

7536975-liquidity-capacity-and-token-survival-in-cryptocurrency-markets

SUMMARY: The study investigates the liquidity capacity and token survival in cryptocurrency markets, proposing a capacity-adjusted illiquidity measure that scales Amihud price impact by the square root of the product of market capitalization and the number of distinct takers. The research uses on-chain decentralized-exchange trade data to create survivorship-complete panels, including every token that clears a weekly volume threshold. The study finds that realized volatility has limited predictive power once liquidity is measured directly, and that the capacity adjustment sharpens this signal into a nearly sufficient statistic. The measure is also related to returns, with portfolios with higher death probabilities earning higher average returns. The findings suggest that token survival is primarily a liquidity-capacity problem, with price impact signals being more fragile when market capitalization is low and participation is thin.

METHOD: The study uses on-chain decentralized-exchange trade data from Dune Analytics and Uniswap liquidity events to reconstruct survivorship-complete panels for the primary Ethereum panel and a parallel Solana panel. The capacity-adjusted illiquidity measure is computed on these panels, scaling Amihud price impact by the square root of the product of market capitalization and the number of distinct takers. The study employs standard microstructure and asset-pricing measures, such as volume-weighted price and return, dollar volume, and the price impact of trading, to analyze the data.

KEY FINDINGS:

- Capacity-adjusted illiquidity is a nearly sufficient statistic for predicting token survival.
- Realized volatility has limited incremental predictive power once liquidity is measured directly.
- The measure is related to returns, with higher death probabilities associated with higher average returns.
- The measure is related to market impact, with trading activity declining, participation narrowing, and price impact rising before token death.

IMPLICATION FOR TRADING: The findings suggest that investors and market venues should focus on liquidity capacity when assessing cryptocurrency tokens. The study's capacity-adjusted illiquidity measure can be used to identify tokens with higher survival probabilities, which can inform investment decisions. Additionally, the measure's relationship to returns indicates that investors may benefit from holding tokens with higher death probabilities, as they tend to offer higher average returns. However, practitioners should be cautious and consider other factors, as the study's findings are based on historical data and may not hold in all market conditions.

Liquidity Capacity and Token Survival in Cryptocurrency Markets

Chang Liu

Department of Information Engineering and Computer Science, University of Trento
Via Sommarive 9, 38123 Povo (TN), Italy

chang.liu-1@unitn.it

August 16, 2026

Corresponding author. Chang Liu, Department of Information Engineering and Computer Science, University of Trento, Via Sommarive 9, 38123 Povo (TN), Italy. Email: chang.liu-1@unitn.it.

Declaration of competing interest. The author declares no competing financial or personal interests that could have influenced the work reported in this paper.

Data availability. All data are reconstructed from publicly available sources: the Dune Analytics `dex.trades` and `dex_solana.trades` datasets and public Uniswap liquidity events. No proprietary or paid data are used. The query definitions and the code that reproduce all tables and figures will be deposited in a public repository upon acceptance.

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Declaration of generative AI and AI-assisted technologies in the writing process. During the preparation of this work the author used AI-assisted tools (Claude, Anthropic) for language editing and code assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

CRedit authorship contribution statement. Chang Liu: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Liquidity Capacity and Token Survival in Cryptocurrency Markets

[Author name(s) on separate title page for anonymized review]*

August 16, 2026

Abstract

Cryptocurrency tokens are created permissionlessly and at near-zero cost, and most eventually stop trading. To predict which tokens survive, we propose capacity-adjusted illiquidity, a survival statistic that scales Amihud price impact by the square root of the market's absorption capacity—the product of its market capitalization and trader breadth. We compute it on survivorship-complete panels reconstructed from on-chain decentralized-exchange trades: the primary Ethereum panel includes every token that clears a weekly volume threshold, with a parallel Solana panel for partial external comparison. Liquidity, not volatility, predicts token death: realized volatility, emphasized in prior survival studies, has limited incremental predictive power once liquidity is measured directly. The capacity adjustment sharpens this signal into a nearly sufficient statistic: it recovers nearly all of the fit of a richer multivariate hazard and improves out-of-sample prediction of the rare weekly death event relative to raw illiquidity, with most of the gain from the market-value component. Its predictive content is consistent with a terminal liquidity deterioration: trading activity declines, participation narrows, and price impact rises before death. The measure is also related to returns: portfolios with higher death probabilities earn higher average returns, tracing a near-linear risk–return frontier. Throughout, we interpret the reported relationships as predictive associations rather than causal effects.

Keywords: cryptocurrency markets; decentralized exchanges; liquidity; token survival; market impact

JEL classification: G12, G14, G23

1 Introduction

Cryptocurrency tokens can be created permissionlessly and at near-zero cost. Many trade for a period and then fade below economically meaningful levels of activity, and prior evidence on survivorship bias indicates that token disappearance is common rather than exceptional (Ammann et al., 2022). Which tokens survive and which characteristics forecast disappearance matter for investors, for the venues that list these assets, and for the interpretation of cryptocurrency cross-sectional evidence. Our central argument is that token survival is primarily a liquidity-capacity

* Author affiliations and contact details are provided on a separate title page. The survivorship-complete on-chain panels are reconstructed from public Dune Analytics `dex.trades` and `dex_solana.trades` records and, for Ethereum, Uniswap liquidity events; the query definitions and the hazard-model code will be deposited in a public repository upon acceptance.

problem: the same price impact signals greater fragility when capitalization is low and participation is thin than when the market is large and broadly traded.

We make this argument precise with a capacity-adjusted illiquidity measure that scales Amihud price impact by the square root of the product of market capitalization and the number of distinct takers, the market’s absorption capacity. The square-root form is motivated by the square-root law of market impact.

Decentralized exchanges (DEXs) provide a comparatively complete record for the survival question. Because swaps are recorded on-chain, DEX trade data include tokens that later stop trading, so a cross-section reconstructed from these records is survivorship-complete by construction. The same record yields standard microstructure and asset-pricing measures at the token-week level: a volume-weighted price and return, dollar volume, the price impact of trading, and the breadth of participation; Uniswap liquidity events further yield the provision of liquidity in the underlying automated market maker (AMM). This design addresses two limitations of existing studies. A growing literature predicts cryptocurrency disappearance with survival and machine-learning models (Gatabazi et al., 2022; Fantazzini, 2022) and reports that low volume, small size, and high volatility raise the death hazard. Much of this work is built on centralized-exchange or aggregator data, which observe tokens only while listed and so condition on a form of survival (Shumway, 1997; Ammann et al., 2022), and predictors are often entered jointly without a direct measure of price impact among them, so an estimated volatility effect may proxy for unmeasured liquidity.

Our evidence comes in four parts. First, as motivation, liquidity rather than volatility predicts token death: on a harmonized sample in which volatility, illiquidity, capitalization, and participation are all measured, illiquidity carries a large death hazard, while realized volatility has limited incremental predictive power. Second, capacity-adjusted illiquidity captures, with a single covariate, nearly all of the fit of a richer multivariate hazard and improves out-of-sample prediction relative to raw illiquidity, capitalization, or participation alone. Third, the weeks before death show a joint deterioration: illiquidity rises while trading activity and participation decline together. Fourth, the measure is related to returns: capacity-adjusted illiquidity predicts the cross-section of subsequent returns, and across portfolios the probability of death and the average return trace a near-linear risk–return frontier. A parallel Solana panel provides a partial external comparison: the volume signal and the market-value adjustment carry over, while the Amihud and participation components are weaker there than on Ethereum.

We make four contributions. First, we give cryptocurrency markets a survival statistic: a liquidity-capacity measure that connects the Amihud price-impact tradition (Amihud, 2002) to the square-root law of market impact (Tóth et al., 2011; Almgren et al., 2005). The statistic is a parsimonious summary of survival risk: its weights are fixed ex ante rather than estimated, so it can be computed from a token’s first weeks of trading and compared across periods and chains, and it matches a freely estimated multivariate hazard out of sample. Second, we validate the statistic in and out of sample on a survivorship-complete on-chain panel constructed from DEX trades, predicting disappearance from directly measured liquidity rather than from survivor-biased listings or a high-dimensional model; in contrast to the emphasis of the crypto survival literature, volatility carries at most a small independent hazard once liquidity is measured. Third, relative to the decentralized-exchange microstructure literature (Lehar and Parlour, 2025; Barbon and Rinaldo, 2026), and to work connecting liquidity to crypto returns (Liu et al., 2022), we connect liquidity to token survival. Fourth, we show that the same capacity-adjusted statistic is also related to the cross-section of returns, aligning death probabilities and average returns along a common risk–return frontier. The analysis is limited to on-chain DEX activity over the sample window, and the

hazard estimates are predictive associations rather than causal effects.

The paper draws on three strands of literature. Cryptocurrencies are now studied as a distinct asset class (Liu et al., 2022; Cong et al., 2021), and much empirical work emphasizes their extreme price volatility (Kim et al., 2021) and return dynamics (Corbet and Katsiampa, 2020); we ask instead whether liquidity is more informative than volatility in explaining which tokens survive. The survival and machine-learning work noted above (Gatabazi et al., 2022; Fantazzini, 2022) models disappearance on centralized-exchange or aggregator data, while a parallel literature studies the financing and post-issuance fate of newly issued tokens (Howell et al., 2020) and the trading and arbitrage environment in which they change hands (Makarov and Schoar, 2020). Conditioning on survival biases inference about returns and characteristics (Ammann et al., 2022), data-design choices materially affect crypto cross-sections (Fieberg et al., 2024a), and economic constraints shape which anomalies appear (Fieberg et al., 2024b). We therefore reconstruct the cross-section from on-chain DEX trades and predict death from measured liquidity itself, in the failure-hazard tradition of Shumway (1997, 2001), Chava and Jarrow (2004), Bharath and Shumway (2008), and Lunde et al. (1999), applied here to a permissionless cross-section in which exit is the norm rather than the exception.

