Ethereum Congestion and the Mean Reversion of DEX-CEX Price Spreads

Cory Wagaman-Eure

April 5, 2026

Abstract

This paper studies whether Ethereum congestion is associated with slower correction of price spreads between centralized and decentralized exchanges. Using minute-level data on Coinbase ETH price series, Uniswap v3 WETH-USDC price series, Ethereum Mainnet gas fees, realized volatility, flow, swap counts, and staleness for two fee tiers (5 bps and 30 bps), I test whether gas weakens the mean-reverting dynamics of the DEX-CEX spread. I first show that both spread series strongly reject a unit root and are therefore consistent with stationary, mean-reverting processes. I then estimate reduced-form spread-magnitude regressions and a dynamic mean-reversion specification in which the key coefficient is the interaction between lagged spread and lagged gas. In the full sample, higher gas is associated with slower correction in the 5 bps pool but not in the 30 bps pool. However, the 30 bps pool is also substantially less active and more stale relative to the CEX price. After restricting the sample to fresh, trade-active minutes, both pools exhibit a strong positive interaction between lagged spread and lagged gas, indicating that higher congestion is associated with slower correction of DEX-CEX mispricing. The results suggest that Ethereum congestion impairs arbitrage-driven price alignment, but that effect is easiest to detect when arbitrage can plausibly operate.

1. Introduction

A central promise of decentralized exchanges (DEXs) is that arbitrage should keep their prices closely linked with the prices on centralized exchanges (CEXs). In practice, however, DEX trading occurs on a blockchain, and therefore closely dependent on the frictions it brings. On Ethereum, traders must pay gas to compete for inclusion, and the base fee rises mechanically with demand when recent blocks are congested. Since a DEX-CEX arbitrage strategy typically combines nearly instantaneous CEX lag with a delayed and fee-sensitive on-chain lag, congestion should matter for how quickly mispricing is corrected.

This paper asks the question: does Ethereum congestion slow the mean reversion of the DEX-CEX spread for the same underlying asset? The question is economically important because the answer helps to distinguish between two stories about DEX inefficiency. One story is that DEX spreads mainly reflect structural liquidity and stale-price frictions. Another is that spreads are corrected by arbitrage, but congestion weakens the correction mechanism when execution becomes expensive.

The empirical design uses minute-level data combining Coinbase prices, Uniswap v3 prices for two fee tiers, Ethereum gas conditions, realized volatility, and several pool-activity controls. I start by ensuring that the spread series are statistically consistent with mean-reverting processes. I then estimate (i) a spread-magnitude regression and (ii) a dynamic mean-reversion regression where the slope on the lagged spread is allowed to vary with the lagged gas cost.

The main finding is that higher gas is associated with slower spread correction. In the full sample, this result is clear in the 5 bps pool but not in the 30 bps pool. However, once the sample is restricted to fresh, trade-active minutes, the effect appears strongly in both pools.

The contribution of this paper is not just that gas matters. It is that the impact of congestion is state dependent: the slowdown in mean reversion is most visible when the pool is actively trading and arbitrage can operate. That interpretation is necessary for interpreting the results of the 30 bps pool.

2. Background

2.1 Ethereum fees and congestion

Ever since EIP-1559, Ethereum transaction fees are made up of two main components: a protocol-determined base fee and a user-chosen priority fee (tip). The base fee rises when recent blocks are above the target size and falls when they are below it. The priority fee is the way for the sender to reflect the urgency of their transaction. This allows for competition for block inclusion. The total fee paid is `gas used x (base fee + priority fee)`. Because the base fee responds to recent block fullness and can increase by as much as 12.5% per block, it is a natural congestion proxy for use of the Ethereum network. Essentially, higher demand for block space raises the priority required for fast inclusion and mechanically pushes the next block's base fee.

2.2 AMM pricing, arbitrage and congestion

Unlike a limit order book, which has dominated traditional financial markets, a Uniswap pool does not update its price by matching bids and asks. Instead, the on-chain price is determined by the pool state and changes only when traders interact with the pool. In Uniswap v3, this takes the form of concentrated liquidity distributed across active ticks instead of a single global order book. Lower fee tiers are associated with tighter tick spacing, while higher fee tiers have wider spacing. As a result, the DEX price is not automatically equal to the external market price at every instant. It is pushed towards the efficient market price by trades.

The simplest intuition comes from the constant-product AMMs. If the efficient price of the underlying asset moves in an external market, the AMM's marginal price becomes temporarily misaligned. Arbitrageurs can then trade against the pool, moving it along its pricing curve until the on-chain price is again aligned with the external price net of arbitrage frictions. In the Uniswap v3 model, the mechanics are implemented through concentrated liquidity across ticks to increase capital efficiency of liquidity providers.

