forecasting

MDPI

Article

Delayed Marks, Funding Memory, and Forecasting Liquidation-Tail Risk in Crypto Perpetual Futures

Edson Pindza * and Hopolang Phillip Mashele

Department of Decision Sciences, School of Economic and Financial Sciences, College of Economic and Management Sciences, University of South Africa (UNISA), Pretoria 0002, South Africa
* Correspondence: edsonpindza@gmail.com

Highlights

What are the main findings?

- The model produces a horizon-specific first-passage breach probability and a conditional catch-up-loss distribution for an isolated-margin position.
- In rolling forecasts, the structural model has lower Brier scores than historical simulation at 10× across 24-, 72- and 168-hour horizons, but not uniformly at high weekly leverage.

What are the implications of the main findings?

- Probability forecasts should jointly condition on mark delay, maintenance margin, funding state and executable depth.
- The benchmark results point to conditionally heteroskedastic or regime-sensitive price densities as the next extension for stressed high-leverage forecasts.

Abstract

Crypto perpetual futures embed liquidation risk in one chain: leverage and funding move the margin boundary, the mark determines when a crossing is observed, and executable depth determines the concession paid after detection. The primary forecasting question is how to quantify both the probability of an isolated-margin boundary breach and the loss hidden by delayed or smoothed detection. This paper develops a first-passage density-forecasting framework in which the executable price is observed through a delayed or smoothed mark, funding is a persistent collateral drain, and liquidation occurs when isolated-margin surplus reaches its maintenance boundary. The output is a joint predictive distribution: a horizon-specific probability of a boundary breach and, conditional on detection, a distribution of catch-up loss. Under a fixed delay, expected overshoot is $\sigma\sqrt{\delta/2\pi}$, and bad-debt probability is a Gaussian tail governed by the latency-to-margin ratio $\sigma\sqrt{\delta}/m$. Its leverage independence is exact only in the constant-maintenance, fixed-non-price-drain benchmark. For time-weighted-average marks, fixed-time variance reduction does not imply liquidation-time safety. A stationary-downcrossing approximation size-biases the stale error in the adverse direction, producing a mean overshoot about 1.6 times the same-window fixed-delay value. A rolling comparison with historical simulation scores model-consistent margin-breach forecasts from public price and funding paths. The structural forecast has lower Brier scores at 10× over 24-, 72- and 168-hour horizons and across the 24-hour grid, but historical simulation performs better for one-week forecasts at 20× and 50×. The evidence supports a conditional risk-forecasting use of the framework, while not establishing uniform forecast dominance.

Keywords: crypto perpetual futures; liquidation risk forecasting; predictive distributions; oracle latency; funding rates; order-book depth; first-passage; market microstructure

Check for updates

Academic Editors: Dimitrios D. Thomakos and Periklis Gogas

Received: 15 July 2026
Revised: 19 August 2026
Accepted: 28 August 2026
Published: 30 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

MSC: 60G22; 60H30; 91G20; 91G60

JEL Classification: C58; G12; G13; G15; G23

1. Introduction

1.1. Motivation and Research Problem

Perpetual futures are a major venue for leveraged crypto exposure without fixed maturity. The contract tracks a mark price, a periodic funding transfer keeps the derivative close to spot, and a margin engine liquidates accounts whose collateral falls below a maintenance requirement. For risk management, these components cannot be treated as separate controls. Funding rates are feedback variables that respond to and influence positioning pressure [1,2]. The mark price is an exchange-defined observation operator that may smooth, delay or discretise information relative to the executable price [3]. The order book determines how shocks convert into slippage and collateral depletion [4]. Liquidation tails depend on the joint behaviour of funding, mark construction and executable depth, rather than on any one component in isolation.

The interaction is modelled as a first-passage problem. Leverage and funding determine when executable margin surplus reaches its maintenance boundary; mark construction determines when that crossing is detected and how much executable price risk remains at liquidation; depth determines the additional execution loss. A trader can appear solvent under a delayed mark while the executable price has already moved adversely; when the mark catches up, liquidation triggers late, at a price that may have overshot the boundary. Persistent funding drains collateral in the direction of crowded positioning and inflates the variance of cumulative drain. Convex depth converts forced liquidation flow into more-than-proportional impact in the cascade regime. The tail event is therefore a boundary-crossing problem for a risk engine, not only a volatility forecast.

1.2. Objective and Primary Research Question

The primary research question is: *How can the probability and conditional severity of an isolated-margin liquidation be forecast when funding changes the account boundary, the mark observes executable prices with delay or smoothing, and liquidation must be executed against finite depth?* Existing approaches usually study one part of this chain—return tails, perpetual-futures pricing, oracle manipulation, or liquidation execution—while taking the other parts as external. The objective here is to construct one first-passage framework that maps those mechanisms into a joint event–severity predictive distribution and to evaluate each component with the public data that can identify it.

This formulation imposes a single architecture on the paper. Price and funding paths determine whether surplus crosses the maintenance boundary; the mark operator determines when that crossing is observed and the hidden price displacement at detection; the order book determines the execution concession after detection. The theory, simulations, and empirical measurements below follow that sequence. Stablecoin and option observations describe the surrounding market state but are not additional causal mechanisms in the central calibration.

1.3. Forecast Object and Contributions

The forecast object is consequently a distribution rather than a point estimate. At a forecast origin, the model assigns a probability to an isolated-margin boundary breach over a stated horizon and, conditional on a detected breach, assigns a continuous distribution to

the catch-up overshoot and bad-debt state. The binary event distribution answers whether a particular margin configuration is likely to fail; the conditional severity distribution answers how much executable loss can be hidden by the mark. This distinction permits a direct forecast evaluation of the first component on later public price and funding paths, while retaining the latency results as a forecast of conditional loss severity.

In plain terms, the framework forecasts two linked outcomes. First, it forecasts the chance that a stated leverage–margin configuration will breach within 24, 72, or 168 h. Second, if the mark detects a breach, it forecasts the extra loss hidden during the delay or smoothing window. The forecasting contribution is this explicit event–severity distribution, not merely a fitted volatility, an in-sample liquidation count, or a descriptive measure of funding persistence.

The main contributions are fourfold:

1. A unified first-passage forecast in which leverage, funding, price observation, maintenance margin, and execution depth have distinct and traceable roles;
2. Stopping-time formulas for the conditional catch-up loss of fixed-delay and TWAP marks, including the distinction between fixed-time smoothing variance and crossing-selected error at liquidation;
3. A validation design that combines analytical checks, path-level mechanism simulation, public-data calibration, and rolling proper-score comparison with historical simulation;
4. Explicit claim boundaries that separate model-consistent hypothetical-account forecasts from unobserved venue-level customer liquidations.

Table 1 makes the cumulative research design explicit. Each empirical object is attached to a defined theoretical output, so the evidence is not pooled into a generic fragility index.