The third strand is the liquidity literature on which the measure builds. How best to measure liquidity in cryptocurrency markets remains an active question (Brauneis et al., 2021; Dong et al., 2022). Amihud illiquidity (Amihud, 2002) measures price impact per dollar traded, and low-frequency proxies such as the Amihud ratio track high-frequency benchmarks closely (Goyenko et al., 2009; Hasbrouck, 2009). Price impact is best interpreted relative to the market that absorbs the trade. The price-impact coefficient of Kyle (1985) falls with the standard deviation of uninformed order flow, and under the square-root law of market impact (Tóth et al., 2011; Almgren et al., 2005), price impact grows only with the square root of order size relative to the scale of trading it meets, so larger, more heavily traded markets absorb a given flow at smaller impact. This motivates measuring price impact against the scale—the capital and participation—of the market behind it. A closely related literature characterizes decentralized-exchange microstructure: Lehar and Parlour (2025) model the constant-product automated market maker and its price impact, Barbon and Rinaldo (2026) compare DEX and centralized-exchange market quality, and work on maximal extractable value (MEV) documents on-chain trading frictions (Daian et al., 2020). In an automated market maker, a token’s market has two sides—trade demand and liquidity provision—echoing the liquidity spirals of Brunnermeier and Pedersen (2009).

The remainder of the paper is organized as follows. Section 2 describes the data, the definition of death, and the hazard model; Section 3 develops capacity-adjusted illiquidity and the evidence on token death; Section 4 relates the measure to the cross-section of returns; Section 5 reports robustness within Ethereum, across venues, and on a second chain; and Section 6 concludes.

2 Data, token death, and the hazard model

2.1 The trading environment

The tokens we study are ERC-20 assets, the standard token format on Ethereum, and the trading we observe takes place on automated market makers rather than on order books. On Uniswap v2, the dominant venue for the long tail of small, thinly traded tokens, each token trades against a pooled counterparty: liquidity providers deposit a pair of assets—the token and a numeraire—into a shared pool, and traders (takers) swap against it at a price that moves along the pool’s constant-

product curve (Lehar and Parlour, 2025). Price impact is therefore mechanical in this market—a swap of a given size moves the price by more when the pool is shallow—so the price impact of trading measures how much order flow a token’s market can absorb, rather than serving only as a statistical proxy. This is the sense in which the Amihud ratio, and the capacity-adjusted statistic built on it, gauge the strain on a token’s market directly. Uniswap v3 refines the same design by letting providers concentrate liquidity within price ranges; the panel aggregates swaps across both versions and the other major Ethereum venues.

Three further institutional features recur. First, the numeraire: long-tail tokens are most often paired against wrapped ether (WETH), an on-chain token that represents ether one for one, or against major dollar stablecoins; Section 5 verifies the illiquidity result across both pairings. Second, liquidity provision is observable: providers add liquidity to a pool through Mint events and withdraw it through Burn events, which allows us to reconstruct a pool’s depth over time (Appendix A). Third, execution frictions: pending transactions are typically publicly visible before they settle, so on-chain traders can be front-run; in a sandwich attack, an attacker brackets a victim’s swap with trades of its own, capturing part of the victim’s price impact (Daian et al., 2020). Appendix A examines whether this friction, rather than illiquidity, is associated with token death.

2.2 The Ethereum survivorship-complete DEX panel

The primary panel is reconstructed from Ethereum DEX trades drawn from the public Dune Analytics `dex.trades` dataset, which decodes swaps across the major Ethereum automated-market-maker protocols, including Uniswap v2 and v3, and covers the tokens that traded on those venues. For every swap we observe the two token legs, their USD value, and the on-chain token amounts. Aggregating to the token-week, we form a volume-weighted average USD price, total DEX volume, and a trade count. Weekly returns are simple returns between consecutive weeks’ volume-weighted prices; a gap in activity leaves the return missing rather than interpolated. A token-week enters the panel when DEX volume exceeds a weekly floor of \$50,000. The panel therefore includes dead tokens for the weeks they were active and is survivorship-complete. Near-stablecoins—tokens with annualized return volatility below ten percent—are excluded, since a pegged price makes return-based liquidity measures and the volume margin of death uninformative; numeraire tokens such as WETH otherwise remain in the panel in their own right. The panel spans December 2019 through December 2025 and contains 6,515 distinct tokens, of which 6,096 permanently exit in sample; 6,049 of these exits can be dated to a death week with an observable one-week-lagged trading state, the requirement described below, and are eligible to enter the hazard as death events. Table 1 traces the construction of the estimation samples.

2.3 The definition of token death

A token is classified as dead at the last week in which it clears the volume floor with an observable one-week-lagged trading state, provided it then stays below the floor for at least four further weeks and does not resume sustained trading before the sample end. Death thus marks a permanent exit from economically meaningful DEX trading. Trading is intermittent, so a token’s final trade may be an isolated observation after a prolonged inactive period. We therefore date death to the last floor-clearing week that is immediately preceded by an active week, rather than to such an isolated trade. This choice measures the hazard’s one-week-lagged covariates from a genuinely adjacent

Table 1: Sample construction, Ethereum. The panel begins with every ERC-20 token recorded on the covered Ethereum DEX venues and is filtered to the hazard-estimation samples. The Deaths column counts dated deaths (6,049 of 6,096 permanent exits). Death is dated to a token’s last active week with an observable one-week-lagged trading state (Section 2), so the illiquidity filter retains every dated death; the capitalization-and-breadth row drops 368 deaths, and the deaths excluded by the final row are those of tokens with fewer than six weekly returns in the trailing window—too few to form a realized-volatility estimate. Deaths fall by more than tokens across the final row because the volatility screen removes token-weeks rather than whole tokens: 848 of the 950 excluded deaths leave with dropped tokens, and 102 are deaths of retained tokens whose death week fails the screen while earlier weeks remain (see Section 2). The final row is the harmonized sample used in Tables 3–5; the row above it is the broader sample of Appendix B, which drops the volatility requirement.

Step	Tokens	Token-weeks	Deaths
ERC-20 tokens ever traded on an Ethereum DEX	107,956	—	—
cleared weekly volume floor (\$50k) at least once	6,515	283,041	6,049
non-missing illiquidity, one-week lag	6,470	236,151	6,049
+ market cap and trader breadth (broader sample; App. B)	6,075	215,086	5,681
+ realized volatility (harmonized sample; Tables 3–5)	5,222	186,964	4,731

week.¹ Tokens that have not met the death criteria by the sample end—including those whose four-week confirmation window is cut off by it—are right-censored at their last active week. The definition pools terminal outcomes with different triggers: gradual fade-outs and creator-initiated exits such as rug pulls (Mazorra et al., 2022) are both treated as death, and the hazard estimates describe the pooled event. Because the death indicator is dated to an active week while every predictor enters with a one-week lag, the hazard’s covariates do not use the volume realizations that define exit. Appendix A decomposes Amihud illiquidity into its volume and return components (Table A.1), and Section 5 raises the volume floor that defines death (Table 7).

2.4 Variables

For each token-week we construct microstructure characteristics grouped by economic concept. Liquidity: Amihud illiquidity (Amihud, 2002), the ratio of absolute return to dollar volume, reported as price impact per billion dollars traded; log DEX volume; and log trade count. Capacity: log market capitalization—the weekly volume-weighted price multiplied by token supply, reconstructed as the running net of on-chain mints less burns—and trader breadth, the number of distinct takers transacting in the token that week. Risk and performance: realized volatility, the standard deviation of weekly returns over a trailing twelve-week window (at least six non-missing returns required); and the recent return, the token’s return over the preceding week; both are entered in levels rather than logs. Lifecycle: token age, in weeks since the token’s first active week. The venue analysis of Section 5 and the appendix analyses draw on additional variables—reconstructed pool depth and net liquidity flow from Mint and Burn events, trading concentration, and exposure to sandwich extraction—whose construction is described where each is first used.

¹Operationally, the one-week-lag requirement sets a covariate to missing whenever the preceding active week is not exactly seven days earlier; the death indicator is placed on the last active week that meets this requirement, and later isolated trades are dropped from the token’s risk set. The four-week confirmation window and this re-dating are designed to address trading intermittency directly.

Table 2: Summary statistics and correlations, harmonized sample. Panel A reports the distribution of the main lagged inputs to the death hazard (186,964 token-weeks, 4,731 deaths, 5,222 tokens). Illiquidity is the log Amihud ratio (price impact per billion dollars traded); capitalization, breadth, DEX volume, and age enter in logs. Panel B reports pairwise Pearson correlations. Realized volatility is nearly uncorrelated with illiquidity, and illiquidity and breadth are negatively correlated yet each retains an independent hazard in Table 4.

Panel A: distribution of lagged inputs						
Variable	Mean	SD	p25	Median	p75	
(1) Illiquidity (log Amihud)	4.86	2.17	3.67	5.20	6.42	
(2) Realized volatility	0.55	9.98	0.12	0.19	0.35	
(3) Log market capitalization	17.42	2.20	15.86	17.36	18.86	
(4) Log trader breadth	4.33	1.14	3.61	4.23	4.99	
(5) Log DEX volume	13.44	1.76	12.10	13.11	14.45	
(6) Log token age	3.76	0.95	3.00	3.78	4.52	
(7) Recent return	0.06	0.76	-0.12	-0.02	0.08	

Panel B: pairwise correlations							
	1	2	3	4	5	6	7
(1)	1.00						
(2)	0.02	1.00					
(3)	-0.50	-0.01	1.00				
(4)	-0.41	0.01	0.30	1.00			
(5)	-0.76	0.00	0.55	0.66	1.00		
(6)	-0.16	-0.02	0.30	-0.03	0.04	1.00	
(7)	0.08	0.05	0.09	0.04	0.12	-0.00	1.00

Table 2 reports the distribution of the main characteristics defined in this subsection and their pairwise correlations on the harmonized estimation sample defined below. The unwinsorized standard deviation of realized volatility is dominated by a small number of extreme return weeks; Section 5 examines log and winsorized treatments. Realized volatility is nearly uncorrelated with Amihud illiquidity, with a correlation of 0.02, suggesting that the two capture distinct cross-sectional variation; price impact and trader breadth are moderately negatively correlated yet each carries an independent hazard in the models below.