One thing to note is that, because price on DEXs lags the leading market price, liquidity providers on DEXs are constantly supplying liquidity at stale pool prices. This means they are constantly facing adverse selection by mechanism design of being passive providers. This constant gap between the value earned by an AMM liquidity provider and the benchmark of a continuously rebalancing benchmark portfolio is known as Loss-Versus-Rebalancing (LVR). In that sense, LVR is the liquidity-provider side of the same price-alignment mechanism studied in this paper.

3. Economic Theory

The economic benchmark for this paper is the law of one price. When the same underlying asset trades on multiple exchanges, their prices should not diverge persistently in an efficient market. When a spread opens, arbitrageurs can buy on the cheaper venue and sell on the more expensive venue, earning profits while pushing the prices back together.

However, price equality need not hold instantaneously because arbitrage is not frictionless. Trading venues have frictions specific to their market structures. When these frictions are high, execution may become more expensive and more uncertain. This is a limits-to-arbitrage setting: spreads may still be corrected, but they should correct more slowly when arbitrage is more costly.

If arbitrage is active, the spread should be a mean-reverting process, so the coefficient on the lagged spread should be negative. When frictions are higher, arbitrage is weakened, and therefore the restoring force becomes smaller in magnitude. This implies that the speed of spread correction should depend on the level of trading frictions facing arbitrageurs.

4. Relevant Literature

This paper sits at the intersection of five main literature strands: (i) cryptoasset arbitrage and market segmentation, (ii) AMM design and DEX microstructure, (iii) DEX price discovery and market quality, (iv) the economics of arbitrage against AMMs, and (v) execution risk and congestion in CEX-DEX arbitrage.

4.1 Arbitrage in crypto markets more broadly

A useful starting point is Makarov and Schoar (2020), who documented large and recurring price deviations across cryptocurrency exchanges, often persisting for hours, days, and even weeks. Their evidence shows that arbitrage in crypto is neither frictionless nor instantaneous, especially across different market structures. Although not about DEXs specifically, that paper lays the groundwork for existence of price misalignment caused by execution frictions.

4.2 AMM design and economics of DEXs

Lehar and Parlour (2025) provide one of the foundational empirical and theoretical papers on Uniswap as an automated market maker. Using the full history of early Uniswap interactions and comparing the AMM with a centralized limit order book, they document the absence of long-lived arbitrage opportunities and describe the conditions under which an AMM can compare favorably to a traditional order-book market. A related theoretical paper by Aoyagi and Ito (2025) studies the coexistence of limit order books and AMMs, showing that liquidity and trader composition interact across the two market forms. These papers together motivate why a DEX-CEX spread is useful for understanding coexistence of both market structures.

4.3 DEX price discovery and market quality

Capponi, Jia, and Yu (2026) show that high-fee DEX trades are more informative and contribute more to price discovery, consistent with a market in which informed trades bid up for faster execution. This connects blockchain fees to informed trader flow.

Alexander, Chen, Deng, and Fu (2025) study price discovery and efficiency in Uniswap liquidity pools. They find that Uniswap v3 is much more efficient than v2 and that some v3 pools approach or even exceed Bitstamp in price discovery ability. They also show that during periods of high uncertainty, traders prefer centralized exchanges, cross-exchange arbitrage can become more prevalent, and Uniswap's price discovery ability weakens relative to Coinbase and Bitstamp. That is especially relevant to this paper because it shows that market stress can increase cross-venue activity while also worsening DEX price efficiency.

Barbon and Ranaldo (2026) compare transaction costs and deviations from no-arbitrage on centralized and decentralized exchanges and give causal evidence that gas fees deteriorate DEX price efficiency and create persistent arbitrage deviations, especially for smaller trades. Their result is very aligned with my hypothesis that gas is not just a nuisance cost, but also a friction that can degrade price alignment.

4.4 Gas competition, execution priority, CEX-DEX arbitrage

Recent papers have moved even closer to the mechanism studied here. He, Yang, and Zhou (2025) model gas-fee competition between arbitrageurs and show that larger arbitrage opportunities and higher liquidity result in higher gas bids and larger trades. Wu, Sui, Thiery, and Pai (2025) provide a large-scale empirical study of CEX-DEX extracted value on Ethereum and estimate over 7.2 million arbitrages between August 2023 and March 2025, showing the centralization of searcher activity and importance of integration with block builders.

Taken all together, this literature implies that CEX-DEX arbitrage is highly sensitive to timing, inclusion priority, and gas costs. The present paper contributes to this literature by asking a narrower question: conditional on observing a spread, is the spread corrected more slowly when congestion is high?

5. Data and Variable Construction

The empirical analysis uses an aligned minute-level panel spanning 2025-01-01 00:00 UTC through 2025-06-30 23:59 UTC. The merged dataset contains 260,640 one-minute observations, corresponding to complete calendar-minute coverage over the full sample window. The panel combines Coinbase ETH price data, Uniswap v3 WETH-USDC minute price series for the 5 bps and 30 bps fee-tier pools, and Ethereum base-fee data. All series are aligned to a common minute timestamp and merged into a unified minute-level dataset.