Table 1. Unified theory-to-evidence map. The three rows are components of one liquidation forecast rather than independent empirical exercises.

Forecast Component	Principal Inputs	Theoretical Output	Validation or Evaluation
Boundary-breach event	Price path, leverage, maintenance margin, funding path	$\pi_{t,h,L} = \mathbb{P}(Y_{t,h,L} = 1 \mid \mathcal{I}_t)$	Calibrated first passage; rolling Brier and log scores against historical simulation
Conditional delayed-mark severity	Volatility, feed delay, maintenance margin	Fixed-delay overshoot and bad-debt distribution	One-minute stale errors and path-level delayed-mark simulation
Conditional smoothing and execution severity	TWAP window, volatility, crossing selection, liquidation size and depth	Crossing-conditioned TWAP loss plus average execution concession	TWAP mechanism simulation and public depth descriptors

1.4. Data, Validation, and Principal Findings

The first analytical result derives the cost of oracle delay at the liquidation time rather than at an arbitrary fixed time. Under a mark delayed by δ, the executable liquidation price overshoots the boundary with mean $\sigma\sqrt{\delta/2\pi}$, and the probability of a bad-debt liquidation is $\Phi(\log(1-m)/(\sigma\sqrt{\delta})) \approx \Phi(-m/(\sigma\sqrt{\delta}))$ (Theorem 1). The dimensionless driver is the latency–margin ratio $\sigma\sqrt{\delta}/m$. The leverage cancellation is exact only for constant maintenance fraction and fixed non-price drains during the delay; tiered margin, funding, fees, or execution can restore leverage and size dependence. For a time-weighted-average-price (TWAP) mark of window δ, the unconditional stale error carries only one third of the fixed-delay variance. Liquidation instead samples the error at a downcrossing of the

boundary by the smoothed mark. A stationary-crossing approximation size-biases the error in the adverse direction, giving a TWAP catch-up overshoot about 1.6 times the same-window fixed-delay value and a numerical latency–margin crossover near 0.89 (Theorem 2). Path-level simulation confirms the fixed-delay result and measures the error induced by applying the stationary-crossing approximation to a discrete first-passage mechanism.

A second result separates breach risk from loss severity. Leverage moves the first-passage boundary, funding dependence changes the variance of cumulative collateral drain, and an execution-cost process raises breach probability whenever it is pointwise larger (Proposition 3). Under fractional driving without mean reversion, cumulative funding drain scales as T^{2H+2} (Proposition 2), a prediction confirmed by direct fractional-Brownian simulation with fitted log–log exponents within 1.1% of the theoretical values (Figure 1). This power-law exponent is distinct from the level multiplier of the tractable AR(1) drain used in the breach experiment, which inflates one-week cumulative-drain variance by approximately $(1+\widehat{\phi})/(1-\widehat{\phi}) \approx 6$ without changing the long-horizon exponent of the AR(1) sum.

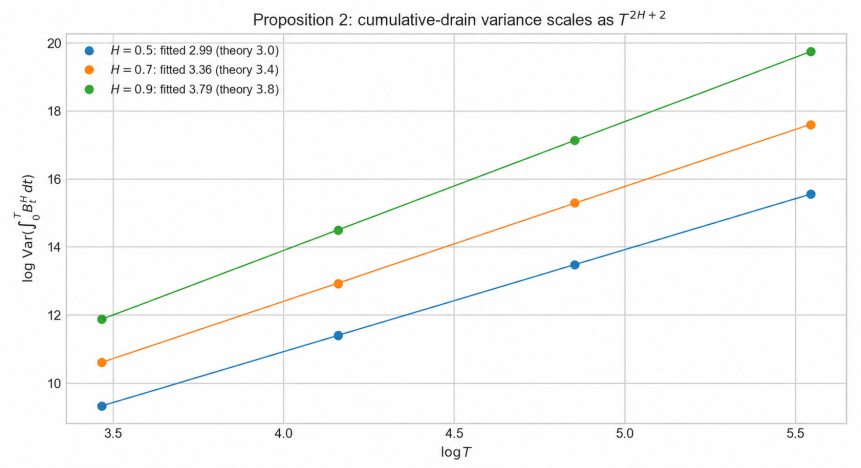


Figure 1. Validation of Proposition 2: $\log \mathrm{Var}(\int_0^T B_t^H \, \mathrm{d}t)$ against $\log T$ for $H = 0.5, 0.7, 0.9$, from exact fractional-Gaussian-noise simulation. Fitted slopes $(2.99, 3.36, 3.79)$ match the theoretical $2H+2$ $(3.0, 3.4, 3.8)$ to within 1.1%.

Empirically, the claims are restricted to objects measurable from public data. Funding persistence is estimated on the native 8-hourly funding series rather than on a forward-filled hourly grid, which would mechanically repeat observations and manufacture persistence. On the native series, R/S Hurst exponents of 0.82–0.94 lie 3.1–4.9 surrogate standard deviations above an i.i.d. shuffle baseline of 0.62. This is evidence of positive serial dependence in the sample, not of a specific fractional law.

The first-passage experiment uses the isolated-margin boundary $d_L = (1/L - m)/(1-m)$ and horizon-matched volatility. With that boundary, one-week breach probabilities vary across leverage, from 0.07–0.23 at $10\times$ to 0.39–0.56 at $20\times$, and funding memory is second-order at calm funding levels but material in crowded regimes. A mis-specified boundary would saturate these probabilities near unity and erase the comparative statics, so the margin algebra is part of the model rather than an implementation detail.

The empirical study uses public data accessed on 14 June 2026 UTC. The principal calibration sample contains Binance USD-M futures hourly and one-minute prices, funding histories and order-book depth for BTCUSDT, ETHUSDT and SOLUSDT: 1000 hourly and 1000 one-minute observations per asset, 200 native eight-hourly funding records per asset, and 15 depth-feature rows. The hourly panel runs from 3 May to 14 June 2026. A series of 379 aggregate stablecoin observations and a cross-section of 1684 public BTC and ETH

option summaries describe the wider liquidity and forward-looking volatility environment. These contextual series do not enter the liquidation-probability calibration and are not treated as user-level collateral or liquidation evidence.

1.5. Paper Organisation

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 defines the margin, mark-price, funding and depth primitives. Section 4 derives the first-passage results. Section 5 describes the data and empirical design, Section 6 reports the calibration and validation evidence, and Sections 7–9 discuss limitations and implications before Section 10 concludes.

2. Literature Review

The relevant literature concerns structural first passage, market impact, persistent financial processes, perpetual-futures funding, liquidation and oracle design.

2.1. First Passage and Margin Boundaries

The margin engine is a first-passage device. In structural credit models, default occurs when a state crosses a boundary: Merton [5] prices default at maturity, while Black and Cox [6] introduce the first-passage barrier that triggers before maturity, and Leland [7] and Longstaff and Schwartz [8] develop endogenous and exogenous-barrier variants. Crypto liquidation is operationally faster and more rule-based than corporate default, but the mathematics is the same: a stochastic surplus reaches a boundary before intervention. We place the boundary inside a perpetual-futures margin engine with funding payments and, critically, a *delayed* observation of the state.

