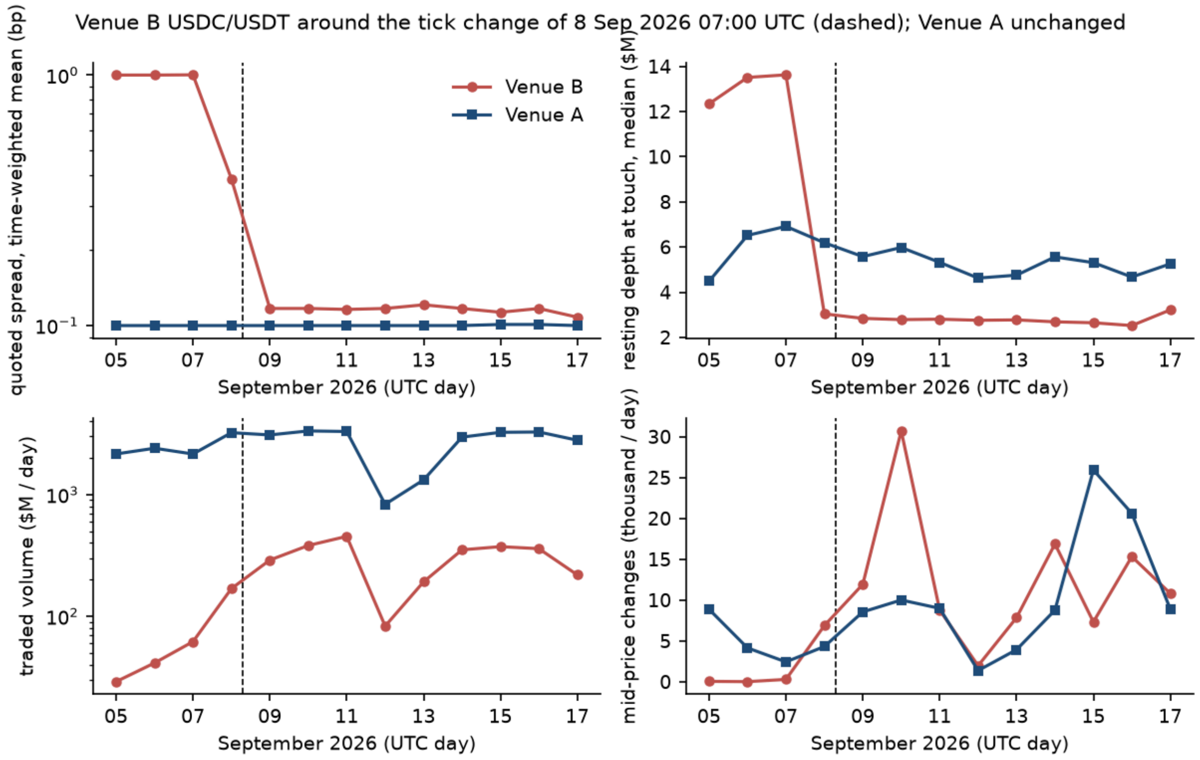

Figure 5. Venue B around the change, with venue A as control. Daily values, 5 to 17 September; the dashed line is 07:00 UTC on 8 September. Spread and volume on log axes; 12 and 13 September are a weekend.

The last row of Table 9 is the one that matters for price discovery. A book whose spread is one tick and whose tick is a full basis point has a mid that cannot move by less than half a basis point, and a stablecoin's fair value almost never moves that much. On the coarse grid venue B's mid changed about 122 times a day; on the fine grid it changes about 12,400 times. The coarse mid was not a price, it was a rounding of venue A's price, which is why the cross-venue gap of Section 4 existed at all and why it carried no information.

With both venues on the same grid, price discovery becomes measurable. Fitting a vector error-correction model to mid prices with the cointegrating vectors imposed at $(1, -1)$, which the peg justifies, gives venue A the large majority of the Hasbrouck (1995) information share on every one of the nine post-change days. Because the three venues update their quotes at very different rates, the sampling interval is the natural robustness check, and Table 10 runs it from one second down to a tenth.