The centralized-exchange series is constructed from Coinbase minute candles retrieved through the Coinbase REST API. The decentralized-exchange series is constructed from the Ethereum mainnet WETH-USDC Uniswap v3 pools, retrieved through The Graph. Ethereum congestion data are constructed from Ethereum mainnet block-level base-fee observations retrieved through an Alchemy-backed JSON-RPC client.

Let P_t^{CEX} denote the Coinbase minute price and let $P_t^{DEX,5}$ and $P_t^{DEX,30}$ denote the corresponding Uniswap v3 minute prices for the 5 bps and 30 bps pools. I construct both dollar price-difference spreads and basis-point spreads. The main spread measure in the paper is the basis-point wedge, defined as

$$wedge_t = 10{,}000 \times (\log P_t^{DEX} - \log P_t^{CEX})$$

I also compute realized annualized volatility from Coinbase log returns using a 30-minute rolling window. Ethereum congestion is measured primarily using the base fee per gas, expressed in both gwei and an estimated USD gas-cost proxy constructed from the base fee, the contemporaneous ETH/USD price, and a 200,000 gas-unit assumption. In addition, I construct a trailing 30-day congestion percentile from the gas-cost series. For each Uniswap pool, I retain measures of dollar flow, swap count, and staleness, where staleness is based on age since last trade.

I drop pool-specific observations flagged as price outliers or spike-patched prints, and I fill missing minute-level DEX flow and swap counts with zero when the interpretation is no swaps occurred this minute.

6. Empirical Strategy

6.1 Stationarity test

In order to talk about mean-reversion speed, we need to ensure that the spread does not behave like a unit root process. I therefore estimate Augmented Dickey-Fuller (ADF) tests for both spread series. The null is that the spread has a unit root. Rejection supports the interpretation that the spread is stationary and therefore mean reverting.

6.2 Spread-magnitude regression

I first estimate a reduced-form regression for the magnitude of the spread:

$$\log(1 + |wedge_t|) = \alpha + \beta \log(gas_t) + \Gamma X_t + \varepsilon_t$$

where X_t includes realized volatility, log flow, log swap count, log staleness, hour-of-day fixed effects, and day-of-week fixed effects. This specification asks whether high gas is associated with larger dislocations, but does not answer whether the dislocation corrects slower.

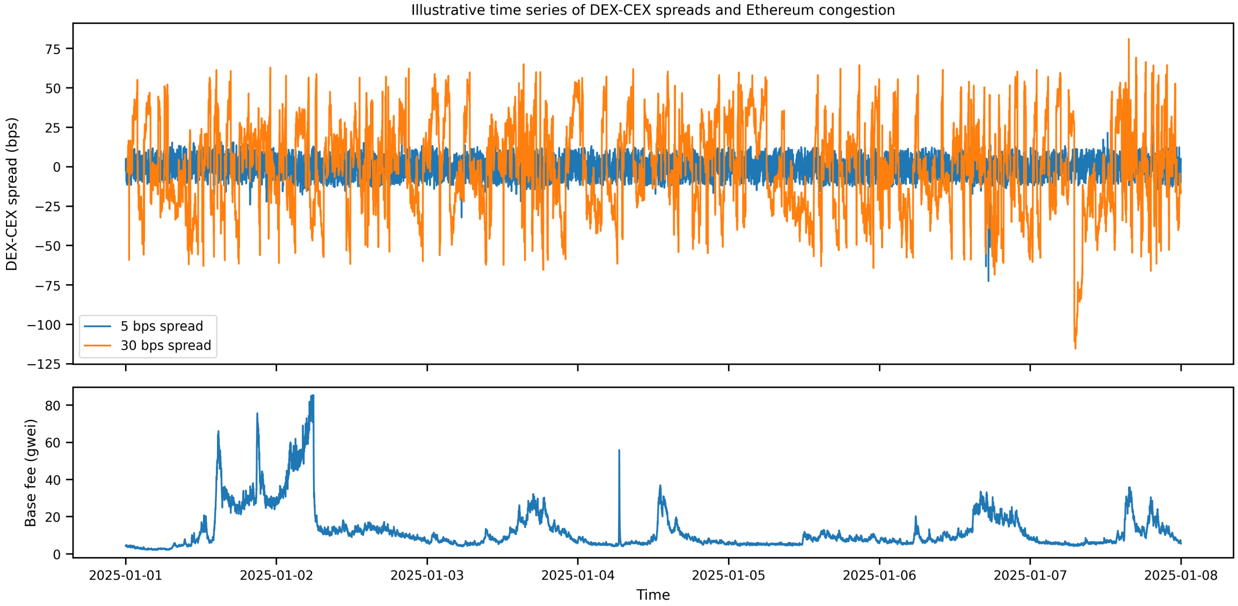

Figure 1: Illustrative time series of DEX-CEX spreads and Ethereum congestion. The top panel plots the 5 bps and 30 bps DEX-CEX spreads in basis points over a representative one-week window. The bottom panel plots Ethereum base fee in gwei over the same interval.