2.2. Microstructure, Impact and Depth

Limit-order-book markets do not convert quantity into price at a constant rate; marginal cost rises as liquidity is consumed, especially under stress [9–14]. In crypto specifically, Makarov and Schoar [15] document segmentation and arbitrage limits, Brauneis et al. [16] benchmark liquidity measures, and Donier and Bonart [4] confirm a square-root impact law on Bitcoin. We use a convex depth function so that a given mark displacement carries higher loss when depth is thin, and we calibrate the cross-asset depth gap directly from public Binance snapshots.

2.3. Long Memory and Rough Dynamics

Long memory was formalised by Mandelbrot and Van Ness [17]; Cont [18] surveys long-range dependence in financial series, and Gatheral et al. [19] show volatility is rough, with Takaishi [20] extending roughness to Bitcoin. The natural object for a persistent funding state is the fractional Ornstein–Uhlenbeck process [21], whose drift can be estimated consistently [22]. The empirical analysis tests for positive serial dependence but does not identify a universal fractional law for funding.

2.4. Perpetual Futures, Collateral Liquidity and DeFi Liquidations

The funding transfer is the perpetual analogue of cost-of-carry in commodity forward pricing [23]. He et al. [2] derive no-arbitrage perpetual prices and document slowly decaying funding-basis deviations; Ackerer et al. [1] provide a pricing theory in which the perpetual equals spot sampled at a funding-determined random time. Liquidation cascades connect to the margin- and loss-spiral mechanism of Brunnermeier and Pedersen [24]. In decentralised markets, Qin et al. [25] study liquidations empirically, Werner et al. [26] systematise liquidation and oracle dependence, Lehar and Parlour [27] measure the price impact of on-chain liquidations, and Klages-Mundt and Minca [28] model deleveraging spi-

rals as a stochastic stability problem. Stablecoins are an important settlement and collateral medium in digital-asset markets, but aggregate supply is not equivalent to collateral posted by perpetual-futures traders. The literature documents both their role in decentralised finance [29,30] and their susceptibility to leverage, run and flight-to-safety dynamics [31,32]. We therefore use stablecoin supply only to describe market-wide liquidity conditions, not to identify the margin state of individual accounts. The broader crypto risk environment is also distinct from conventional asset classes [33], and transaction activity may be related to stablecoin flows [34]; neither relationship is imposed as a causal channel in the model.

2.5. Oracle Risk

Price feeds may be manipulated or stale. Qin et al. [35] document flash-loan oracle manipulation, and Mackinga et al. [3] show that time-weighted oracles can be cheaper to manipulate than previously assumed. Those studies concern adversarial price formation. The present model concerns the stopping-time consequences of latency even when the observed price is not manipulated.

2.6. Probabilistic Risk Forecasting

A risk forecast is properly evaluated as a predictive distribution, not only through a fitted volatility or a realised tail event. Proper scoring rules provide incentives for honest probabilistic forecasts [36], and density-forecast evaluation has a long financial-risk-management application [37]. In the present setting, the first-passage event produces a Bernoulli predictive distribution, while the fixed-delay and TWAP results provide conditional continuous distributions for catch-up severity. This framing differs from a conventional return or VaR forecast: the forecasted state is the future crossing of an isolated-margin boundary after funding payments, and the observation rule determines the loss distribution once the mark detects that crossing.

2.7. Closest Related Work

Three recent strands sit nearest to the present model. On *crypto liquidation and margin*, Cheng et al. [38] characterise Bitcoin-futures liquidation risk and optimal margins through generalised-extreme-value tail theory on observed returns, computing margin-call probabilities across leverage and horizons. They share the margin/leverage/breach object but model neither a delayed mark nor funding memory; the additional object here is the delayed-observation operator inside the margin stopping rule. On *perpetual-futures mechanics*, He et al. [2] and Ackerer et al. [1] price the contract and characterise funding, and Kim and Park [39] design funding rates via path-dependent BSDEs; these papers do not treat latency or the liquidation stopping problem under a coarsened filtration. On *liquidation execution*, Donnelly et al. [40] solve the optimal *active* liquidation of a perpetual position by stochastic control, which is complementary to the passive margin first passage studied here. Finally, the mathematical primitive of a first passage observed through delay appears in detection theory [41], where the problem is signal estimation rather than margin loss.

2.8. Critical Comparison of Forecast Targets

The differences are methodological, not only topical. Extreme-value margin models estimate a tail from realised returns and then map that tail to a margin rule; they can be well suited to observed return extremes but do not identify the loss hidden between an executable crossing and a delayed mark. Perpetual-pricing models discipline the no-arbitrage relation between spot, the contract, and funding, but a pricing identity does not specify an account's stopping boundary. Oracle-manipulation studies quantify adversarial distortion, whereas the present latency question remains even when the feed is honest. Optimal-execution models begin with a decision to liquidate, while the present event is

the passive trigger that precedes that decision. These distinct forecast targets explain why none of the existing strands can be substituted unchanged for the joint boundary-breach and catch-up-loss distribution studied here.

The comparison also identifies the assumptions that limit the present extension. The analytical fixed-delay result relies on a local Gaussian strong-Markov benchmark after the boundary crossing; the TWAP result is a stationary-downcrossing approximation rather than an exact discrete first-passage law; and the public calibration scores a hypothetical isolated-margin path rather than observed customer liquidations. In exchange, the framework makes the observation rule endogenous to the risk forecast and evaluates probabilities with proper scores. The contribution is therefore narrower than a general theory of crypto liquidation but more explicit about the stopping-time information set and the conditional severity target.

Against this literature, the model treats mark construction as part of the stopping rule rather than as an external price input. The delayed or smoothed mark is embedded inside the liquidation boundary, so the relevant object is the executable price error at the liquidation time. This leads to the fixed-delay and TWAP catch-up formulas derived below and to the calibration based on the same liquidation mechanism.

3. Mathematical Framework

3.1. Executable Price, Delayed Mark and Filtration

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ satisfy the usual conditions. The executable log price of the reference asset is $X_t = \log S_t$, $S_t > 0$. The exchange mark is not assumed equal to S_t; it is an observation functional

$$M_t = \mathcal{O}_{\delta,K,\eta}(S)_t, \tag{1}$$

where $\delta \geq 0$ is a discrete feed delay, K is a normalised smoothing kernel or other adapted aggregation rule, and η_t is reserved for additive measurement disturbance. Smoothing is represented by K, not by η_t. The pure fixed-delay benchmark has no averaging kernel and is

$$\log M_t = X_{(t-\delta)^+} + \eta_t. \tag{2}$$

A trailing TWAP instead uses a uniform kernel on its averaging window; it is conceptually different from additive measurement noise. Section 4.2 studies that kernel explicitly, and Appendix A records the general kernel form. The risk manager observes the mark filtration $\mathcal{F}_t^M = \sigma(M_u : u \leq t)$, coarser than the executable filtration $\mathcal{F}_t$. The gap between the two filtrations is the mathematical content of oracle latency.