Sampling interval	Venue A			Venue B			Venue C		
	Information share	GG weight	Stale	Information share	GG weight	Stale	Information share	GG weight	Stale
1 second	0.67 to 0.85	0.61	97%	0.15 to 0.31	0.33	97%	0.01 to 0.02	0.06	99%
500 ms	0.69 to 0.83	0.61	98%	0.17 to 0.30	0.33	98%	0.01 to 0.02	0.06	99%
200 ms	0.67 to 0.77	0.53	99%	0.21 to 0.30	0.33	99%	0.02 to 0.03	0.14	100%
100 ms	**0.66 to 0.73**	0.49	100%	**0.24 to 0.30**	0.35	99%	**0.03 to 0.04**	0.16	100%

Table 10. Price discovery by sampling interval, nine post-change days. Hasbrouck information shares are the daily means of the lower and upper bounds over Cholesky orderings; GG is the Gonzalo and Granger (1995) component weight; "stale" is the share of grid points at which that venue's mid had not changed since the previous point. Five lags on differences throughout. The two-venue system excluding venue C gives venue A 0.68 to 0.85 at one second and 0.65 to 0.72 at 100 ms. The contemporaneous correlation of the A and B innovations is 0.20, 0.16, 0.11 and 0.07 at the four intervals.

The ranking does not move: venue A leads at every interval, venue B follows, and venue C contributes essentially nothing to the common price. What does move is the width of the bounds and the split between the two large venues, which narrows from roughly three to one at one second to roughly five to two at a tenth of a second. That is the Epps (1979) effect: correlation measured over very short intervals is biased towards zero because venues do not update at the same instant, and here the contemporaneous correlation of venue A's and venue B's innovations falls from 0.20 at one second to 0.16, 0.11 and 0.07 at the finer intervals. At one second the two venues' innovations overlap and the Cholesky ordering matters, so the bounds are wide; at 100 ms they separate. It is not evidence that venue B gains ground at high frequency, because the stale column shows what the finer grid is actually made of. On a pegged pair the mid moves a few thousand times a day at most, so at 100 ms essentially every grid point repeats the last one and the estimate rests on the same few thousand moves however finely they are sliced.

Venue C is the instructive case, and it needs care. Its information share stays at three or four hundredths at every interval, while its Gonzalo-Granger weight rises from 0.06 to 0.16. Read naively, a rising component weight says venue C matters more to the common trend at high frequency. It does not. The weight is built from the error-correction coefficients, and a venue's weight rises when the others are measured to correct more, not when it corrects less. Expressed per second of clock time, venue C closes its gap to venue A at the same rate at every resolution, 0.78% to 0.80% of the gap a second, a half-life of about a minute and a half. What changes is venues A and B. Their measured correction speeds rise from 0.40% and 0.73% a second at one-second sampling to 0.85% and 1.18% at 100 ms, because the sub-second adjustment between the two large venues, which the one-second grid books as contemporaneous correlation, is booked as lagged error correction on the finer grid: the same Epps effect, seen from the other side. Venue C's rising weight is the arithmetic of a ratio whose other terms grew. The substantive fact is its flat and slow adjustment, and Section 7 shows the behavioural counterpart: the same venue withdrawing its book entirely under a macroeconomic release.

The same model applied to the coarse-grid days returns shares that swing between 3% and 47% from day to day, which is not a measurement of anything: a series that changes 122 times a day is, to an error-correction model, a series that never needs correcting.

7. The windows that do carry money, and a stress test

If grids do not create capturable cross-venue gaps, lag might. From 12 September the live detector described in Section 2 has logged every interval in which one venue's resting quote could be traded against another's at half a basis point or more. Table 11 summarises the 57 windows recorded in the first six days.

57 windows, 13 to 18 September	25th pct	Median	75th pct	Max
Duration	17 s	43 s	88 s	622 s
Peak gap	0.7 bp	0.9 bp	1.8 bp	6.2 bp
Move that opened it	0.3 bp	0.6 bp	1.6 bp	6.4 bp
Lag before the stale quote moved	8 s	30 s	51 s	345 s
Resting size behind the stale quote	$49	$108	$1,121	$794,000
One-shot capture	$0.00	$0.05	$0.19	$55.60

Table 11. Lag windows. Forty-nine of the 57 involve venue C. The sum of one-shot capture across all 57 is $83. The peak gap scales with the size of the originating move, with a correlation of 0.66, but the money does not: the large gaps are the ones with the emptiest quote behind them.