2.5 Estimation samples and coverage

Except for the Kaplan–Meier curves noted below, the analysis in Section 3 uses a single harmonized sample: the token-weeks for which realized volatility, Amihud illiquidity, market capitalization, and trader breadth are all available with a valid one-week lag, so that every model is estimated on identical rows. This sample contains 186,964 token-weeks, 4,731 deaths, and 5,222 tokens. Because volatility requires several weeks of history, the sample conditions on tokens that survive long enough to establish a trading record; the Kaplan–Meier curves below, which require only a duration and a single early-life characteristic, use the full panel rather than the harmonized sample. A broader sample that drops the volatility requirement, retaining token-weeks with illiquidity, capitalization, and breadth, contains 215,086 token-weeks; Appendix B shows the capacity-adjustment result also holds there. Market capitalization is the binding coverage constraint; trader breadth is available

for almost the entire lagged sample. Covered token-weeks carry a higher weekly death rate than uncovered ones (2.5 versus 1.4 percent per week), so the harmonized sample has a higher death rate than the full panel. Because death is dated to a week with an observable lagged trading state, the illiquidity filter retains every dated death. Requiring lagged capitalization and breadth drops 368 deaths, and the deaths excluded by the volatility requirement are those of tokens with fewer than six weekly returns in the trailing window—too few to form a volatility estimate.²

2.6 The discrete-time hazard

We model death in discrete time. The unit of observation is a token-week, and the event indicator equals one in a token’s last active week if it dies in sample and zero otherwise; surviving tokens are right-censored at their last active week. The weekly hazard is estimated with a logit link (Shumway, 2001),

$$\Pr(\text{death}_{i,t} = 1 \mid x_{i,t-1}) = \frac{\exp(\alpha + \beta' x_{i,t-1})}{1 + \exp(\alpha + \beta' x_{i,t-1})},$$

where $x_{i,t-1}$ is the vector of standardized characteristics measured in the prior week, α is a common intercept, and β the vector of hazard coefficients; specifications with year fixed effects replace α with calendar-year intercepts, as indicated in the tables. Death is thus predicted from lagged state, and the prediction uses no information from the final active week itself. Standardizing each predictor makes coefficients comparable, and $\exp(\beta)$ is the multiplicative change in the weekly odds of death per one-standard-deviation increase.³ Reported pseudo- R^2 values are McFadden’s, and in the out-of-sample tests both the fitted coefficients and the standardization moments come from the training window alone. Standard errors are two-way clustered by token and week (Petersen, 2009; Cameron et al., 2011); Section 5 reports a token-level Cox cross-check and alternative clusterings.

3 Capacity-adjusted illiquidity

This section presents the core evidence. A first pass establishes that liquidity, not volatility, predicts token death; we then argue that price impact alone is an incomplete measure of fragility, define capacity-adjusted illiquidity, evaluate it as a death predictor in and out of sample, and characterize the terminal deterioration it captures.

3.1 Liquidity, not volatility

Table 3 reports the baseline hazard on the harmonized sample, moving from volatility alone to a multivariate baseline (trader breadth enters in Table 4). Realized volatility on its own carries

²The token and death counts of Table 1 sit at different levels: a token is retained if *any* of its weeks passes a filter, while a death is retained only if the death week itself passes. Because the volatility screen removes token-weeks rather than whole tokens, a token can remain in the sample while its death event is excluded—which is why deaths fall by more than tokens in the final step. Of the 950 excluded deaths, 848 leave with the 853 tokens dropped entirely (five of these tokens had no datable death), and 102 are deaths of long-lived but intermittently trading tokens: the death-dating rule guarantees one adjacent prior active week, enough for the lagged trading state but not for the six returns volatility requires, so their death week fails the screen while their earlier, continuously traded weeks remain. These tokens enter the harmonized hazard as censored at their last valid week; the broader sample of Appendix B, which drops the volatility requirement, is unaffected by this loss.

³In the discrete-time logit, $\exp(\beta)$ is an odds ratio for the weekly hazard; following discrete-hazard convention, the tables label it a hazard ratio.

Table 3: Baseline death hazard: volatility versus illiquidity. Discrete-time logit hazard of token death on the harmonized sample (token-weeks for which realized volatility, Amihud illiquidity, market capitalization, and trader breadth are all available with a valid one-week lag). Death is the confirmed final active week (Section 2); all covariates are lagged one week and standardized. Entries are the hazard ratio per one-standard-deviation increase with two-way (token and week) clustered z -statistics in parentheses. Column 6 adds year fixed effects. The final row and the first table note report, respectively, out-of-sample ROC-AUC (train before 2024, test 2024–2025; the out-of-sample fits exclude the year effects) and an early-life token-level Cox check (first-eight-week medians). Column (5), labeled “Full,” is the multivariate baseline referenced in the text. Sample: 186,964 token-weeks, 4,731 deaths.

Predictor ($t - 1$)	(1)	(2)	(3)	(4)	(5)	(6)
	Volatility	Amihud	Vol.+Am.	Vol.+size+age	Full	+Year FE
Realized volatility	1.01 (1.6)		0.63*** (−3.9)	1.00 (0.2)	1.00 (1.0)	1.01 (1.6)
Illiquidity (Amihud)		4.98*** (27.0)	5.19*** (26.6)		1.78*** (7.8)	1.83*** (8.8)
Log DEX volume					0.32*** (−8.0)	0.38*** (−8.4)
Log trade count					0.83* (−1.8)	0.70*** (−5.0)
Log market cap				0.29*** (−32.0)	0.50*** (−17.1)	0.51*** (−17.0)
Recent return					1.10*** (5.4)	1.09*** (4.6)
Log token age				0.85*** (−3.1)	0.75*** (−6.5)	0.69*** (−11.0)
Year fixed effects	No	No	No	No	No	Yes
Pseudo- R^2	0.000	0.128	0.129	0.136	0.210	0.219
OOS ROC-AUC	0.643	0.796	0.795	0.786	0.846	0.846

First-8wk token-level Cox on the full panel ($n = 6,114$ tokens): illiquidity $z = +16.2^{***}$, volatility $z = -0.5$. Hazard ratio per SD; z two-way clustered by token and week. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

only a weak hazard, while Amihud illiquidity carries a large one. When the two enter together, illiquidity’s hazard is essentially unchanged, while volatility’s coefficient flips to significantly below one—the specification dependence that Section 5 traces to the level scaling. In the multivariate baseline model the illiquidity hazard remains large alongside volume, size, and age, and the result is unaffected by year fixed effects. The death hazard is far more strongly associated with the price impact of trading than with the volatility of price. Nor is the illiquidity hazard a restatement of the volume floor that defines death: it remains large and significant conditional on log volume and trade count themselves (Table 3, columns (5) and (6)). Two checks guard against the concern that the illiquidity hazard merely reflects measurement in the weeks adjacent to exit. First, a token-level Cox model on each token’s first four weeks—using no information beyond each token’s first month—reproduces the illiquidity hazard on lifespans (Table 7, Panel C; an eight-week version appears in the note to Table 3). Second, out-of-sample discrimination of the liquidity signal—measured by the capacity-adjusted statistic developed below—stays above chance when the covariate lag is lengthened to twelve weeks (Appendix Table B.5).

Realized volatility, the death predictor emphasized by the prior survival literature, is predictively weak here: in levels its hazard ratios are close to one both univariately and in the multivari-




Figure 1: Kaplan–Meier survival curves by Amihud illiquidity tercile, where each token is assigned to a tercile by the median of its available illiquidity observations over its first four weeks of trading, so the grouping uses no information beyond each token’s first month; tokens with no illiquidity observation in that window are excluded. The horizontal axis is weeks since a token’s first active week and the vertical axis is the probability of remaining active. Less liquid tokens exit substantially sooner.

ate baseline, its out-of-sample area under the receiver operating characteristic (ROC) curve is only 0.643, and its sign is not stable across specifications. Nor is the contrast an artifact of collinearity or of scaling: realized volatility and Amihud illiquidity are nearly orthogonal in the cross-section (Table 2), and under log or winsorized transformations volatility carries a modest but significant independent hazard, while illiquidity remains dominant and volatility’s out-of-sample discrimination is unchanged (Section 5, Appendix Table B.7).

Figure 1 shows the illiquidity–survival ranking nonparametrically. Sorting tokens into terciles by their Amihud illiquidity over their first four weeks of trading, so that the grouping uses no information beyond a token’s first month, produces a wide and monotone separation in Kaplan–Meier survival curves. The least liquid tercile falls to a small surviving fraction within months of first trading, the middle tercile decays steadily, and the most liquid tercile retains a substantial surviving fraction after years. The separation requires no model and no controls.

Price impact thus provides a strong and robust ranking of survival risk, one that is not driven by measurement in the weeks adjacent to exit. What it omits is the capacity of the market on the other side of the trade; the next subsections build that dimension into a single survival statistic.

3.2 Why capacity matters

The first-pass evidence identifies illiquidity as a central survival predictor, but price impact alone does not fully characterize token fragility. Amihud illiquidity measures how far the price moves per dollar traded. The same price impact signals greater fragility in a small, thinly traded market—little

capitalization, few active participants standing behind the market—than in a large, broadly traded one. Two tokens with identical Amihud illiquidity but different capitalization and participation need not carry the same survival risk. Two distinct dimensions are associated with whether a token continues to trade: the price impact of a given trade, and the capacity—capitalization and breadth of participation—that absorbs it. In our data these quantities are related but far from collinear—the correlation between price impact and breadth is -0.41 on the harmonized sample (Table 2)—yet each carries an independent death hazard when entered jointly below (Table 4). A scalar liquidity measure that captures price impact but omits capacity may miss a dimension of fragility. Capitalization captures economic scale and trader breadth captures participation capacity, and a measure of survival risk should combine price impact with both.

3.3 The measure

We define capacity-adjusted illiquidity (CAI) as

$$\Lambda_{i,t}^{\text{adj}} = \frac{\text{Amihud}_{i,t}}{\sqrt{M_{i,t} N_{i,t}}},$$

where, for token i in week t , $M_{i,t}$ is market capitalization and $N_{i,t}$ the number of distinct takers. The numerator is realized price impact; the denominator is the market’s absorption capacity. The square-root scaling is motivated by the square-root law of market impact (Tóth et al., 2011; Almgren et al., 2005), under which price impact grows with the square root of order size relative to the scale of trading that absorbs it; scaling by $\sqrt{M_{i,t} N_{i,t}}$ treats market capitalization and participation as that scale. Out-of-sample performance is similar across a grid of exponents (Section 5), indicating weak identification of the exact scaling parameters; the square-root form lies within the best-performing region. The measure is built entirely from liquidity and scale quantities, with no realized-volatility input beyond the absolute return inside the Amihud numerator, whose secondary role the decomposition of Appendix A documents. Empirically we use the log form,

$$\log \Lambda_{i,t}^{\text{adj}} = \log \text{Amihud}_{i,t} - \frac{1}{2} \log M_{i,t} - \frac{1}{2} \log N_{i,t},$$

standardized before estimation, so it enters on the same log scale as the inputs of Tables 2 and 4.