6.3 Dynamic mean-reversion regression

The core specification is:

$$\Delta wedge_t = \alpha + \beta \cdot wedge_{t-1} + \delta \cdot (wedge_{t-1} \times \log gas_{t-1}) + \Gamma X_{t-1}$$

The interpretation is that, if $\beta < 0$, the wedge mean reverts and, if $\delta > 0$, higher gas makes the effective reversion coefficient less negative meaning slower correction.

The effective restoring-force coefficient at a given gas level g is:

$$\beta_{eff}(g) = \beta + \delta \log g$$

All regressions are estimated with HAC standard errors using 60 lags.

6.4 Active-sample restriction

The descriptive statistics show that the 30 bps pool is much less active and much more stale relative to the 5 bps pool. That raises an identification problem: if the 30 bps pool often has no recent trades, then full-sample dynamics can be dominated by inactivity rather than arbitrage. To address this, I re-estimated the dynamic model on a restricted sample that keeps only fresh, trade-active minutes.

The active sample restricts attention to minutes in which pool staleness is at most 1 minute, age since last trade is at most 1 minute, and at least one trade occurs in the minute.

7. Results

7.1 ADF tests

Both spread series strongly reject the unit-root null. The ADF statistic is approximately -46.53 for the 5 bps spread and -48.60 for the 30 bps spread, with effectively zero p-values in both cases. These results do not imply a perfect Ornstein-Uhlenbeck structure, but they do provide strong evidence that the spread is not behaving like a random walk.

7.2 Descriptive difference across fee tiers

Table 1: Descriptive statistics by pool and sample

Sample	N	Mean spread (bps)	Mean |spread| (bps)	Median |spread| (bps)	Mean staleness (mi
5 bps full	260634	-0.751000	4.706000	3.575000	0.0650
5 bps active	248238	-0.757000	4.686000	3.526000	0.0000
30 bps full	260622	2.029000	21.773000	19.409000	8.4870
30 bps active	55471	0.161000	12.936000	6.117000	0.0000

Before turning to regressions, the two pools already look quite different. In raw summaries, the 30 bps pool exhibits much larger average spread magnitudes, much higher staleness, and many more no-trade minutes than the 5 bps pool. This immediately suggests that any comparison of coefficients across fee tiers must account for the fact that the 30 bps pool is not simply “the same market with a higher fee”. It is also a less active market state.

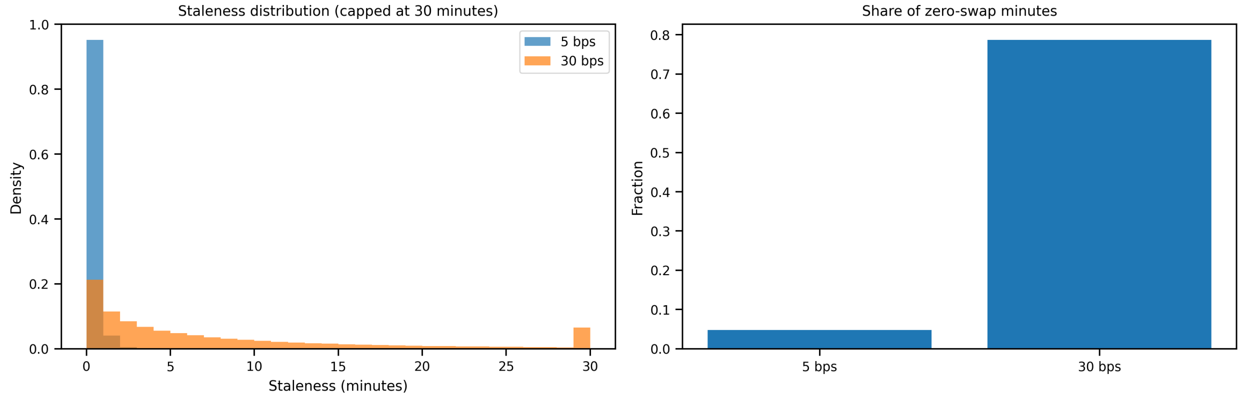

Figure 2: Pool activity and staleness by fee tier. The left panel shows the distribution of staleness for the 5 bps and 30 bps pools. The right panel shows the share of zero-swap minutes in each pool. The figure illustrates that the 30 bps pool is less active and more stale in the full sample.