Definition 1 (Stale-mark error)**.** *For the pure-delay benchmark* (2)*, the stale-mark log error is* $D_t(\delta) = X_t - X_{(t-\delta)^+} - \eta_t$*. Its absolute value is the log-price displacement hidden by the delayed mark. For a general kernel, the corresponding definition is* $D_t(K, \eta) = X_t - \log M_t$*.*

The empirical counterpart of $D_t(\delta)$ is computed from public one-minute prices by replacing $X_t - X_{t-\delta}$ with the observed δ-minute log return, yielding an evidence-based stale-error distribution without proprietary mark internals.

3.2. Funding as a Persistent Feedback State

Let F_t be the instantaneous funding state. Funding depends on premium, interest and exchange rules, but the modelling-relevant feature is persistence [2]. We specify (possibly fractional) Ornstein–Uhlenbeck dynamics [21],

$$\mathrm{d}F_t = -\kappa F_t \,\mathrm{d}t + \sigma_F \,\mathrm{d}B_t^H + \chi \,\mathrm{d}Q_t, \tag{3}$$

with mean reversion $\kappa \geq 0$, fractional Brownian motion B^H of Hurst index $H \in (1/2, 1)$ under long memory, and a net-positioning process Q_t. The model nests short-memory funding at $H = 1/2$, $\chi = 0$; the empirical layer works at $\chi = 0$ and retains the positioning term for interpreting directional stress through $q_t F_t$ below. A position q_t pays funding according to $\mathrm{d}C_t^F = -q_t S_t F_t \, \mathrm{d}t$: for a long, positive funding drains collateral; for a short, negative funding does. The stress variable is $q_t F_t$, not F_t alone.

3.3. Depth, Impact and Collateral Loss

Let $L_t(y)$ be dollar liquidity available between the mid and terminal fractional price displacement $y \geq 0$. Cumulative depth is concave at the inside and thins under stress. We represent the total execution cost by a convex function $\Phi_t(z)$ of liquidation flow z, with

$$\Phi_t(0) = 0, \qquad \Phi_t'(z) > 0, \qquad \Phi_t''(z) \geq 0. \tag{4}$$

A structural form follows from a power-law book $L(y) = ay^\gamma$, $0 < \gamma \leq 1$. If $D = L(b) = ab^\gamma$ is cumulative dollar depth inside band b, executing notional z reaches the *terminal* book displacement

$$y_\gamma(z; D) = b \, (z/D)^{1/\gamma}. \tag{5}$$

This endpoint is not the average price paid. Integrating across the consumed levels gives the average fractional execution-price concession

$$\bar{\phi}_\gamma(z; D) = \frac{1}{z} \int_0^z y_\gamma(u; D) \, \mathrm{d}u = \frac{\gamma}{1+\gamma} \, b \, (z/D)^{1/\gamma}, \tag{6}$$

and the log-equivalent concession used with log-return thresholds is

$$\ell_\gamma(z; D) = -\log(1 - \bar{\phi}_\gamma(z; D)), \qquad 0 \leq \bar{\phi}_\gamma(z; D) < 1. \tag{7}$$

Thus, y_γ describes the terminal level reached in the book, $\bar{\phi}_\gamma$ describes the average execution price, and ℓ_γ is its log-return representation. Thin depth raises all three. A smaller γ raises terminal displacement when $z/D > 1$; for average concession the ordering requires sufficiently stressed utilisation, with the local condition $\log(z/D) > \gamma/(1+\gamma)$. Appendix A gives the corresponding total-cost identity.

3.4. Margin Surplus, the Isolated-Margin Boundary, and Liquidation Time

Consider an isolated-margin long with leverage $L > 1$, collateral C_0, notional $N_0 = LC_0$ and entry P_0. At price P_t, the equity is $E_t = C_0 + N_0(P_t/P_0 - 1)$, and the maintenance requirement is $mN_0(P_t/P_0)$ for maintenance fraction $m \in (0, 1)$. Liquidation occurs when $E_t \leq mN_0(P_t/P_0)$, which solving gives the liquidation price ratio and simple-return drop

$$\rho^* = \frac{1 - 1/L}{1 - m}, \qquad d_L = 1 - \rho^* = \frac{1/L - m}{1 - m}, \qquad b_L = \log(1 - d_L), \tag{8}$$

so the long is liquidated when the cumulative executable log return reaches $b_L < 0$. The equity remaining at the boundary, in collateral units, is

$$\xi(L) \;=\; \frac{E_\tau}{C_0}\bigg|_{P_t/P_0 = \rho^*} = \frac{m(L-1)}{1-m}, \tag{9}$$

which equals the maintenance margin held at the boundary, $mL\rho^*$.

Let G_t^F denote cumulative funding and fee drain since entry, and write $g_t = G_t^F / C_0$. When equity and maintenance are evaluated on the mark, normalised surplus is

$$\frac{Y_t^{\delta,G}}{C_0} = 1 + L\left(\frac{M_t}{P_0} - 1\right) - g_t - mL\frac{M_t}{P_0}. \tag{10}$$

For the pure-delay mark (2), this is equivalent to a generally time-varying log-price boundary

$$b_L(g_t) = \log\left(\frac{1 - 1/L + g_t/L}{1 - m}\right), \qquad \tau_m(\delta, G) = \inf\{t \geq 0 : X_{(t-\delta)^+} + \eta_t \leq b_L(g_t)\}. \tag{11}$$

The central risk object is $\Pi(T; m, \delta, H, \Phi) = \mathbb{P}(\tau_m(\delta, G) \leq T)$.

The simple time shift requires more than a delayed price. If $\eta_t = 0$ and funding, fees, maintenance margin, and the position are fixed over the relevant interval so that $b_L(g_t) = b_L$ is constant, continuity gives the price-channel benchmark

$$\tau_m^{\text{price}}(\delta) = \tau_m^{\text{price}}(0) + \delta, \qquad \mathbb{P}\{\tau_m^{\text{price}}(\delta) \leq T\} = \mathbb{P}\{\tau_m^{\text{price}}(0) \leq T - \delta\}, \quad T \geq \delta. \tag{12}$$

If g_t evolves during the delay, then $b_L(g_t)$ moves, and (12) need not hold; no exact time-shift claim is made for the full funded surplus (10). The first-passage simulation evaluates funding contemporaneously through the full surplus. The fixed-delay experiment and Theorem 1 deliberately isolate the conditional price-latency channel by holding the boundary and non-price drains fixed during the short detection interval. Under that benchmark, delay can lower the number of liquidations recorded by a fixed horizon by postponing detection while increasing conditional loss severity once detection occurs.