3.4 Predicting token death

Table 4 shows that the components carry independent hazards and that the single ratio captures most of their joint information. Entered freely, illiquidity, capitalization, and trader breadth each predict death with the expected sign, and all three remain significant when entered jointly, indicating that the components capture distinct variation. Capacity-adjusted illiquidity enters as one standardized covariate and attains a pseudo- R^2 of 0.195, against 0.201 for the additive model with the three components (column (6)) and 0.134 for Amihud illiquidity alone: a single statistic recovers nearly all of the in-sample fit of the additive liquidity-capacity model. The component pattern is unchanged when volatility and lifecycle controls are added (column (7)).

In out-of-sample tests, the measure adds predictive value beyond raw illiquidity. Table 5 estimates each predictor on token-weeks before 2024 and evaluates it on 2024–2025. Capacity-adjusted illiquidity raises the precision-recall area, an important metric for the rare death event, by roughly thirty-five percent relative to Amihud illiquidity and improves the area under the ROC curve (AUC) as well; a week-block bootstrap places both gains above zero. Its out-of-sample performance is comparable to that of a hazard that estimates the liquidity, capacity, age, and volatility loadings freely:

Table 4: Capacity-adjusted illiquidity: in-sample hazard. Discrete-time logit hazard of token death on the harmonized sample; covariates lagged one week and standardized, hazard ratio per standard deviation with two-way (token and week) clustered z in parentheses, year fixed effects in all columns. Year fixed effects are included in every column so that the pseudo- R^2 comparison across columns shares a common calendar baseline; because of the fixed effects, and because columns (6)–(7) omit the volume and trade-count controls of Table 3’s multivariate baseline, the Amihud hazards here differ from Table 3’s. $\Lambda^{\text{adj}} = \text{Amihud}/\sqrt{MN}$; column (4) scales by capitalization alone ($A/\sqrt{M}$) and column (5) is capacity-adjusted illiquidity (CAI). With a single covariate, capacity-adjusted illiquidity recovers nearly all of the in-sample fit of the additive liquidity-capacity model (column (6)). Sample: 186,964 token-weeks, 4,731 deaths.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Amihud	Size	Breadth	$A/\sqrt{M}$	CAI	Additive	Free
Illiquidity (Amihud)	5.25*** (27.6)					2.92*** (20.0)	2.85*** (20.7)
Log market cap		0.26*** (−27.3)				0.39*** (−20.0)	0.45*** (−20.7)
Log takers (N)			0.50*** (−20.9)			0.60*** (−11.5)	0.59*** (−13.2)
Amihud/ $\sqrt{M}$				6.75*** (28.3)			
Λ^{adj}					7.87*** (28.6)		
Realized volatility							1.00 (0.3)
Log token age							0.74*** (−7.7)
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo- R^2	0.134	0.140	0.051	0.182	0.195	0.201	0.207

Hazard ratio per SD; two-way clustered z . * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the two ROC areas differ by 0.002, within bootstrap noise, and the two bootstrap gains over Amihud are of similar size. The measure’s main advantage is parsimony: a single covariate performs similarly to the free multivariate model and better than raw price impact. Appendix Table B.4 repeats the comparison with rolling training origins on the broader sample, testing on 2023, 2024, and 2025 in turn; the capacity-adjusted measure improves on raw Amihud in every fold. Table 5 also shows that the market-value component of the adjustment accounts for most of the gain, while trader breadth mainly refines the ranking, a decomposition we take up next.

Capacity adjustment captures information beyond size alone (Table 5). Capitalization alone improves precision-recall relative to raw Amihud but not rank discrimination, and it does not by itself recover the joint price-impact-and-capacity signal. Both capacity adjustments, scaling by $\sqrt{M}$ and by $\sqrt{MN}$, improve on raw Amihud in both the precision-recall and ROC areas and remain above the capitalization-only benchmark. This suggests that the predictive signal comes from combining price impact with capacity rather than from scale alone.

Across the components, market-value capacity is the dominant adjustment and trader breadth is a second-order refinement. Scaling by capitalization alone, $\text{Amihud}/\sqrt{M}$, already captures most of the gain: its out-of-sample precision-recall is comparable to that of the full measure, and its area under the ROC curve is only marginally lower (Table 5). Adding trader breadth sharpens

Table 5: Out-of-sample horse race. Each row is a discrete-time logit trained on token-weeks before 2024 and evaluated on 2024–2025 (harmonized sample); the final row enters the listed covariates jointly. Reported: out-of-sample AUC and precision-recall area (PR-AUC), differences versus Amihud illiquidity with 95% week-block bootstrap intervals. Test sample: 81,317 token-weeks, 2,502 deaths.

Predictor	AUC	PR-AUC	Δ PR-AUC vs Amihud	Δ AUC vs Amihud
Volatility	0.643	0.049	−0.070 [−0.081, −0.060]	−0.153 [−0.175, −0.134]
Amihud $ r /V$	0.796	0.119	—	—
Size ($-\log M$)	0.795	0.143	+0.025 [+0.012, +0.038]	−0.001 [−0.014, +0.011]
Breadth ($-\log N$)	0.675	0.052	−0.068 [−0.077, −0.057]	−0.122 [−0.139, −0.105]
Amihud/ $\sqrt{M}$	0.829	0.166	+0.047 [+0.038, +0.056]	+0.033 [+0.026, +0.039]
$\Lambda^{\text{adj}} = \text{Amihud}/\sqrt{MN}$	0.838	0.162	+0.043 [+0.035, +0.050]	+0.042 [+0.035, +0.048]
Free multivariate (+vol, age)	0.836	0.152	+0.033 [+0.024, +0.044]	+0.040 [+0.026, +0.051]

Week-block bootstrap (150 resamples); differences whose intervals exclude zero are significant at 5%.

Differences are computed before rounding.

rank discrimination (area under the ROC curve 0.829 to 0.838) at a slight cost in out-of-sample precision-recall (0.166 to 0.162) on the harmonized sample. In a paired week-block bootstrap, the rank-discrimination gain of the full measure over Amihud/ $\sqrt{M}$ excludes zero while the precision-recall difference does not; on the broader sample of Appendix B both differences favor the full measure and exclude zero (both untabulated). This distinction helps interpret the cross-chain evidence of Section 5.3, where the capitalization adjustment carries over to Solana more clearly than the participation refinement. On first-four-week values alone, by contrast, the three illiquidity variants rank eventual deaths about equally well (Appendix Table B.6, Panel B): the capacity adjustment contributes through week-ahead timing rather than early-life discrimination.

3.5 What the measure captures

Capacity-adjusted illiquidity is a predictive statistic; this subsection examines what it captures economically. Price impact and participation—two of the three quantities it combines—deteriorate jointly in the weeks around exit, alongside collapsing volume.⁴ Figure 2 traces the median dying token over the twelve weeks before its last active week: Amihud illiquidity rises, DEX volume declines, and participation narrows in the final weeks, with the median number of distinct weekly takers falling by roughly a third. The pattern recalls the logic of the liquidity spirals of Brunnermeier and Pedersen (2009), although the evidence points to concurrent deterioration rather than an estimated feedback loop. The level of illiquidity, rather than its recent change, predicts death: conditional on the level, the recent change adds little predictive content (untabulated).

Appendix A documents the mechanics behind this pattern. A decomposition of the Amihud ratio shows that low dollar volume carries most of its univariate signal, with the ratio serving as a compact price-impact summary (Table A.1). Reconstructed liquidity provision on Uniswap v2 adds no clean lagged withdrawal channel beyond illiquidity itself (Table A.2). And exposure to sandwich extraction marks liquid, actively traded tokens rather than tokens close to exit, so the trading friction most closely associated with death is illiquidity, not measured front-running (Table A.3).

⁴These analyses use contemporaneous and event-time measures rather than the lagged covariates of the predictive models, because these analyses characterize the deterioration around exit rather than predict it one week ahead.




Figure 2: The terminal deterioration before token death. Median log Amihud illiquidity (left axis) and median log DEX volume (right axis) in the top panel, and the median number of distinct weekly takers in the bottom panel, over the twelve weeks preceding a token’s death, across all dying tokens with data at each horizon. Price impact rises while trading activity and participation decline as death approaches.

4 The capacity-adjusted illiquidity premium

If capacity-adjusted illiquidity measures a token’s exposure to the risk of terminal liquidity deterioration, it may also be priced. Tokens with high capacity-adjusted illiquidity are more likely to die; if investors are compensated for bearing this risk, such tokens should earn higher expected returns (Amihud and Mendelson, 1986; Acharya and Pedersen, 2005). We test this on the cross-section of weekly token returns.

4.1 Cross-sectional return predictability

Table 6 reports Fama–MacBeth (Fama and MacBeth, 1973) weekly cross-sectional regressions of the next-week return on lagged, standardized capacity-adjusted illiquidity. The slope is positive and strongly significant, with a Newey–West (1987) t (four lags) above eight univariately and above five once Amihud illiquidity, size, and volatility are controlled: a one-standard-deviation increase in capacity-adjusted illiquidity is followed by a higher return. The conditional slope exceeds the univariate one because, with Amihud illiquidity and size held fixed, the remaining variation in the

Table 6: Capacity-adjusted illiquidity and the cross-section of returns. Panel A: Fama–MacBeth weekly regressions of the next-week return on standardized lagged predictors, average slope per SD with Newey–West t (4 lags). “Multivariate” adds Amihud illiquidity, log market cap, and realized volatility; the next row adds reversal and momentum; “skip-week” lags CAI an extra week. Panel B sorts tokens weekly into deciles by lagged CAI (equal-weighted next-week return, % per week); D1 and D10 are pooled token-week means, while the spread is the time-series average of the weekly D10–D1 spread, so it differs from the difference of the pooled decile means. Forward returns winsorized at 1/99%; the return levels are inflated by micro-capitalization microstructure and the portfolios are not tradable—the content is the sign, the significance, and the overall gradient.