7.3 Wedge-magnitude regressions

The level regressions show that in the 5 bps pool, higher gas is positively associated with a larger absolute spread. This is consistent with the idea that congestion hampers the speed or intensity of cross-venue price alignment. In the 30 bps pool, however, the gas coefficient is negative in the full sample, while staleness strongly predicts larger wedge magnitudes. That pattern initially suggests that the 30 bps pool may be dominated by stale-price frictions rather than congestion.

Table 2: Spread-magnitude regressions

Variable	5 bps	30 bps
log(Gas)	0.0819***	-0.0355***
	(0.0029)	(0.0064)
Realized volatility	0.1834***	0.2551***
	(0.0128)	(0.0119)
log(Flow USD)	-0.0511***	-0.1283***
	(0.0013)	(0.0016)
log(Swap count)	0.1588***	0.5101***
	(0.0049)	(0.0144)
log(Staleness)	-0.0434	0.2084***
	(0.0414)	(0.0045)
Hour fixed effects	Yes	Yes
Day fixed effects	Yes	Yes
Observations	260506	260469
R-squared	0.0439	0.2515

Notes: HAC standard errors in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.10$. All specifications include hour-of-day and day-of-week fixed effects.

7.4 Full-sample dynamic regressions

Table 3: Dynamic mean-reversion regressions

Variable	5 bps full	5 bps active	30 bps full	30 bps active
Lagged spread	-1.0410***	-1.0721***	-0.1823***	-1.0685***
	(0.0123)	(0.0110)	(0.0349)	(0.0161)
Lagged spread × lagged log(Gas)	0.1459***	0.1512***	-0.0044	0.1343***
	(0.0046)	(0.0045)	(0.0099)	(0.0069)
Lagged realized volatility	-1.3498*	-1.3527*	-1.5307**	-3.4533***
	(0.7852)	(0.8026)	(0.7569)	(1.3236)
Lagged log(Flow USD)	0.0687***	0.0542***	0.0852***	-0.0455
	(0.0191)	(0.0135)	(0.0307)	(0.0616)
Lagged log(Swap count)	0.2550	0.2568	-0.6373**	-0.0390
	(0.1858)	(0.1917)	(0.3120)	(0.6196)
Lagged log(Staleness)	1.2390**	0.7433**	-0.0280	0.0489
	(0.6273)	(0.3684)	(0.0917)	(0.2065)
Lagged log(Gas) included	Yes	Yes	Yes	Yes
Hour fixed effects	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes
Observations	260502	248114	260462	55454
R-squared	0.4156	0.4286	0.0941	0.5122

Notes: HAC standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.10. All specifications include lagged log gas, hour-of-day fixed effects, and day-of-week fixed effects. The active sample restricts attention to fresh, trade-active minutes.

For the 5 bps pool, the lagged spread coefficient is strongly negative (-1.0410), confirming fast mean reversion, and the interaction between lagged spread and lagged gas is strongly positive (0.1459). This means that as gas increases, the effective restoring force becomes less negative. Essentially, the spread still mean reverts, but it reverts more slowly when congestion is high.

For the 30 bps pool, the lagged spread coefficient is also negative (-0.1823), so the spread mean reverts in the full sample, but the interaction with lagged gas is close to zero and statistically insignificant (-0.0044). At face value, this would suggest that gas slows correction in the 5 bps pool but not in the 30 bps pool.

7.5 Active sample dynamic regression

That interpretation changes sharply once the sample is restricted to fresh trades, trade-active minutes.

In the 5 bps active sample, the dynamic result is essentially unchanged. The interaction stays strongly positive (0.1512), which confirms that higher gas weakens the restoring force.

In the 30 bps active sample, the result changes completely. The interaction becomes strongly positive and highly significant (0.1343). In other words, once stale and inactive minutes are removed, the 30 bps pool also shows slower correction when gas costs are high.

This is the central result of this paper. The full-sample null result for the 30 bps pool is not evidence that congestion is irrelevant here. Rather, it appears that inactivity and stale pricing mask the congestion effect when the sample includes long stretches in which arbitrage is not plausibly operating.