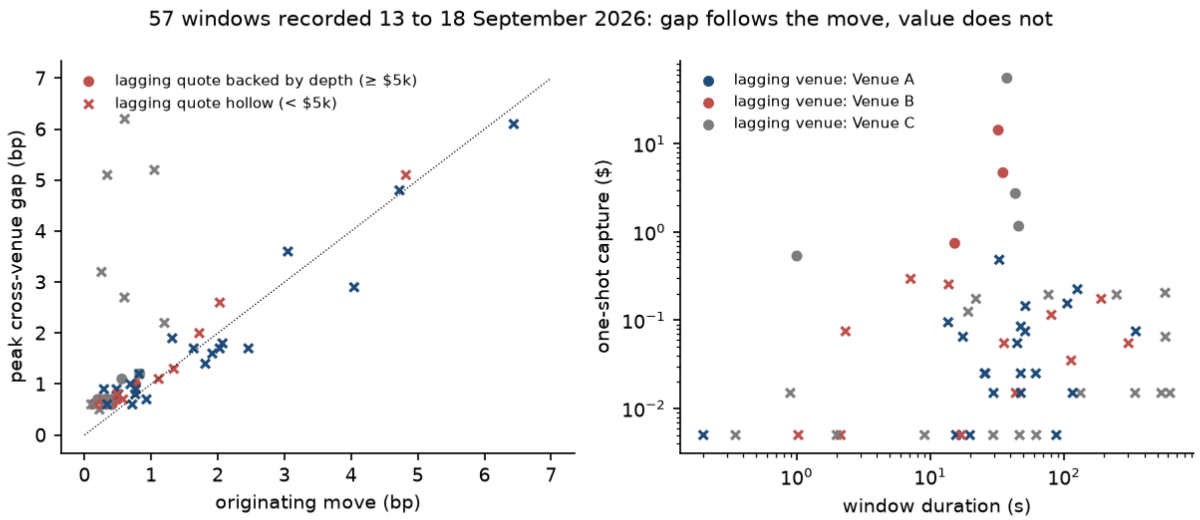

Figure 6. The 57 windows. Left: peak gap against the originating move; the dotted line is equality; crosses mark windows whose lagging quote held less than $5,000. Right: one-shot capture against duration, log axes, coloured by the lagging venue.

Two features of these windows are worth stating because they invert the usual intuition. They are *long*: the shortest lasted 0.2 seconds, the tenth percentile is 2.1 seconds, and all 57 were still open 200 milliseconds after they opened. Against a 15 millisecond order path the nominal window and the executable window are the same window, which is the opposite of the equities latency race described by Budish, Cramton and Shim (2015) and measured by Aquilina, Budish and O'Neill (2022). What limits capture here is not the speed of the second-fastest participant but the size of the order that failed to move, and that size is usually nothing. They are also *fee-fragile*: charging a half-basis-point taker fee on the crossing leg leaves 26 of the 57 windows nominally open but only three with as much as $5,000 behind them, and at one basis point nothing with money behind it survives.

The largest lag window in the sample predates the detector and is worth anatomising because it is the best case. On 9 September at 01:06 UTC venue A stepped down 0.6 bp and venue B's bid at the old price held for about 70 seconds while $4.9 million of aggressive selling hit it. The one-shot capture at the peak was about $58 and the whole-wall ceiling about $194.