4. Theoretical Results

We first record that delay enlarges hidden displacement, then establish the stale-boundary and funding-variance results, give a comparison principle for pointwise ordered losses, and derive the cost of delay for fixed-lag and TWAP marks.

The results form a cumulative argument. The dispersion lemma identifies how much executable movement can be hidden by observation; the stale-boundary proposition translates that hidden movement into overstated surplus; the funding result describes how the boundary itself can move through cumulative collateral drain; and the comparison result identifies when execution cost raises the crossing probability. The fixed-delay and TWAP theorems then condition on a crossing and supply the severity component of the same forecast. This sequence separates event frequency from conditional loss without treating the mechanisms as unrelated outcomes.

Assumption 1 (Locally sub-Gaussian executable increments)**.** *For each $\delta > 0$, conditional on $\mathcal{F}_{t-\delta}$, the increment $X_t - X_{t-\delta}$ is centred with conditional variance $v_t(\delta)$ and satisfies $\mathbb{P}(|X_t - X_{t-\delta}| > u \mid \mathcal{F}_{t-\delta}) \leq 2\exp\{-u^2/[2v_t(\delta)]\}$.*

This is a tractable benchmark, not a denial of heavy tails; the simulation stresses the boundary beyond the Gaussian core.

Lemma 1 (Delay increases stale-price dispersion)**.** *Suppose that, over the intraday window, X is a square-integrable local martingale. Then, $\mathbb{E}[D_t(\delta)^2]$ is non-decreasing in δ under the benchmark $\eta_t = 0$. Stationary increments are not required for this ordering.*

Proof. With $\eta_t = 0$, $D_t(\delta) = X_t - X_{t-\delta}$. The martingale isometry gives

$$\mathbb{E}[D_t(\delta)^2] = \mathbb{E}[\langle X \rangle_t - \langle X \rangle_{t-\delta}].$$

For $0 < \delta_1 < \delta_2$,

$$\begin{aligned}\mathbb{E}[D_t(\delta_2)^2] - \mathbb{E}[D_t(\delta_1)^2] &= \mathbb{E}\big[\langle X\rangle_t - \langle X\rangle_{t-\delta_2}\big] - \mathbb{E}\big[\langle X\rangle_t - \langle X\rangle_{t-\delta_1}\big] \\ &= \mathbb{E}\big[\langle X\rangle_{t-\delta_1} - \langle X\rangle_{t-\delta_2}\big] \geq 0.\end{aligned}$$

The inequality follows because quadratic variation is increasing in time. Stationary increments would make the expected dispersion a function of elapsed lag alone, but they are unnecessary for the nested-interval inequality. If a centred measurement disturbance η_t has the same variance for every compared lag and is orthogonal to $X_t - X_{t-\delta}$ for every δ (independence is sufficient), then

$$\mathbb{E}\Big[(X_t - X_{t-\delta} - \eta_t)^2\Big] = \mathbb{E}\Big[(X_t - X_{t-\delta})^2\Big] + \mathbb{E}[\eta_t^2],$$

so it adds the same term and preserves the ordering. Without that orthogonality condition, the covariance term can depend on δ, and the ordering need not be preserved. □

Proposition 1 (Boundary shrinking by a stale mark)**.** *Let the long be marked at the stale price M_t with $S_t = M_t e^{D_t}$ and adverse staleness $D_t < 0$. Then, to first order, executable equity is below mark-reported equity by the positive amount $-\tilde{N}_t D_t$, where $\tilde{N}_t = N_0 M_t / P_0$. The maintenance-requirement error is $-m\tilde{N}_t D_t$, smaller by the factor m. Hence, the dominant effect of a stale mark is to overstate equity, not to mis-state the requirement.*

Proof. Equity evaluated on the executable price and on the mark is

$$E_t = C_0 + N_0\left(\frac{S_t}{P_0} - 1\right), \qquad \tilde{E}_t = C_0 + N_0\left(\frac{M_t}{P_0} - 1\right).$$

Since $S_t = M_t e^{D_t}$ and $\tilde{N}_t = N_0 M_t / P_0$,

$$E_t - \tilde{E}_t = \frac{N_0 M_t}{P_0}(e^{D_t} - 1) = \tilde{N}_t D_t + O(D_t^2).$$

For $D_t < 0$, the executable equity is therefore lower than the mark-reported equity by

$$\tilde{E}_t - E_t = -\tilde{N}_t D_t + O(D_t^2) = \tilde{N}_t |D_t| + O(D_t^2).$$

The maintenance requirement evaluated on S_t rather than M_t changes by

$$m\frac{N_0 M_t}{P_0}(e^{D_t} - 1) = m\tilde{N}_t D_t + O(D_t^2),$$

a factor m smaller. Equivalently, if $Y_t = E_t - mN_0 S_t / P_0$ and $\tilde{Y}_t = \tilde{E}_t - mN_0 M_t / P_0$, then

$$Y_t - \tilde{Y}_t = (1-m)\tilde{N}_t(e^{D_t} - 1) = (1-m)\tilde{N}_t D_t + O(D_t^2).$$

Thus, $\tilde{Y}_t - Y_t = (1-m)\tilde{N}_t |D_t| + O(D_t^2) > 0$ for adverse staleness, so the stale mark overstates surplus. The leading term comes from equity because the requirement term is multiplied by $m < 1$. □

Remark 1. *This corrects a common understatement: because exchanges mark equity at the stale price, a delayed mark hides adverse moves at full mark-notional sensitivity $\tilde{N}_t = N_0 M_t / P_0$. The requirement channel, of order $m\tilde{N}_t$, is secondary. The practical consequence is the catch-up risk formalised in Theorem 1.*

Proposition 2 (Funding memory raises cumulative-drain variance)**.** *Let funding be driven by fractional noise with $H \in (1/2, 1)$ and no mean reversion ($\kappa = 0$), so $F_t = \sigma_F B_t^H$. For constant exposure over $[0, T]$,*

$$\mathrm{Var}\left(\int_0^T F_t \, \mathrm{d}t\right) = \frac{\sigma_F^2}{2H+2} T^{2H+2} \propto T^{2H+2}. \tag{13}$$

Since $2H + 2 > 3$ for $H > 1/2$, the persistent case has strictly stronger horizon scaling than the Brownian benchmark (T^3). With $\kappa > 0$ and a non-stationary initial condition, this local scaling is relevant only for $T \ll 1/\kappa$. A stationary fractional Ornstein–Uhlenbeck process has covariance of order t^{2H-2}, so the variance of its time integral grows as T^{2H} at long horizons rather than as T^{2H+2}.