Panel A: Fama–MacBeth slopes (per SD)		
Specification	Avg. slope	NW t
CAI, univariate	+0.020	+8.50
CAI, multivariate	+0.075	+5.34
CAI, + reversal, momentum	+0.063	+4.54
CAI, skip-week	+0.016	+6.10
Panel B: decile sort on capacity-adjusted illiquidity (% per week)		
D1 (lowest CAI)	+0.12	
D10 (highest CAI)	+10.56	
D10 – D1 spread	+9.44	($t = +7.95$)

statistic comes from its participation component. The premium remains after adding reversal and momentum, which makes a purely mechanical bid-ask-bounce explanation less likely. It also remains in a skip-week specification that lags the predictor an additional week, though the magnitude attenuates by roughly a fifth relative to the univariate slope, consistent with a small short-horizon microstructure component and a larger persistent one. Dying tokens require a stricter treatment. A dead token’s final active week has no subsequent in-panel price, so its forward return is missing by construction; Appendix Table D.2 therefore re-estimates the premium, assigning a forward return of –100 percent to every death week. Under this bound the decile spread is roughly halved and the Fama–MacBeth slope attenuates, but both remain significant, so the premium is not solely due to the missing forward returns of dying tokens.

4.2 The risk–return frontier

A portfolio sort provides a clear economic interpretation. Sorting tokens each week into deciles by lagged capacity-adjusted illiquidity, the mean next-week return rises steeply from the lowest- to the highest-CAI decile (Table 6, Panel B reports the extreme deciles; Figure 3 plots all ten). The figure charts each decile’s mean forward return against its weekly death probability, and returns rise with it across the deciles. The relation is close to linear: a regression of decile return on decile death probability has an R^2 of 0.97. The liquidity-capacity statistic orders both outcomes across portfolios—the probability that a token dies and the return earned by those that survive—and aligns them along a common risk–return gradient. The sign of the alignment contrasts with the negative distress-risk premium documented for equities (Campbell et al., 2008). Capacity-adjusted illiquidity is thus not only a survival predictor; it also identifies a return-related dimension along which death risk and average returns move together.




Figure 3: The capacity-adjusted illiquidity risk–return frontier. Tokens are sorted each week into deciles by lagged capacity-adjusted illiquidity; the horizontal axis is the decile’s average weekly death probability and the vertical axis its average next-week return. Higher capacity-adjusted illiquidity is associated with both a higher death probability and a higher average next-week return, and the deciles fall on a single near-linear frontier ($R^2 = 0.97$). Forward returns are winsorized at the 1st and 99th percentiles; the return levels are inflated by micro-capitalization microstructure and the portfolios are not tradable.

4.3 Robustness of the premium

The premium is robust to factor adjustment and across size groups and subperiods, and has a Solana analogue based on $\text{Amihud}/\sqrt{N}$ (Appendix Table D.1). The CAI-sorted long-short earns a statistically significant alpha after controlling for an equal-weighted market return and size and momentum long-short factors, so it is not spanned by in-panel analogues of the standard crypto factors. The Fama–MacBeth slope is positive and significant in every calendar subperiod. On the Solana chain, where market capitalization is unavailable for the full token universe (Section 5.3), the participation-scaled measure $\text{Amihud}/\sqrt{N}$ stands in and again predicts the cross-section of returns. The premium is also present outside the micro-capitalization tail. When the weekly sort is restricted to tokens above the median market capitalization, the quintile long-short earns 5.6 percent per week with a Newey–West t of 8.6, and the Fama–MacBeth slope remains strongly significant; restricting to the top size tercile leaves both intact at similar magnitudes. As with the decile returns of Figure 3, the levels of these long-short returns are inflated by thin-market microstructure (Blume and Stambaugh, 1983; Asparouhova et al., 2013) and should not be read as tradable net of costs; the informative content is the sign, the statistical significance, and the alignment with death risk.

5 Robustness and external validity

5.1 Robustness of the Ethereum result

Table 7 reports subperiod, threshold, and specification checks. In the parsimonious subperiod models, illiquidity is significant in every tabulated subperiod and volatility in two of the three—in 2021 with strength comparable to illiquidity’s; in the multivariate baseline model and across the specifications of Table 8, illiquidity remains dominant while volatility is not robustly signed. The illiquidity ranking is preserved as the volume floor defining death is raised, and a token-level Cox proportional-hazards model (Cox, 1972) reproduces the illiquidity hazard on lifespans; the eight-week Cox in the note to Table 3 also ranks illiquidity far above volatility. Because realized volatility enters the baseline in levels while the liquidity and capacity inputs are logged, Appendix Table B.7 re-estimates the hazard with log and winsorized volatility. Under log and winsorized transformations, volatility carries a modest but significant independent hazard (z of +6.0 and +4.5, against +1.2 in levels, in the breadth-augmented specification of that table), indicating that the level scaling partly accounts for the weak independent volatility coefficient. Across all three treatments, however, illiquidity remains the stronger predictor, and volatility’s out-of-sample discrimination is weak and unchanged, because the area under the ROC curve depends only on ranks, which the log transform preserves exactly; the winsorized AUC is unchanged here as well. Table 8 shows the ranking is unchanged across lagged and contemporaneous covariates, year fixed effects, clustering by token alone, and restriction to single-spell tokens (tokens whose active weeks form a single contiguous run). The capacity-adjusted measure is robust in univariate form (Appendix Table B.3): its hazard ratio is stable across subperiods, two-way clustering by token and week gives the most conservative standard errors, and a look-ahead-free token-level Cox model on first-four-week values is strongly significant. Out-of-sample performance is also similar across the exponent grid (Appendix Table B.2), which identifies the exponents only weakly; the square-root form remains among the best-performing cells.

5.2 Cross-venue robustness

The illiquidity result should not hinge on a particular market design or pairing asset. We therefore repeat the analysis across Uniswap v2 and v3—the two dominant automated-market-maker designs—and across pools paired with WETH and with major stablecoins, the two numeraires against which these tokens most often trade. Table 9 reports, for each venue cell, the hazard on Amihud illiquidity and on reconstructed pool depth, defined as the cumulative net liquidity supplied on the pool’s numeraire side (WETH or stablecoin, per cell) following the construction of Appendix A. Illiquidity predicts death in all four cells, so the illiquidity result is present across market designs and numeraires. The reconstructed depth measure, by contrast, gives no consistent signal—insignificant in the WETH cells and significantly positive in the stablecoin cells—which we read as a limitation of the proxy: a cumulative net-flow measure cannot capture the concentrated, range-bound liquidity of v3, and a cumulative net numeraire-side flow is at best a coarse stand-in for effective depth. Depth is uninformative even in the v2–WETH cell, where the proxy is best measured. The venue cells therefore speak to the robustness of the illiquidity result, not to a depth channel.

Table 7: Robustness across subperiods, volume floors, and survival-model choice (lagged $t - 1$ covariates). Panels A and B report hazard z -statistics on illiquidity and volatility (controlling for volume and size); the sparse 2020 start of the sample (3,992 token-weeks, 80 deaths) is not tabulated as a separate subperiod. Panel A splits the harmonized sample by calendar year (39,566, 33,109, and 110,297 token-weeks; 861, 729, and 3,061 deaths); Panel B re-filters it at the stated floors (186,964, 163,295, and 120,880 token-weeks; 4,731, 1,790, and 485 deaths). Panel C is a token-level Cox proportional-hazards model on lifespans with covariates computed from each token’s first four weeks (6,043 tokens, 5,656 deaths).

	Illiquidity z	Volatility z
Panel A: subperiods		
2021	+6.22***	+6.72***
2022	+5.87***	+0.68
2023–2025	+6.42***	+3.44***
Panel B: volume floor		
\$50k floor	+9.36***	+1.43
\$100k floor	+5.92***	+1.42
\$250k floor	+3.47***	+1.65*
Panel C: token-level Cox (baseline covariates)		
Cox PH, baseline illiquidity	+14.48***	—

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 External validity: Solana

To assess whether the liquidity results are specific to Ethereum, we rebuild the panel on Solana, a chain with different venues, token standard, and fee environment. The Solana panel is reconstructed from the public Dune Analytics `dex_solana.trades` dataset, which decodes swaps across the major Solana venues, and spans April 2021 through December 2025. Because Solana’s fee environment produces a heavy tail of economically negligible micro-swaps, trades below \$1,000 are excluded before aggregation, and a token-week enters the panel at a weekly volume floor of \$250,000; token death is defined analogously to the Ethereum definition—a confirmed permanent drop below the panel’s volume floor. The panel contains 154,399 token-weeks from 104,306 distinct tokens, of which 103,182 permanently exit in sample; the hazard estimation sample, which requires a return and volatility history, contains 25,033 token-weeks, 1,155 deaths, and 2,230 tokens. The hazard estimates therefore describe the minority of tokens that outlive their first weeks. Table 10 reports the cross-chain comparison. Two features transfer: trading volume predicts death with the same sign on each chain, strongly in the univariate hazard, and realized volatility is not robustly predictive on either, so the liquidity-versus-volatility ranking holds across chains, with the liquidity signal carried by volume on Solana. The chains differ in which liquidity measure is most informative and in the nature of death itself. On Ethereum, Amihud illiquidity carries a large hazard and death is gradual, with the median token active for twenty weeks. On Solana the modal token falls below the floor within a single week (ninety percent of tokens are active for exactly one week), a rapid issuance-and-exit pattern in which raw volume carries most of the liquidity signal and Amihud is weaker.