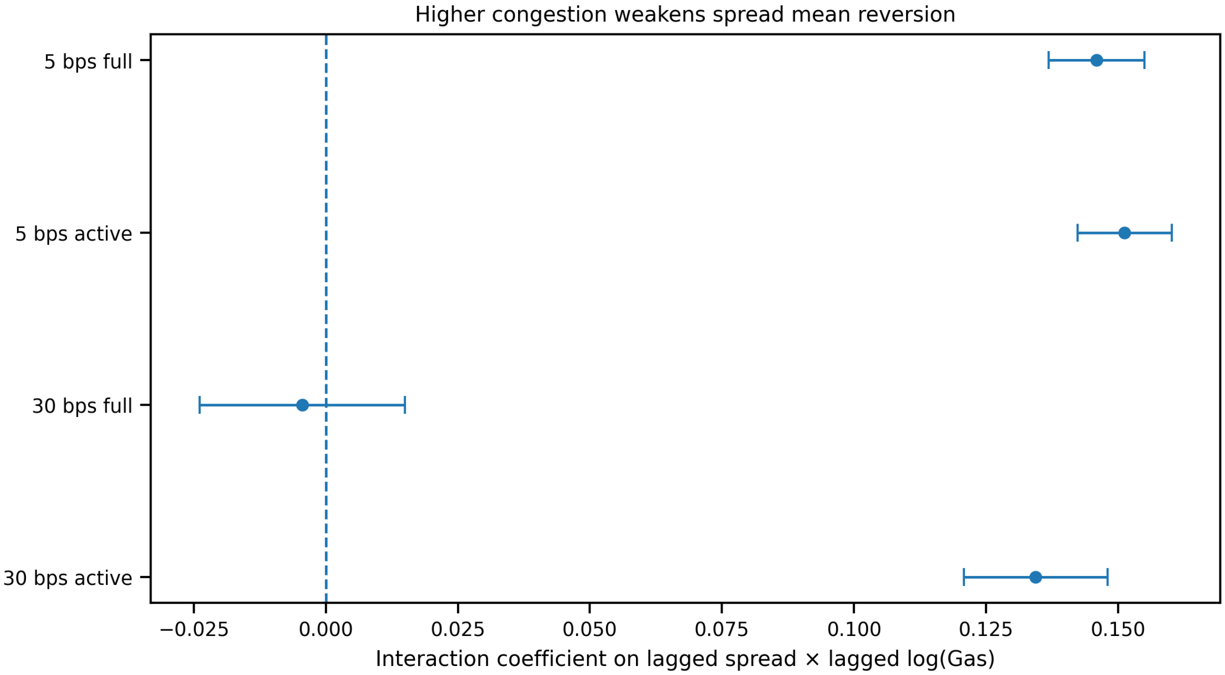

Figure 3: Congestion weakens spread mean reversion. The figure plots the estimated coefficient on the interaction between lagged spread and lagged log gas across the four main specifications, with 95% confidence intervals. Positive coefficients indicate that higher congestion makes the restoring force less negative.

7.6 Interpreting the effective coefficients

The active-sample coefficients are especially intuitive when evaluated at different gas quantiles.

For the 5 bps active sample:

- at the 25th percentile of gas, $\beta_{eff} \approx -1.0067$;
- at the 75th percentile of gas, $\beta_{eff} \approx -0.8645$.

For the 30 bps active sample:

- at the 25th percentile of gas, $\beta_{eff} \approx -1.0031$;
- at the 75th percentile of gas, $\beta_{eff} \approx -0.8704$.

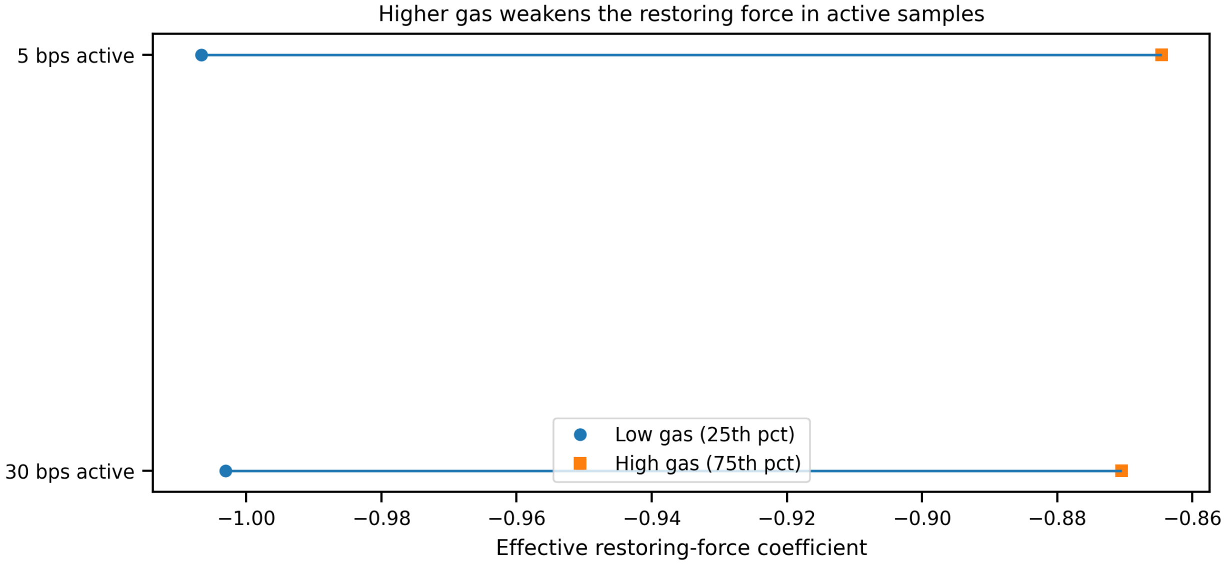

Figure 4: Effective restoring-force coefficient at low and high gas. For each active-sample specification, the figure plots the implied restoring-force coefficient evaluated at the 25th and 75th percentiles of lagged log gas. Higher gas shifts the coefficient toward zero, implying slower spread correction.