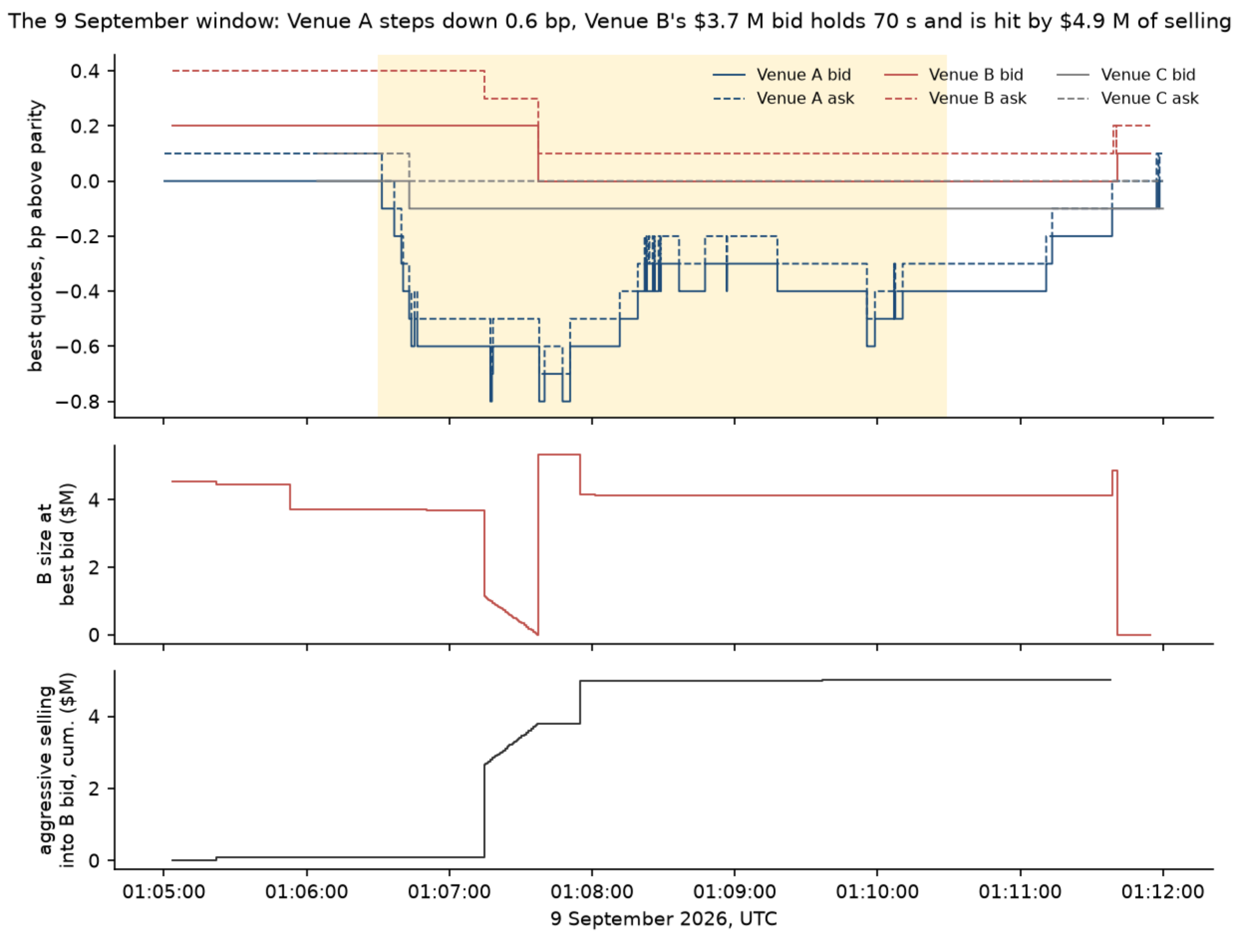

Figure 7. The largest lag window recorded, 9 September, 01:05 to 01:12 UTC. Top: best bid solid and ask dashed on the three venues in basis points above parity; the shaded band is the window. Middle: dollar size resting at venue B's best bid. Bottom: cumulative aggressive selling into that bid. The bid holds, is replenished, and the competing flow exceeds the wall itself.

Finally, the one macroeconomic shock in the sample separates the venues in a way that grids do not. The Federal Open Market Committee raised its target range on 16 September at 18:00 UTC (Board of Governors of the Federal Reserve System, 2026). Over the following 85 minutes venues A and B never differed by more than 0.65 bp and both held millions at the touch. Venue C's resting size collapsed from $3 to 4 million to a few thousand dollars, and for long stretches to under a hundred, while its mid wandered as far as 7.3 bp below the other two. A scan of every recorded venue-day finds that venue C does this only around United States macroeconomic releases: the consumer price index release on 11 September and the two 18:00 UTC afternoons of the Federal Open Market Committee meeting. Venue A has no hollow minute in 38 recorded days.