Proof. By self-similarity, $B_{Tv}^H \overset{d}{=} T^H B_v^H$. Hence,

$$\int_0^T B_t^H \, \mathrm{d}t = T \int_0^1 B_{Tv}^H \, \mathrm{d}v \overset{d}{=} T^{H+1} \int_0^1 B_v^H \, \mathrm{d}v.$$

It remains only to compute the unit-interval variance. With

$$\mathrm{Cov}(B_u^H, B_v^H) = \frac{1}{2}\{u^{2H} + v^{2H} - |u - v|^{2H}\},$$

$$\begin{aligned} \mathrm{Var}\left(\int_0^1 B_v^H \, \mathrm{d}v\right) &= \int_0^1 \int_0^1 \mathrm{Cov}(B_u^H, B_v^H) \, \mathrm{d}u \, \mathrm{d}\, v \\ &= \frac{1}{2H+1} - \frac{1}{(2H+1)(2H+2)} = \frac{1}{2H+2}. \end{aligned}$$

Multiplying by $\sigma_F^2 T^{2H+2}$ gives the stated formula. At $H = 1/2$, this becomes $\sigma_F^2 T^3/3$, the Brownian benchmark. For $\kappa > 0$, mean reversion changes the covariance kernel after the relaxation scale $1/\kappa$. In the stationary long-memory regime, integrating a covariance proportional to t^{2H-2} gives long-horizon variance proportional to T^{2H}. □

Remark 2 (Numerical validation of the scaling exponent)**.** *We validate Proposition 2 directly: we simulate exact fractional Gaussian noise (Cholesky factorisation of the fractional-noise covariance), form $\int_0^T B_t^H \, \mathrm{d}t$ on a unit grid, and regress* $\log \mathrm{Var}$ *on* $\log T$ *over $T \in \{32, 64, 128, 256\}$. The fitted exponents are* 2.99, 3.36 *and* 3.79 *at $H = 0.5, 0.7, 0.9$, against the predicted $2H + 2 = 3.0, 3.4, 3.8$ (Figure 1). This power-law exponent is distinct from the variance leve multiplier $(1 + \phi)/(1 - \phi)$ of the short-memory AR(1) drain used in the first-passage experiment: AR(1) raises the level of cumulative-drain variance (to $\approx 6\times$ the i.i.d. value at $\widehat{\phi} = 0.71$) without steepening the horizon scaling. The two are separate statements, and only the fractional simulation tests Proposition 2.*

Proposition 3 (Comparison principle for ordered losses)**.** *Let $Y_t^{(j)} = Y_0 - A_t - G_t - I_t^{(j)}$, $j \in \{1, 2\}$, be two surplus processes coupled on the same price and funding paths. If $I_t^{(2)} \geq I_t^{(1)}$ almost surely for every $t \leq T$, then*

$$\mathbb{P}\left(\inf_{0 \leq t \leq T} Y_t^{(2)} \leq 0\right) \geq \mathbb{P}\left(\inf_{0 \leq t \leq T} Y_t^{(1)} \leq 0\right).$$

For terminal displacement (5), decreasing γ produces this order when $z/D > 1$ and reverses it when $z/D < 1$. For the average concession (6), the $\gamma/(1+\gamma)$ factor means that $z/D \geq 1$ alone is insufficient; the actual concessions must be ordered, with the local sufficient direction

$\log(z/D) > \gamma/(1+\gamma)$. *No analogous pathwise order is asserted for two observation delays or two Hurst indices.*

Proof. The assumed loss order gives

$$Y_t^{(2)} = Y_0 - A_t - G_t - I_t^{(2)} \leq Y_0 - A_t - G_t - I_t^{(1)} = Y_t^{(1)}$$

for every $t \leq T$. Hence,

$$\mathbf{1}_{\{\inf_t Y_t^{(1)} \leq 0\}} \leq \mathbf{1}_{\{\inf_t Y_t^{(2)} \leq 0\}}.$$

Taking expectations proves the probability order. For $u = z/D$, terminal displacement (5) is $bu^{1/\gamma}$, which increases as γ decreases when $u > 1$ and decreases when $0 < u < 1$. Average concession is $[\gamma/(1+\gamma)]bu^{1/\gamma}$, and

$$\frac{\partial}{\partial\gamma} \log \bar{\phi}_\gamma = \frac{1}{\gamma(1+\gamma)} - \frac{\log u}{\gamma^2}.$$

Thus, reducing γ locally raises the average concession only when $\log u > \gamma/(1+\gamma)$. A pathwise probability order follows only after the relevant execution-loss functions are actually ordered over the utilisation range. □

Remark 3. *Variance growth alone does not imply stochastic ordering of a running maximum against a moving boundary. The funding experiment therefore reports paired differences rather than treating Proposition 2 as a monotonicity theorem. Delay is also not an additive pathwise loss: it postpones the recorded crossing according to* (12) *while increasing conditional loss severity according to Theorem 1.*

4.1. The Cost of Oracle Delay: Catch-Up Overshoot and Bad-Debt Probability

Let the pure-delay mark satisfy $\log M_t = X_{t-\delta}$, and let the executable first-passage time be $\tau = \inf\{t : X_t \leq b_L\}$, with b_L from (8).

Assumption 2 (Fixed-delay Gaussian stopping benchmark)**.** *The executable log price is continuous and, locally over $[\tau, \tau+\delta]$, has the strong-Markov representation*

$$X_{\tau+u} - X_\tau = \sigma(W_{\tau+u} - W_\tau), \qquad 0 \leq u \leq \delta,$$

with constant $\sigma > 0$ and Brownian motion W. The maintenance fraction m is constant and does not depend on leverage; measurement noise is absent; funding, fees, position size, and the price boundary are held fixed or absorbed into the boundary during the detection interval. Execution loss is applied separately through the log-equivalent average concession (7).

Under Assumption 2, the delayed mark reaches the same fixed boundary at $\sigma_\delta = \tau + \delta$, and $\Delta X_\delta = X_{\tau+\delta} - X_\tau$ is independent of $\mathcal{F}_\tau$ with distribution $N(0, \sigma^2\delta)$. These properties are stronger than Assumption 1; the general sub-Gaussian condition alone does not imply the theorem below.