The capacity adjustment carries over to Solana only partially. On the pump.fun cohort—tokens issued on Solana’s dominant token-launch platform, a cohort that accounts for most of recent Solana issuance and whose fixed supply makes market capitalization recoverable—capacity-

Table 8: Robustness of the illiquidity-versus-volatility ranking (lagged $t - 1$ baseline). Logit hazard z -statistics on Amihud illiquidity and realized volatility from a multivariate hazard controlling for volume, size, and age. Rows vary covariate timing, add calendar-year fixed effects, cluster by token, and restrict to single-spell tokens. Illiquidity is dominant throughout and volatility is not robustly signed. The lagged rows are estimated on the harmonized sample (186,964 token-weeks, 4,731 deaths); the contemporaneous row uses 201,674 token-weeks (5,064 deaths) and the single-spell row 52,354 (1,657).

Specification	Illiquidity z	Volatility z
Lagged baseline ($t - 1$)	+9.39***	+1.11
Contemporaneous covariates	+11.23***	+1.64
+ calendar-year fixed effects	+9.47***	+1.67*
Cluster by token	+14.81***	+1.14
Single-spell tokens only	+7.02***	-0.81

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Liquidity and death across AMM venues and numeraires (lagged $t - 1$ covariates). Logit hazard z -statistics (two-way clustered by token and week) on Amihud illiquidity and reconstructed pool depth, for tokens paired on Uniswap v2 and v3 with WETH and with major stablecoins; each model also controls for volume, volatility, and age. Illiquidity predicts death in every cell. The cumulative-net-flow depth measure does not provide a consistent signal across venue-numeraire cells; the significant positive depth coefficients in the stablecoin cells likely reflect limitations of the proxy rather than a clean depth channel. Cell samples (token-weeks; deaths): v2-WETH 158,262 (3,734); v2-stablecoin 72,752 (892); v3-WETH 144,333 (1,909); v3-stablecoin 101,662 (905).

AMM	Nueraire	Illiquidity z	Pool-depth z
Uniswap v2	WETH	+14.26***	+1.25
Uniswap v2	Stablecoin	+8.02***	+3.29***
Uniswap v3	WETH	+5.72***	+1.33
Uniswap v3	Stablecoin	+4.62***	+4.36***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

adjusted illiquidity improves out-of-sample prediction over Amihud illiquidity, and the bootstrap gain excludes zero (Appendix Table C.1). Two limitations qualify this evidence. The improvement comes from the market-value component of the adjustment; the participation refinement that helps on Ethereum adds nothing on the pump.fun cohort—it slightly reduces both out-of-sample areas—though breadth scaling alone still improves rank discrimination on the full Solana universe (Appendix Table C.1, Panel A). On Solana the taker count may simply be a poorer proxy for genuine breadth of participation. And because the pump.fun cohort has a fixed supply, its market capitalization moves with price, so part of the market-value adjustment reflects the price decline that precedes death rather than capacity alone—a caveat that applies, in attenuated form, to any price-times-supply capitalization measure. On Ethereum, reconstructed supply varies separately from price, diluting the mechanical price component. These results suggest that the market-value component of the capacity adjustment carries cross-chain signal, while the participation refinement is specific to settings in which the observed taker count tracks genuine breadth of participation.

Table 10: Cross-chain death hazard: Solana versus Ethereum (lagged $t - 1$ covariates). Logit hazard z -statistics; “Univariate” enters each predictor alone, and “Joint” enters all six jointly. Market capitalization is omitted (Solana token supply is not reconstructed). Trading volume and age predict death with the same sign on both chains; recent return is not a robust cross-chain predictor (the univariate signs differ and the joint coefficients are insignificant); volatility is not a robust predictor on either chain; Amihud illiquidity dominates on Ethereum but not Solana. Joint-model samples: Solana 25,033 token-weeks (1,155 deaths; the hazard sample of Section 5.3), Ethereum 206,162 (5,053; larger than the harmonized sample because size is not required here); univariate rows use each predictor’s maximal sample. Realized volatility on Solana is measured over a trailing eight-week window (at least three returns), reflecting Solana’s shorter token histories.

Predictor	Solana		Ethereum	
	Univariate	Joint	Univariate	Joint
Log trading volume	−12.10***	−2.56**	−39.32***	−11.36***
Token age	−31.81***	−9.81***	−6.52***	−10.83***
Recent return	+3.04***	+0.06	−2.80***	+1.35
Realized volatility	+3.01***	−0.96	+1.72*	+1.07
Amihud illiquidity	+2.25**	+2.46**	+31.06***	+9.34***
Log trade count	−8.92***	−3.70***	−19.44***	−0.22

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 Conclusion

In our survivorship-complete on-chain panels, token exit is more closely associated with the liquidity capacity of a token’s market than with the volatility of its price. Realized volatility, which prior survival studies flag as a death predictor, carries an independent hazard that is at most modest (hazard ratio 1.17 per standard deviation under the most favorable scaling) once liquidity and capacity are measured directly. It also discriminates the rare death event poorly across alternative scalings.

The paper’s central contribution is the capacity-adjusted illiquidity measure. It scales the price impact of trading by the square root of the market’s absorption capacity, and it is a parsimonious statistic, with no estimated parameters, that improves out-of-sample prediction beyond raw illiquidity—with the market-value component of the adjustment carrying most of that gain—and recovers nearly all of the fit of a richer multivariate hazard. Token fragility is two-dimensional in this predictive sense, combining price impact with the scale and breadth of the market that absorbs the trading; the measure summarizes both dimensions in a single statistic. The same statistic is related to the cross-section of returns: across portfolios, death probabilities and average returns align along a common risk–return frontier.

These findings suggest that empirical studies of the cryptocurrency cross-section should account for failed tokens and for liquidity capacity, and that volatility alone is an incomplete measure of token fragility. Token exit is accompanied by a joint deterioration in which trading activity and participation decline and price impact rises, which points to liquidity rather than price stability as the more informative margin for token longevity.

Several limitations qualify the interpretation. The harmonized-sample hazard regressions condition on a token surviving long enough to form a volatility estimate, so they describe tokens that establish a trading record rather than the very short-lived; the Kaplan–Meier evidence, which uses

only information from a token's first four weeks, shows the same illiquidity separation. Prices are DEX volume-weighted averages and therefore inherit measurement noise from decentralized-exchange microstructure. A token whose trading migrates entirely to centralized venues is classified as dead; because such migrations occur among larger, more liquid tokens, this misclassification biases against the estimated liquidity–death relation. Wash trading (Cong et al., 2023) may make a token appear more liquid than it is; because wash volume lowers measured illiquidity, wash activity among tokens close to exit would mask rather than create the estimated liquidity–death relation, and only wash activity concentrated among surviving tokens could inflate it. The estimates are predictive hazard associations rather than causal effects.

The setting is primarily Ethereum decentralized-exchange activity. The Solana evidence is consistent with the liquidity interpretation, most clearly through the volume and market-value components of the measure, while suggesting that participation-based refinements depend on market design. Extending the capacity adjustment to other settings, such as equity delisting, is a natural direction for future work.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the author used Claude (Anthropic) in order to assist with language editing and code development. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

References

- Acharya, V. V. and Pedersen, L. H. (2005). Asset pricing with liquidity risk. *Journal of Financial Economics*, 77(2):375–410.
- Almgren, R., Thum, C., Hauptmann, E., and Li, H. (2005). Direct estimation of equity market impact. *Risk*, 18(7):58–62.
- Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.
- Amihud, Y. and Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2):223–249.
- Ammann, M., Burdorf, T., Liebi, L., and Stöckl, S. (2022). Survivorship and delisting bias in cryptocurrency markets. *Working paper, University of St. Gallen*. SSRN 4287573.
- Asparouhova, E., Bessembinder, H., and Kalcheva, I. (2013). Noisy prices and inference regarding returns. *Journal of Finance*, 68(2):665–714.
- Barbon, A. and Rinaldo, A. (2026). On the quality of cryptocurrency markets: Centralized vs. decentralized exchanges. *Management Science*. Articles in Advance, DOI: 10.1287/mnsc.2024.07703.
- Bharath, S. T. and Shumway, T. (2008). Forecasting default with the Merton distance to default model. *Review of Financial Studies*, 21(3):1339–1369.
- Blume, M. E. and Stambaugh, R. F. (1983). Biases in computed returns: An application to the size effect. *Journal of Financial Economics*, 12(3):387–404.
- Brauneis, A., Mestel, R., Riordan, R., and Theissen, E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124:106041.
- Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238.
- Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, 29(2):238–249.
- Campbell, J. Y., Hilscher, J., and Szilagyi, J. (2008). In search of distress risk. *Journal of Finance*, 63(6):2899–2939.
- Chava, S. and Jarrow, R. A. (2004). Bankruptcy prediction with industry effects. *Review of Finance*, 8(4):537–569.
- Cong, L. W., Li, X., Tang, K., and Yang, Y. (2023). Crypto wash trading. *Management Science*, 69(11):6427–6454.
- Cong, L. W., Li, Y., and Wang, N. (2021). Tokenomics: Dynamic adoption and valuation. *Review of Financial Studies*, 34(3):1105–1155.
- Corbet, S. and Katsiampa, P. (2020). Asymmetric mean reversion of Bitcoin price returns. *International Review of Financial Analysis*, 71:101267.

- Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B*, 34(2):187–202.
- Daian, P., Goldfeder, S., Kell, T., Li, Y., Zhao, X., Bentov, I., Breidenbach, L., and Juels, A. (2020). Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 910–927.
- Dong, B., Jiang, L., Liu, J., and Zhu, Y. (2022). Liquidity in the cryptocurrency market and commonalities across anomalies. *International Review of Financial Analysis*, 81:102097.
- Fama, E. F. and MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636.
- Fantazzini, D. (2022). Crypto-coins and credit risk: Modelling and forecasting their probability of death. *Journal of Risk and Financial Management*, 15(7):304.
- Fieberg, C., Günther, S., Poddig, T., and Zarembo, A. (2024a). Non-standard errors in the cryptocurrency world. *International Review of Financial Analysis*, 92:103106.
- Fieberg, C., Liedtke, G., and Zarembo, A. (2024b). Cryptocurrency anomalies and economic constraints. *International Review of Financial Analysis*, 94:103218.
- Gatabazi, P., Kabera, G., Mba, J. C., Pindza, E., and Melesse, S. F. (2022). Cryptocurrencies and tokens lifetime analysis from 2009 to 2021. *Economics*, 10(3):60.
- Goyenko, R. Y., Holden, C. W., and Trzcinka, C. A. (2009). Do liquidity measures measure liquidity? *Journal of Financial Economics*, 92(2):153–181.
- Hasbrouck, J. (2009). Trading costs and returns for U.S. equities: Estimating effective costs from daily data. *Journal of Finance*, 64(3):1445–1477.
- Howell, S. T., Niessner, M., and Yermack, D. (2020). Initial coin offerings: Financing growth with cryptocurrency token sales. *Review of Financial Studies*, 33(9):3925–3974.
- Kim, A., Trimborn, S., and Härdle, W. K. (2021). VCRIX—a volatility index for crypto-currencies. *International Review of Financial Analysis*, 78:101915.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335.
- Lehar, A. and Parlour, C. A. (2025). Decentralized exchange: The Uniswap automated market maker. *Journal of Finance*, 80(1):321–374.
- Liu, Y., Tsyvinski, A., and Wu, X. (2022). Common risk factors in cryptocurrency. *Journal of Finance*, 77(2):1133–1177.
- Lunde, A., Timmermann, A., and Blake, D. (1999). The hazards of mutual fund underperformance: A Cox regression analysis. *Journal of Empirical Finance*, 6(2):121–152.
- Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319.
- Mazorra, B., Adan, V., and Daza, V. (2022). Do not rug on me: Leveraging machine learning techniques for automated scam detection. *Mathematics*, 10(6):949.

- Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.
- Petersen, M. A. (2009). Estimating standard errors in finance panel data sets: Comparing approaches. *Review of Financial Studies*, 22(1):435–480.
- Shumway, T. (1997). The delisting bias in CRSP data. *Journal of Finance*, 52(1):327–340.
- Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. *Journal of Business*, 74(1):101–124.
- Tóth, B., Lempérière, Y., Deremble, C., de Lataillade, J., Kockelkoren, J., and Bouchaud, J.-P. (2011). Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006.

A Mechanism details

This appendix documents the mechanism evidence summarized in Section 3.5: the decomposition of the Amihud ratio into its components, the information in reconstructed liquidity provision, and whether sandwich extraction rather than illiquidity carries the hazard.

The Amihud measure combines two moving parts—the size of a token’s price change and its dollar volume—so we test which one the hazard loads on. Because Amihud illiquidity is absolute return divided by dollar volume, its logarithm equals log absolute return minus log volume. Table A.1 enters the ratio and its two components. On its own, low dollar volume is a strong predictor of death—volume enters with a large negative z -statistic—while a larger absolute return is a weaker, positive predictor. Entered jointly, neither the ratio’s coefficient nor the components’ coefficients are separately interpretable—a mechanical consequence of including a ratio alongside its own numerator and denominator—so the univariate rows carry the information. We therefore read the Amihud ratio as a parsimonious price-impact summary rather than as separately identified volume and return channels.

Trading is only one side of an automated market maker; the other is liquidity provision, which we examine next. In a constant-product automated market maker, liquidity providers deposit a pair of tokens into a shared pool that traders swap against—adding liquidity through Mint events and removing it through Burn events—and a token’s price impact falls as its pool grows deeper. We reconstruct each token’s pool depth on Uniswap v2, a major venue for long-tail tokens, as the cumulative net liquidity supplied on the WETH (wrapped-ether) side, the asset these tokens are most often paired against. Table A.2 adds the provision measures to the one-week-lagged hazard. Illiquidity retains its large hazard, but the lagged provision proxies do not isolate a clean withdrawal channel: pool depth is insignificant, and lagged withdrawal enters with a negative sign, likely because Burn events can occur only in pools that still hold liquidity, so they proxy for pool size rather than distress. These proxies therefore identify no distress signal beyond illiquidity; the deterioration that Figure 2 documents around exit is in trading conditions—price impact, volume, and participation—not in measured provision. We do not reconstruct swap-driven reserve changes or impermanent loss (the shortfall of pool value relative to holding the deposited assets), so pool depth is a net-provision proxy; in this appendix its coverage is limited to WETH pairs, while the venue analysis of Section 5 applies the same cumulative construction on each cell’s numeraire side.

A final exercise tests whether a trading friction specific to decentralized exchanges, rather than illiquidity, carries the death hazard. Table A.3 measures each token-week’s MEV exposure as the share of its DEX volume executed as sandwich-victim trades, where a sandwich attack brackets a victim’s swap with an attacker’s own trades. The association with death is negative after controlling for illiquidity and the other characteristics, consistent with sandwich activity occurring in markets with sustained, predictable order flow. In this setting, the friction most closely associated with token death is illiquidity itself, not measured sandwich extraction.

Table A.1: Amihud decomposition (lagged $t - 1$ covariates). Amihud illiquidity is absolute return over dollar volume; this table separates the price-impact numerator and the volume denominator in the death hazard. Logit hazard z -statistics, two-way clustered by token and week. On their own, low volume and higher absolute return each predict death; when the ratio and its components are entered jointly, the coefficients are not separately interpretable (a ratio beside its own numerator and denominator), and the univariate rows carry the information. The joint row is estimated on the lagged-illiquidity panel of Table 1 (236,151 token-weeks, 6,049 deaths); univariate rows use each covariate's maximal sample (up to 255,359 token-weeks for volume).

Specification	Illiquidity z	Volume z	Return z
Volume only	—	−39.32***	—
Return only	—	—	+9.12***
Illiquidity only	+31.06***	—	—
Illiquidity + volume + return	+5.29***	−2.14**	−2.72***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.2: Liquidity provision and the death hazard (lagged $t - 1$ covariates). Logit hazard z -statistics (two-way clustered by token and week). Pool depth is cumulative net WETH-side liquidity; liquidity-provider (LP) withdrawal is log liquidity removed via Burn events. Illiquidity retains a large hazard. The lagged provision proxies do not identify a clean withdrawal channel: pool depth is insignificant and withdrawal enters with a negative sign, consistent with Burn events proxying for pool size rather than distress. Sample: the lagged-illiquidity panel of Table 1 (236,151 token-weeks, 6,049 deaths).

Predictor	Hazard z
Pool depth	+1.19
LP withdrawal (Burn)	−3.19***
Amihud illiquidity	+13.23***
Log volume	−26.39***
Token age	−9.64***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: MEV exposure and token death (lagged $t - 1$ covariates). Sandwich-MEV exposure is the share of a token's weekly DEX volume executed as sandwich-victim trades. Logit hazard z -statistics (two-way clustered by token and week); the extended model also controls for volume, volatility, size, and age. MEV exposure is negatively associated with death after controlling for illiquidity and other characteristics, consistent with sandwich activity occurring in markets with sustained order flow. Samples: 255,359 token-weeks for the univariate row, 236,151 with illiquidity, and the harmonized sample (186,964 token-weeks, 4,731 deaths) for the extended model.

Specification	MEV exposure z	Illiquidity z
MEV exposure only	+0.75	—
MEV exposure + illiquidity	−0.22	+30.92***
Extended model	−3.23***	+9.59***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Additional robustness

Table B.1 repeats the out-of-sample horse race on the broader sample that drops the volatility-availability requirement of the main analysis; the capacity-adjusted measure again improves on Amihud ($\Delta\text{AUC} +0.040$, $\Delta\text{PR-AUC} +0.031$, both excluding zero), so the result is not driven by the harmonized-sample restriction. Table B.2 reports, over a grid of exponents (a, b) and on the same broader sample, the out-of-sample performance of the generalized statistic $\text{Amihud}/(M^a N^b)$. Performance is similar across all cells with a positive market-value exponent, so the exponents are only weakly identified; the square-root form adopted in the text lies within a few thousandths of the best cell on both metrics and retains the market-impact motivation.

Table B.3 reports additional robustness checks for the capacity-adjusted measure, including subperiod stability, two-way token-and-week clustered inference, and a token-level Cox model based only on first-four-week values.

Table B.4 repeats the Amihud–CAI comparison with rolling training origins, and Table B.5 lengthens the covariate lag from one week to twelve; the capacity-adjusted measure improves on raw Amihud in every fold, and discrimination remains above chance at a twelve-week lag.

Table B.1: Out-of-sample horse race on the broader sample that drops the volatility-availability requirement of Table 5, retaining token-weeks with Amihud illiquidity, capitalization, and breadth. Train before 2024, test 2024–2025; differences versus Amihud with 95% week-block bootstrap intervals. Test sample: 90,431 token-weeks, 2,888 deaths.

Predictor	AUC	PR-AUC	$\Delta\text{PR-AUC}$ vs Amihud	ΔAUC vs Amihud
Amihud $ r /V$	0.786	0.109	—	—
Size $(-\log M)$	0.774	0.113	+0.004 $[-0.005, +0.013]$	-0.012 $[-0.024, +0.000]$
Breadth $(-\log N)$	0.674	0.054	-0.056 $[-0.065, -0.048]$	-0.113 $[-0.130, -0.095]$
Amihud/ $\sqrt{M}$	0.814	0.136	+0.027 $[+0.021, +0.032]$	+0.028 $[+0.021, +0.034]$
$\Lambda^{\text{adj}} = \text{Amihud}/\sqrt{MN}$	0.826	0.140	+0.031 $[+0.026, +0.036]$	+0.040 $[+0.033, +0.046]$
Free multivariate (+age)	0.831	0.132	+0.023 $[+0.018, +0.029]$	+0.045 $[+0.035, +0.053]$

Week-block bootstrap (150 resamples); differences whose intervals exclude zero are significant at 5%.

Differences are computed before rounding.

Table B.2: Exponent robustness for capacity-adjusted illiquidity. Out-of-sample AUC and PR-AUC (training before 2024, testing 2024–2025), on the broader sample of Table B.1, for the generalized statistic $\text{Amihud}/(M^a N^b)$, reported as AUC / PR-AUC. The cell $(a, b) = (0, 0)$ is raw Amihud illiquidity. Performance is similar across all cells with a positive market-value exponent; the square-root form $(a, b) = (\frac{1}{2}, \frac{1}{2})$ adopted in the text lies within a few thousandths of the best cell on both metrics.