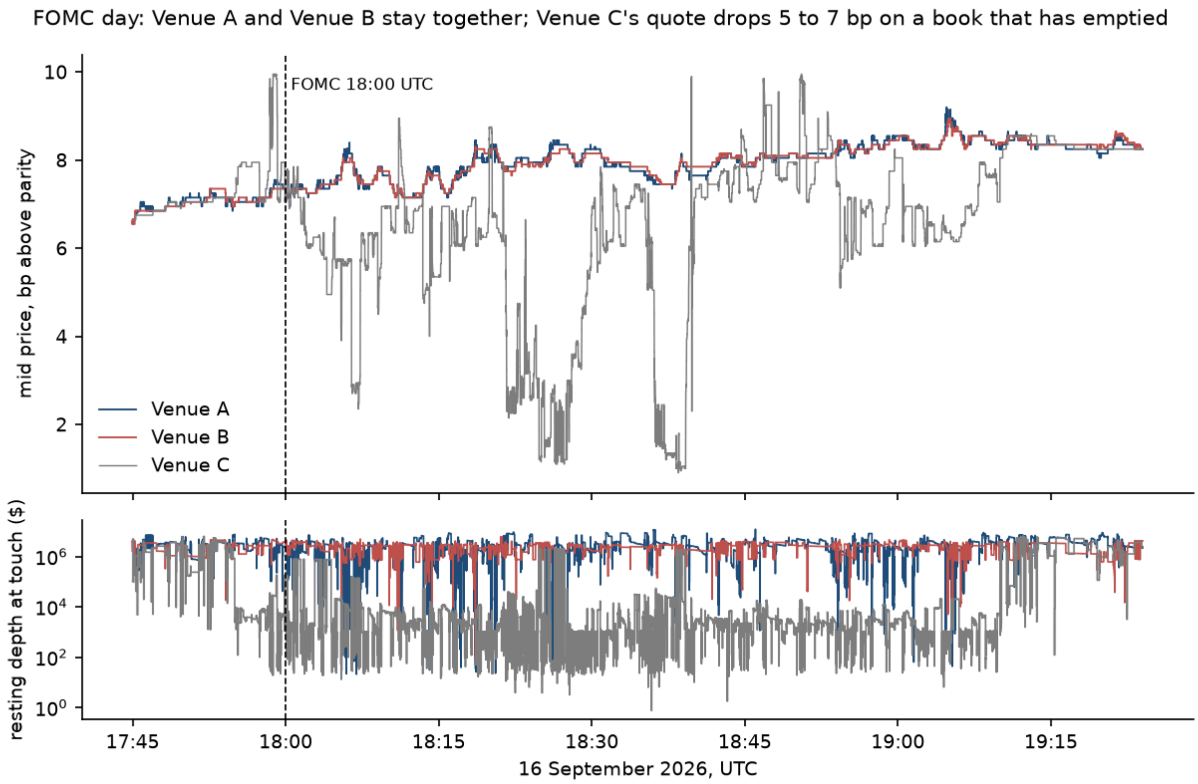

Figure 8. The FOMC afternoon, 16 September. Top: mid price on each venue in basis points above parity. Bottom: dollar depth at the touch, log scale. Two books stay one market; the third withdraws. A quote with $70 behind it is not a price, and the 29 windows the detector logged that afternoon are 29 readings of that absence.

8. Limits

One event per venue. Section 3.1 establishes that neither venue changed this tick at any other point in its published history, so the window cannot be replicated on another date for this pair. Every estimate here is conditional on a single realisation, and the difference-in-differences of Section 6.2 borrows its identification from the fact that the two venues changed on different dates rather than from repetition. What generalises is the mechanism, not the magnitudes.

Ninety-one days of book inside the window. The queue-realistic replay of Section 5.2 starts on 9 June because that is as far back as venue B's published order book is retained; the first eight weeks of the window have no resting-depth data, and venue A publishes no equivalent at all. Two modelling choices in that replay were found to dominate the result and are stated so that it can be reproduced. A quote must be re-priced whenever the touch moves: an order left behind the touch is not a strategy but an un-cancelled order that fills at a stale price, and leaving it there turns a profit of $407 a day into a loss of $1,070. And venue B's daily trade files are cut on a UTC+8 boundary, so each must be merged with its successor before being split into UTC days, or the last eight hours of every day are silently dropped.

The tape-reconstructed mid is a proxy, but a tested one. It is exact when the spread is one tick, which holds 99% of the time on the two regimes that matter. Checked against venue B's own published quotes on the same trades over 86 days, it differs by 0.005 bp on the realised spread and 0.016 bp on price impact, with a daily median difference of zero. It has not been checked on venue A's coarse-grid period, which is thirteen days in April that no recording and no published book reaches; those rows of Table 7 rest on the identity alone.

The hedge is modelled, not executed. The newcomer's result in Section 5.2 is insensitive to hedge delay up to a second and loses 4% when venue A's displayed depth is enforced, but no replay can model other participants hedging the same flow at the same instant, which would thin venue A's touch exactly when it is needed. That figure should be read as an upper bound for this reason, not for the sweep channel, which is measured.

The taker bound is a bound. Table 4 assumes the participant could have been the counterparty to prints that other people actually got, that the two legs execute at once, and that no one else competes. Every one of those assumptions pushes the estimate up, and it is still a tenth of a percent a year.

Inventory, not execution, is the real constraint. Every capture figure assumes inventory already sits on both venues. Closing a window leaves the participant longer on one and shorter on the other, and rebalancing means an on-chain transfer that takes minutes at best and, in exactly the dislocations that would make a window valuable, can be suspended (Kozhan and Viswanath-Natraj, 2021; Ma, Zeng and Zhang, 2025). The riskless arbitrage is riskless only for the first window.