In both pools, higher gas shifts the coefficient toward zero. The spread is still mean reverting, but the pull back toward zero is materially weaker in high-congestion states.

8. Interpretation and Discussion

The evidence supports the interpretation that Ethereum congestion does not simply widen spreads mechanically. More importantly, it weakens the dynamic correction process that would normally pull the DEX prices back toward the CEX prices.

The active-sample result is especially important because it clarifies the role of the 30 bps pool. Once focusing on fresh, trade-active minutes, both pools exhibit the same qualitative congestion effect. This suggests that the apparent asymmetry across fee tiers was largely an artifact of state composition: the 30 bps full sample included far more inactive intervals.

This does not mean fee tiers are irrelevant. Higher-fee pools can still differ systematically in liquidity-provider incentives, routing behavior, and trade composition. But the results imply that one should be cautious about attributing a full-sample cross-pool difference to the fee tier alone. In this setting, the stronger claim is:

> conditional on the pool being active, higher Ethereum congestion is associated with slower correction of DEX-CEX mispricing in both fee tiers.

9. Limitations

This paper has four main limitations.

First, the analysis is reduced form. The results show that higher gas is associated with slower spread correction, but they do not identify the exact structural channel, such as builder integration, failed inclusion risk, or changes to searcher competition.

Second, fee tier and pool characteristics are partly confounded. The 30 bps pool differs not only in its fee, but also in activity, staleness, and likely liquidity composition. The active-sample design greatly improves interpretation, but it does not fully isolate the pure effect of fee tier.

Third, the congestion proxy is based on base fee rather than a full decomposition of transaction cost into base fee, priority fee, and realized execution success. Base fee is a natural and economically meaningful congestion measure, but it is not the entire cost of urgent DEX arbitrage.

Fourth, the dynamic coefficients imply very rapid adjustment in active samples. This should be interpreted as evidence of very fast correction, not as a precise structural half-life estimate under a fully specified continuous-time arbitrage model.

10. Conclusion and Future Research

This paper studies whether Ethereum congestion slows the correction of DEX-CEX mispricing. Using minute-level data on Coinbase, Uniswap v3, gas fees, and pool activity, I find that both spread series are stationary and mean reverting. In the full sample, higher gas weakens mean reversion in the 5 bps pool but not in the 30 bps pool. However, the 30 bps pool also contains many more stale and inactive intervals. Once the analysis is restricted to fresh, trade-active minutes, both fee tiers exhibit a strong positive interaction between lagged spread and lagged gas.

The main conclusion is therefore: Ethereum congestion is associated with slower correction of DEX-CEX spreads when arbitrage is actively operating. The result is not just a static spread-size effect. It is a dynamic effect on the speed of price alignment.

This finding matters for how DEX efficiency should be measured. Full-sample pool comparisons can understate the role of congestion if they mix together active arbitrage states and stale-price states. Conditioning on active trading reveals a more economically meaningful relationship between blockchain congestion and market quality.

A natural direction for future research is to move beyond reduced-form spread dynamics and study the execution channel more directly. Since the congestion effect in this paper is strongest when the pool is fresh and trade-active, future work should identify arbitrage episodes more explicitly and connect spread correction to realized execution costs, priority fees, and transaction inclusion outcomes rather than base fee alone. It would also be useful to test whether the same mechanism appears across other token pairs, AMM designs, and lower-latency blockchains. Doing so would help distinguish whether congestion slows price alignment primarily through higher trading costs, greater landing uncertainty, or broader changes in arbitrage participation.

References

Adadurov, Aleksei, Sergey Barseghyan, Anton Chtepine, Antero Eloranta, Andrei Sebyakin, and Arsenii Valitov. 2026. “Second Thoughts: How 1-second Subslots Transform CEX–DEX Arbitrage on Ethereum.” *arXiv preprint* arXiv:2601.00738.

Alexander, Carol, Xi Chen, Jun Deng, and Qi Fu. 2025. “Price Discovery and Efficiency in Uniswap Liquidity Pools.” *Journal of Futures Markets* 45 (8): 1023–1048.

Aoyagi, Jun, and Yuki Ito. 2025. “Limit Order Books and Automated Market Makers.” *Journal of Political Economy Microeconomics* 3 (3): 611–648.

Barbon, Andrea, and Angelo Ranaldo. 2026. “On the Quality of Cryptocurrency Markets: Centralized vs. Decentralized Exchanges.” *Management Science*, Articles in Advance.