Theorem 1 (Catch-up overshoot and bad-debt probability)**.** *Under Assumption 2, Define the catch-up overshoot $O(\delta) = \max(0, -\Delta X_\delta)$, the additional adverse executable move realised between executable crossing and mark liquidation. Then, to leading order in the move size,*

$$\mathbb{E}[O(\delta)] = \sigma\sqrt{\frac{\delta}{2\pi}}. \tag{14}$$

Working the isolated-margin profit and loss exactly, the equity at mark liquidation is $E_{\sigma_\delta}/C_0 = 1 + L(\rho^* e^{\Delta X_\delta} - 1)$*, and it is non-negative iff* $\Delta X_\delta \geq \log(1-m)$*. A bad-debt (negative-equity) liquidation therefore occurs iff* $\Delta X_\delta < \log(1-m)$*, with probability*

$$P_{\mathrm{bd}}(\delta) = \Phi\left(\frac{\log(1-m)}{\sigma\sqrt{\delta}}\right) = \Phi\left(-\frac{m}{\sigma\sqrt{\delta}}\right)(1+O(m)), \tag{15}$$

exactly independent of leverage within the constant-m benchmark: the return-space cushion between the liquidation price and zero equity is always the same maintenance fraction m*.* P_{bd} *is strictly increasing in* δ *and* σ *and is governed by the dimensionless latency–margin ratio* $\sigma\sqrt{\delta}/m$*. With average execution concession on liquidated notional* $z = N_\tau$*, the threshold tightens consistently in log units to* $\Delta X_\delta < \log(1-m) + \ell_\gamma(z; D)$*. Thinner depth raises* P_{bd}*; a smaller* γ *raises it only over utilisation ranges for which (6) is actually larger, and execution reintroduces size dependence.*

Proof. Equity at delayed liquidation. At the executable crossing, $P_\tau/P_0 = \rho^*$. At the delayed mark liquidation time $\sigma_\delta = \tau + \delta$,

$$\frac{P_{\sigma_\delta}}{P_0} = \rho^* e^{\Delta X_\delta}, \qquad \frac{E_{\sigma_\delta}}{C_0} = 1 + L\left(\rho^* e^{\Delta X_\delta} - 1\right).$$

Using $\rho^* = (1 - 1/L)/(1-m)$,

$$\begin{aligned} E_{\sigma_\delta} \geq 0 &\iff 1 + L(\rho^* e^{\Delta X_\delta} - 1) \geq 0 \\ &\iff \rho^* e^{\Delta X_\delta} \geq 1 - \frac{1}{L} \\ &\iff e^{\Delta X_\delta} \geq \frac{1 - 1/L}{\rho^*} = 1 - m \\ &\iff \Delta X_\delta \geq \log(1-m). \end{aligned}$$

The leverage term cancels in the third line.

Overshoot. Put $s = \sigma\sqrt{\delta}$ and $\Delta X_\delta = sZ$, with $Z \sim N(0,1)$. Then,

$$\mathbb{E}[O(\delta)] = \mathbb{E}[(-sZ)^+] = s\int_{-\infty}^{0} (-z)\varphi(z)\,\mathrm{d}z = s\varphi(0) = \sigma\sqrt{\frac{\delta}{2\pi}}.$$

Bad debt. The previous threshold gives

$$P_{\mathrm{bd}}(\delta) = \mathbb{P}\{\Delta X_\delta < \log(1-m)\} = \Phi\left(\frac{\log(1-m)}{\sigma\sqrt{\delta}}\right).$$

Since $\log(1-m) = -m + O(m^2)$, (15) follows. For $a = \log(1-m) < 0$, the map $s \mapsto \Phi(a/s)$ is increasing for $s > 0$, so P_{bd} is increasing in both δ and σ.

Depth impact. If liquidation pays the average fractional execution concession $\bar{\phi}_\gamma(z; D)$, its log-equivalent concession is $\ell_\gamma(z; D)$ from (7). The effective post-liquidation log move is $\Delta X_\delta - \ell_\gamma(z; D)$. Negative equity then occurs when

$$\Delta X_\delta - \ell_\gamma(z; D) < \log(1-m), \quad \text{or} \quad \Delta X_\delta < \log(1-m) + \ell_\gamma(z; D).$$

The impact term is increasing when depth D falls. Its ordering in γ is utilisation-dependent, as specified in Proposition 3. □

Remark 4 (What breaks leverage independence)**.** *The leverage cancellation is conditional, not universal. If* $m = m(L)$ *is tiered, then the threshold contains* $\log(1 - m(L))$*. If an additional*

normalised funding-and-fee drain $g_\delta = (G^F_{\tau+\delta} - G^F_\tau)/C_0$ *accrues between crossing and detection, equity after the average execution concession is*

$$\frac{E^{\text{net}}_{\sigma_\delta}}{C_0} = 1 + L\left(\rho^* e^{\Delta X_\delta - \ell_\gamma(z;D)} - 1\right) - g_\delta,$$

and bad debt occurs when

$$\Delta X_\delta < \log(1-m) + \ell_\gamma(z;D) + \log\left(1 + \frac{g_\delta}{L-1}\right). \tag{16}$$

Incremental funding therefore moves the threshold and generally reintroduces leverage dependence. Equation (15) *is the isolated price-latency benchmark with* $g_\delta = 0$*, constant* m*, and zero execution concession.*

Corollary 1 (Latency–margin scaling)**.** *A protocol holds bad-debt probability from delay below a target* $p^* < 1/2$ *iff*

$$\sigma\sqrt{\delta} \le \frac{-\log(1-m)}{\Phi^{-1}(1-p^*)}$$

and, to first order in m*, iff* $\sigma\sqrt{\delta} \le m/\Phi^{-1}(1-p^*)$*. This is leverage-independent only under Assumption 2: constant* m*, fixed non-price drains, and zero execution concession. Mark delay and maintenance margin are substitutes within that benchmark: doubling tolerated delay requires raising* m *by a factor* $\sqrt{2}$ *to hold bad-debt risk fixed.*

Proof. By Theorem 1, $P_{\text{bd}} \le p^*$ is equivalent to

$$\Phi\left(\frac{\log(1-m)}{\sigma\sqrt{\delta}}\right) \le p^*.$$

Since $p^* < 1/2$, $\Phi^{-1}(p^*) < 0$. Therefore,

$$\frac{\log(1-m)}{\sigma\sqrt{\delta}} \le \Phi^{-1}(p^*).$$

Multiplying by the positive denominator and using that both $\log(1-m)$ and $\Phi^{-1}(p^*)$ are negative gives

$$\sigma\sqrt{\delta} \le \frac{\log(1-m)}{\Phi^{-1}(p^*)} = \frac{-\log(1-m)}{\Phi^{-1}(1-p^*)}.$$

Finally,

$$-\log(1-m) = m + \frac{m^2}{2} + O(m^3),$$

so the first-order condition is $\sigma\sqrt{\delta} \le m/\Phi^{-1}(1-p^*)$. Holding p^* fixed keeps $m/\sqrt{\delta}$ constant to first order; hence, $\delta \mapsto 2\delta$ requires $m \mapsto \sqrt{2}\,m$. □

Equations (14) and (15) are the operational core: the cost of oracle delay is a $\sqrt{\delta}$ overshoot and a Gaussian-tail bad-debt probability set by $\sigma\sqrt{\delta}/m$. Section 6 confirms both against simulation and one-minute data.