Fee schedules move. Zero fees on all three venues hold under promotional or conditional schedules, and venue C's terms exclude API users. A gap of half a basis point is one taker fee away from zero, as Section 7 shows.

No dislocation in the sample, and none is coming. The window table contains no event of the 3 bp-or-larger kind counted in Zhang (2026), so the conjecture that lag and wall size scale with the size of the move is tested here only at drift scale. This gap cannot be closed later. The asymmetry ended on 8 September and Section 3.1 shows that neither venue has ever reversed a tick change, so any future dislocation will be quoted by both venues on the fine grid: an observation with no treatment period to compare it against. What a large dislocation does to cross-venue lag is a question about dislocations, not about grid asymmetry, and it belongs to a different paper.

9. Conclusion

For 147 days the same pegged asset traded on two zero-fee venues whose price grids differed by a factor of ten, and for four days out of ten by time the two prices differed by half a basis point or more. The asymmetry ended overnight on 8 September and the gap fell by a factor of four hundred. Nothing about the gap was hidden, and almost nothing about it was capturable. Crossing it with market orders yields a tenth of a percent a year once the two legs must be simultaneous, and less than the same measurement yields after the asymmetry ended. Quoting into it from the back of the queue earns a median of a tenth of a basis point per fill, loses on one fill in five, and comes to about $3 a day on a $10,000 quote; the same quote held at the front of the queue would have earned 130 times as much. The reason is that a coarse tick does not misprice an asset relative to another venue; it levies a rent on its own takers and pays it to whoever holds the front of a thirteen-million-dollar queue. Measured across the whole window on both venues' tapes, that rent was about two thirds of the coarse spread and came to $0.64 million on the smaller venue over 147 days and $1.0 million on the larger one in the thirteen coarse-grid days before it cut its own tick; and the adverse-selection share on the coarse book, a third to two fifths, was little more than half what it was on the fine book next door. What the coarse grid did cost was participation: on the fine grid the same venue trades four times the volume in trades half again as large, and its mid price, which used to move a hundred times a day, now moves twelve thousand times and can be measured as a contributor to price discovery for the first time.

Disclosure. The author operates a small proprietary testbed in the market studied and holds no position that depends on the findings. The recorders and the window detector are the author's own; the event definitions and thresholds are stated in full so that a reader with the same public feeds can reproduce the tables. Comments are welcome at accounts@moyaes.com.

References

Aquilina, M., Budish, E. and O'Neill, P. (2022). Quantifying the high-frequency trading "arms race". *Quarterly Journal of Economics*, 137(1), 493–564.

Board of Governors of the Federal Reserve System (2026). Federal Reserve issues FOMC statement, 16 September 2026.

Budish, E., Cramton, P. and Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *Quarterly Journal of Economics*, 130(4), 1547–1621.

Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical Association*, 74(366), 291–298.

Glosten, L. R. and Harris, L. E. (1988). Estimating the components of the bid/ask spread. *Journal of Financial Economics*, 21(1), 123–142.

Gonzalo, J. and Granger, C. (1995). Estimation of common long-memory components in cointegrated systems. *Journal of Business & Economic Statistics*, 13(1), 27–35.

Harris, L. (1994). Minimum price variations, discrete bid-ask spreads, and quotation sizes. *Review of Financial Studies*, 7(1), 149–178.

Hasbrouck, J. (1995). One security, many markets: Determining the contributions to price discovery. *Journal of Finance*, 50(4), 1175–1199.

Kozhan, R. and Viswanath-Natraj, G. (2021). Decentralized stablecoins and collateral risk. WBS Finance Group Research Paper.

Ma, Y., Zeng, Y. and Zhang, A. L. (2025). Stablecoin Runs and the Centralization of Arbitrage. NBER Working Paper No. 33882.

U.S. Securities and Exchange Commission (2018). Assessment of the Plan to Implement a Tick Size Pilot Program. Division of Economic and Risk Analysis.

Yao, C. and Ye, M. (2018). Why trading speed matters: A tale of queue rationing under price controls. *Review of Financial Studies*, 31(6), 2157–2183.

Zhang, W. (2026). How Much Money Fits Inside a Depeg? A Census of 1,136 Stablecoin Dislocations. SSRN Working Paper 7415499. https://ssrn.com/abstract=7415499