Capponi, Agostino, Ruizhe Jia, and Shihao Yu. 2026. “Price Discovery on Decentralized Exchanges.” *Review of Financial Studies*, advance article.

Fritsch, Robin, and Andrea Canidio. 2024. “Measuring Arbitrage Losses and Profitability of AMM Liquidity.” In *Companion Proceedings of the ACM Web Conference 2024*, 1761–1767.

Harvey, Campbell R., Joel S. Hasbrouck, and Fahad Saleh. 2026. “The Evolution of Decentralized Exchange: Risks, Benefits, and Oversight.” *Research Policy* 55 (3).

Hasbrouck, Joel, Thomas J. Rivera, and Fahad Saleh. 2025. “An Economic Model of a Decentralized Exchange with Concentrated Liquidity.” *Management Science*, Articles in Advance.

He, Xue Dong, Chen Yang, and Yutian Zhou. 2025. “Arbitrage on Decentralized Exchanges.” SSRN Working Paper 5347630.

Heimbach, Lioba, Eric Schertenleib, and Roger Wattenhofer. 2022. “Exploring Price Accuracy on Uniswap V3 in Times of Distress.” In *Proceedings of the 2022 ACM CCS Workshop on Decentralized Finance and Security.*

Lehar, Alfred, and Christine A. Parlour. 2025. “Decentralized Exchange: The Uniswap Automated Market Maker.” *Journal of Finance* 80 (1): 321–374.

Lehar, Alfred, Christine A. Parlour, and Marius Zoican. 2024. “Fragmentation and Optimal Liquidity Supply on Decentralized Exchanges.” Working paper / SSRN 4267429.

Makarov, Igor, and Antoinette Schoar. 2020. “Trading and Arbitrage in Cryptocurrency Markets.” *Journal of Financial Economics* 135 (2): 293–319.

Milionis, Jason, Ciamac C. Moallemi, and Tim Roughgarden. 2025. “Automated Market Making and Arbitrage Profits in the Presence of Fees.” *arXiv preprint* arXiv:2305.14604, revised 2025.

Wu, Fei, Danning Sui, Thomas Thiery, and Mallesh Pai. 2025. “Measuring CEX–DEX Extracted Value and Searcher Profitability: The Darkest of the MEV Dark Forest.” *arXiv preprint* arXiv:2507.13023.

Appendix

Table 4: Augmented Dickey-Fuller test results

Series	ADF statistic	p-value	Lags used	Observations
5 bps spread	-46.525800	0.000000	86	260509
30 bps spread	-48.604800	0.000000	79	260463

Table 5: Full descriptive statistics

Sample	N	Mean spread (bps)	SD spread (bps)	Mean |spread| (bps)	Median |spread| (bps)
5 bps full	260634	-0.751000	10.509000	4.706000	3.575000
5 bps active	248238	-0.757000	10.657000	4.686000	3.526000
30 bps full	260622	2.029000	29.022000	21.773000	19.409000
30 bps active	55471	0.161000	28.042000	12.936000	6.117000

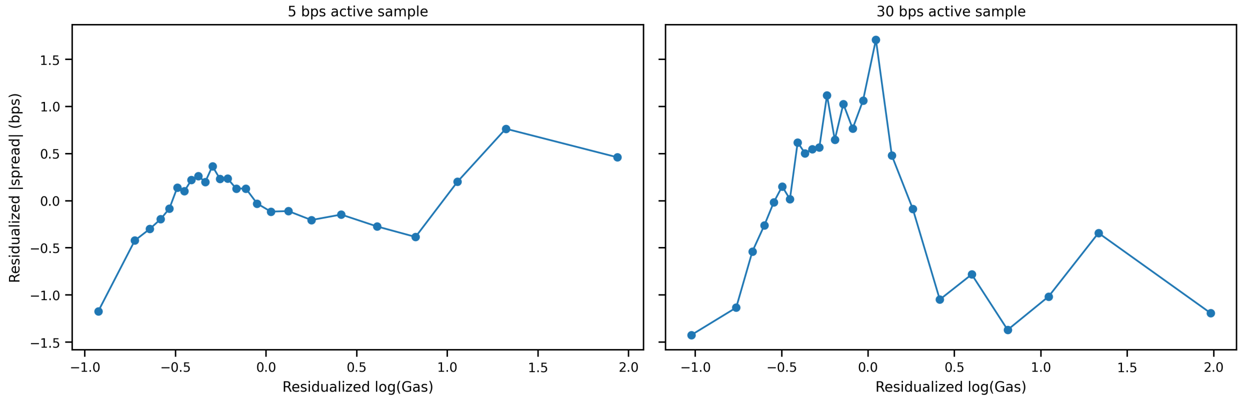

Figure 5: Residualized relationship between congestion and spread magnitude in the active sample. The figure plots binned averages of residualized absolute DEX-CEX spread against residualized log gas after partialling out realized volatility, flow, swap count, and time fixed effects.