	$b = 0$	$b = \frac{1}{2}$	$b = 1$
$a = 0$	0.786 / 0.109	0.796 / 0.109	0.795 / 0.101
$a = \frac{1}{2}$	0.814 / 0.136	0.826 / 0.140	0.830 / 0.134
$a = 1$	0.812 / 0.135	0.824 / 0.142	0.830 / 0.142

Table B.3: Robustness of capacity-adjusted illiquidity. Discrete-time logit hazard of token death on the standardized, one-week-lagged $\Lambda^{\text{adj}} = \text{Amihud}/\sqrt{MN}$, harmonized sample, with no additional controls or fixed effects. Panel A re-estimates within subperiods (week-clustered). Panel B compares week clustering, two-way token-and-week clustering, token clustering, and HC1 standard errors on the full sample; the hazard ratio is identical across rows because clustering changes only the standard errors, and two-way clustering yields the most conservative inference. Panel C is a look-ahead-free token-level Cox on each token's mean CAI over its first four weeks.

	Hazard ratio / SD	z
Panel A: subperiods (week-clustered logit)		
2020–2021	6.66	+32.7
2022–2023	6.67	+40.1
2024–2025	5.96	+16.7
Panel B: standard-error clustering (full sample)		
Week-clustered	6.41	+33.3
Token-and-week (two-way)	6.41	+32.6
Token-clustered	6.41	+68.0
Robust (HC1)	6.41	+75.1
Panel C: token-level Cox PH (first-four-week baseline, no look-ahead)		
Baseline Λ^{adj}	1.63	+34.5

Hazard ratio per SD; z from the stated standard errors.

Table B.4: Rolling-origin out-of-sample folds. Univariate discrete-time logits on lagged Amihud illiquidity and on lagged capacity-adjusted illiquidity (CAI), trained through year y and evaluated on $y+1$, using the broader sample of Table B.1. Out-of-sample AUC and PR-AUC in each test year and the PR-AUC gain of CAI over Amihud, computed before rounding. The ranking is unchanged in every fold.

Test year	Amihud		CAI		$\Delta\text{PR-AUC}$	Test deaths
	AUC	PR-AUC	AUC	PR-AUC		
2023	0.789	0.082	0.846	0.124	+0.041	787
2024	0.803	0.101	0.850	0.133	+0.032	1,570
2025	0.769	0.128	0.804	0.162	+0.035	1,318

Table B.5: Discrimination as the covariate lag lengthens. Univariate discrete-time logit on capacity-adjusted illiquidity lagged h weeks, trained before 2024 and evaluated on 2024–2025 using the broader sample of Table B.1; the sample shrinks as h grows because longer lags require longer histories. Discrimination remains above chance at horizons of two to three months, indicating the measure contains a persistent component beyond exit-week mechanics.

Lag h (weeks)	AUC	PR-AUC	Test token-weeks	Test deaths
1	0.826	0.140	90,431	2,888
4	0.802	0.080	76,754	1,698
8	0.767	0.054	65,599	1,052
12	0.736	0.034	57,913	709

Table B.6 reports two further checks. First, replacing the raw taker count with a concentration-based effective number of traders, the reciprocal of the within-week Herfindahl index, leaves out-of-sample performance essentially unchanged, suggesting that the result is not solely due to unusually high participation counts. Second, when tokens are ranked on first-four-week values alone, all three illiquidity variants separate eventual deaths about equally well (AUC 0.877–0.880): the early-life ranking is already carried by price impact, and the capacity adjustment contributes through week-ahead timing rather than early-life discrimination. Separately, a confirmed-death definition that calls a token dead only after at least eight weeks of inactivity (4,441 confirmed deaths) preserves the ranking (untabulated), with illiquidity strongly predictive ($z = +13.8$) and volatility insignificant ($z = +1.0$).

Table B.6: Robustness of the capacity adjustment to breadth measurement and early life. Panel A replaces the raw taker count with a concentration-based effective number of traders ($N = 1/\text{HHI}$) in the full out-of-sample test (train before 2024, test 2024–2025); Panel A is recomputed on its own concentration-merged panel, so its Amihud benchmark differs from Table B.1’s in the third decimal. Panel B ranks tokens by first-four-week baseline values and reports token-level AUC for eventual death.

Specification	AUC	PR-AUC
Panel A: effective breadth ($N = 1/\text{HHI}$)		
Amihud $ r /V$	0.787	0.107
Λ^{adj} (raw breadth)	0.826	0.140
Λ^{adj} (effective breadth)	0.816	0.135
Panel B: first-four-week baseline, token-level (AUC only)		
Amihud $ r /V$	0.877	—
Amihud/ $\sqrt{M}$	0.880	—
Λ^{adj} (CAI)	0.877	—

Table B.7 re-estimates the breadth-augmented hazard with realized volatility entered in logs and winsorized: volatility then carries a modest but significant independent hazard ($z = +6.0$ and $+4.5$), while illiquidity remains the stronger predictor ($z = +8.6$ and $+8.7$), and volatility’s univariate out-of-sample AUC is unchanged at 0.643, as the AUC depends only on ranks.

Table B.7: Realized volatility under alternative transformations. Baseline enters volatility in levels; this table re-enters it in logs and winsorized (1%/99%), harmonized sample. Columns: volatility hazard ratio per SD and two-way clustered z in the breadth-augmented multivariate hazard, volatility z alone, and out-of-sample univariate AUC. Under the log and winsorized transformations, volatility carries a modest but significant independent hazard, while illiquidity remains the stronger predictor throughout (illiquidity z for the same model in the note) and volatility’s AUC is unchanged, as the AUC depends only on ranks.

Volatility treatment	HR/SD (joint)	z (joint)	z (alone)	OOS AUC
Levels (baseline)	1.00	+1.2	+1.6	0.643
Logarithm	1.17	+6.0	+11.9	0.643
Winsorized (1%/99%)	1.08	+4.5	+9.7	0.643
Illiquidity z in the same model: +9.1 (levels), +8.6 (logarithm), +8.7 (winsorized).				

C Solana external validity

Table C.1 documents the Solana capacity-adjustment result of Section 5.3: on the pump.fun cohort the capacity adjustment improves out-of-sample prediction over Amihud illiquidity with a bootstrap gain excluding zero; the market-value component carries the improvement, and the participation refinement contributes nothing on that cohort.

Table C.1: Solana validation of the capacity adjustment. Out-of-sample discrimination of token death on Solana, training before 2025 and testing on 2025, with predictors lagged one week as in the Ethereum test. Panel A uses the full Solana universe and isolates the participation term, comparing Amihud illiquidity with its breadth-scaled version $\text{Amihud}/\sqrt{N}$ (17,638 test token-weeks, 1,688 deaths). Panel B uses the pump.fun cohort, whose fixed supply of one billion tokens makes market capitalization recoverable without a supply query, and compares the full set of adjustments (5,716 test token-weeks, 623 deaths). The market-value adjustment carries the cross-chain signal on the pump.fun cohort, where adding the participation term slightly reduces both areas; on the full universe, breadth scaling alone improves rank discrimination. Bootstrap 95% intervals (week blocks): in Panel A, $\text{Amihud}/\sqrt{N}$ minus Amihud has $\Delta\text{AUC} +0.017 [+0.003, +0.029]$ and $\Delta\text{PR-AUC} -0.006 [-0.020, +0.009]$; in Panel B, $\text{Amihud}/\sqrt{MN}$ minus Amihud has $\Delta\text{AUC} +0.032 [+0.027, +0.037]$ and $\Delta\text{PR-AUC} +0.030 [+0.016, +0.041]$, both excluding zero. Differences are computed before rounding.

Predictor	AUC	PR-AUC
Panel A: full Solana universe (participation term)		
Amihud $ r /V$	0.597	0.191
$\text{Amihud}/\sqrt{N}$	0.615	0.185
Panel B: pump.fun cohort (market capitalization recoverable)		
Amihud $ r /V$	0.739	0.268
$\text{Amihud}/\sqrt{N}$	0.720	0.242
$\text{Amihud}/\sqrt{M}$	0.786	0.322
$\text{Amihud}/\sqrt{MN}$	0.771	0.299

D Pricing robustness

Table D.1 reports the robustness of the capacity-adjusted illiquidity return premium: its alpha against market, size, and momentum factors, its stability across subperiods, its Solana counterpart based on $\text{Amihud}/\sqrt{N}$, and its persistence among larger tokens. Table D.2 re-estimates the premium, assigning an explicit -100 percent forward return to every death week.

Table D.1: Robustness of the capacity-adjusted illiquidity premium. Panel A regresses the weekly decile long-short (highest minus lowest CAI) on an equal-weighted market return and on size and momentum long-short factors; the intercept is the factor-adjusted premium. The raw long-short here is measured on the weeks with all three factors available, so its level differs slightly from Panel B of Table 6. Panel B reports the univariate Fama–MacBeth slope per SD by subperiod. Panel C reports the same slope on Solana using Amihud/ $\sqrt{N}$, because market capitalization is unavailable for the full Solana universe. Panel D restricts the sort and the Fama–MacBeth regression to larger tokens. Newey–West t (4 lags); returns winsorized 1/99%.

	Estimate	NW t
Panel A: factor-adjusted premium (long-short, % per week)		
Raw long-short	+8.6	+7.94
Alpha vs market, size, momentum	+7.1	+9.11
Panel B: Fama–MacBeth slope per SD, by subperiod		
2020–2021	+0.024	+4.93
2022–2023	+0.019	+8.89
2024–2025	+0.017	+3.70
Panel C: Solana cross-chain		
Solana, Amihud/ $\sqrt{N}$	+0.033	+5.86
Panel D: larger tokens only		
Above-median cap., quintile long-short (% per week)	+5.6	+8.59
Above-median cap., FM slope per SD	+0.019	+8.22
Top size tercile, quintile long-short (% per week)	+7.3	+7.43
Top size tercile, FM slope per SD	+0.023	+7.31

Table D.2: Return premium under an explicit exit return. A forward return of -100% is assigned at the death week, the token’s final active week under the death definition of Section 2 (a dead token has no subsequent in-panel price); Fama–MacBeth univariate slope per SD and the decile D10–D1 spread (% per week), Newey–West t (4 lags). The premium is roughly halved but remains strongly significant, so it is not driven solely by the missing returns of dying tokens.

	Estimate	NW t
FM slope per SD (exit = -100%)	+0.012	+4.49
D10 – D1 spread (% per week)	+4.89	+3.73