Corollary 2 (Leverage and the breach probability)**.** *The liquidation boundary distance* (8) *satisfies* $|b_L| \approx 1/L - m$*, decreasing in* L*, so the breach probability* $\Pi(T;\cdot)$ *increases in leverage. Conditional on delayed liquidation, with constant* m *and before incremental funding, fees, or execution costs, the bad-debt probability in Theorem 1 is leverage-independent. Leverage changes the frequency of reaching liquidation, whereas delay changes conditional severity in that benchmark.*

Proof. On the admissible range $1 < L < 1/m$,

$$d_L = \frac{1/L - m}{1 - m}, \qquad \frac{\partial d_L}{\partial L} = -\frac{1}{L^2(1-m)} < 0.$$

The log boundary is

$$b_L = \log(1 - d_L) < 0, \qquad |b_L| = -\log(1 - d_L) = d_L + O(d_L^2).$$

Thus, $|b_L|$ decreases with L: the boundary moves closer to the initial point. Under the usual coupling of first-passage experiments, if $L_2 > L_1$, then $b_{L_2} > b_{L_1}$, and

$$\{\inf_{t \le T} X_t \le b_{L_1}\} \subseteq \{\inf_{t \le T} X_t \le b_{L_2}\},$$

so the breach probability is non-decreasing in leverage. By contrast, Theorem 1 gives the bad-debt condition

$$\Delta X_\delta < \log(1 - m),$$

which contains no L. Breach risk and bad-debt risk therefore depend on different parameters. □

4.2. The Cost of a Smoothed (TWAP) Mark at the Liquidation Time

Production mark engines rarely use a pure fixed delay; they smooth, most commonly through a trailing time-weighted average price (TWAP), partly because a smoothed mark is harder to manipulate [3]. Two heuristics suggest smoothing is benign. The mean-lag heuristic treats a TWAP of window δ as a fixed delay at the kernel's mean lag $\int_0^\delta u\, K_\delta(u)\, \mathrm{d}u = \delta/2$; the variance heuristic (Lemma 2 below) notes that the unconditional stale error carries only the dispersion of a fixed delay $\delta/3$. Both heuristics evaluate the stale error at a *fixed* time. Liquidation does not: it samples the error at the first *downcrossing* of the boundary by the smoothed mark, an event that requires the trailing window to have fallen. That conditioning is first-order, and it reverses the comparison.

Lemma 2 (TWAP stale-error law)**.** *Let the mark be the trailing uniform TWAP* $M_t = \exp\left(\frac{1}{\delta}\int_{t-\delta}^t X_u\, \mathrm{d}u\right)$ *of window* δ*, let* X *be a driftless local martingale with local volatility* σ*, and write* $D_t^{\mathrm{TWAP}} = X_t - \frac{1}{\delta}\int_{t-\delta}^t X_u\, \mathrm{d}u$ *for the stale error and* $\Delta_t = X_t - X_{t-\delta}$ *for the window increment, so that* $\log M_t$ *is differentiable with derivative* Δ_t/δ*. To leading order* $(D_t^{\mathrm{TWAP}}, \Delta_t)$ *is a centred bivariate Gaussian with*

$$\mathrm{Var}\left(D_t^{\mathrm{TWAP}}\right) = \frac{\sigma^2\delta}{3}, \qquad \mathrm{Var}(\Delta_t) = \sigma^2\delta, \qquad \mathrm{Cov}\left(D_t^{\mathrm{TWAP}}, \Delta_t\right) = \frac{\sigma^2\delta}{2}, \tag{17}$$

and hence, $\mathbb{E}[D_t^{\mathrm{TWAP}} \mid \Delta_t] = \Delta_t/2$*, and* $\mathrm{Var}(D_t^{\mathrm{TWAP}} \mid \Delta_t) = \sigma^2\delta/12$*. The unconditional dispersion is one-third—not one-half—of the fixed-delay dispersion* $\mathrm{Var}(D_t(\delta)) = \sigma^2\delta$*; a K-sample discrete TWAP with grid spacing* $h = \delta/K$ *has exact variance* $\sigma^2\delta(K-1)(2K-1)/(6K^2) \to \sigma^2\delta/3$*.*

Proof. Write

$$D_t^{\mathrm{TWAP}} = \frac{1}{\delta}\int_0^\delta (X_t - X_{t-a})\, \mathrm{d}a, \qquad \Delta_t = X_t - X_{t-\delta}.$$

For a martingale with local volatility σ,

$$\mathrm{Cov}(X_t - X_{t-a}, X_t - X_{t-b}) = \sigma^2 \min(a, b), \qquad \mathrm{Cov}(X_t - X_{t-a}, \Delta_t) = \sigma^2 a$$

for $0 \le a, b \le \delta$. Hence,

$$\mathrm{Var}(D_t^{\mathrm{TWAP}}) = \frac{\sigma^2}{\delta^2}\int_0^\delta\int_0^\delta \min(a,b)\,\mathrm{d}a\,\mathrm{d}b = \frac{\sigma^2\delta}{3},$$
$$\mathrm{Cov}(D_t^{\mathrm{TWAP}}, \Delta_t) = \frac{\sigma^2}{\delta}\int_0^\delta a\,\mathrm{d}a = \frac{\sigma^2\delta}{2},$$

because

$$\int_0^\delta\int_0^\delta \min(a,b)\,\mathrm{d}a\,\mathrm{d}b = 2\int_0^\delta\int_0^b a\,\mathrm{d}a\,\mathrm{d}b = \frac{\delta^3}{3}.$$

Also, $\mathrm{Var}(\Delta_t) = \sigma^2\delta$. Therefore, the covariance matrix of $(D_t^{\mathrm{TWAP}}, \Delta_t)$ is

$$\Sigma = \sigma^2\delta\begin{pmatrix} 1/3 & 1/2 \\ 1/2 & 1 \end{pmatrix}.$$

The Gaussian conditioning formula gives

$$\mathbb{E}[D_t^{\mathrm{TWAP}} \mid \Delta_t] = \frac{\mathrm{Cov}(D_t^{\mathrm{TWAP}}, \Delta_t)}{\mathrm{Var}(\Delta_t)}\Delta_t = \frac{\Delta_t}{2},$$

and

$$\mathrm{Var}(D_t^{\mathrm{TWAP}} \mid \Delta_t) = \frac{\sigma^2\delta}{3} - \frac{(\sigma^2\delta/2)^2}{\sigma^2\delta} = \frac{\sigma^2\delta}{12}.$$

For the discrete identity, write $D = \frac{1}{K}\sum_{i=1}^K (i-1)g_i$, where the i.i.d. increments g_i have variance $\sigma^2 h$ and $h = \delta/K$. Hence,

$$\mathrm{Var}(D) = \frac{\sigma^2 h}{K^2}\sum_{j=0}^{K-1} j^2 = \frac{\sigma^2\delta(K-1)(2K-1)}{6K^2} \to \frac{\sigma^2\delta}{3}.$$

□

Theorem 2 (Stationary-crossing approximation for a TWAP mark)**.** *Let liquidation occur at the first downcrossing τ_M of the boundary by the TWAP mark, at a crossing time large relative to the window. Approximating the first downcrossing by a stationary downcrossing, the leading-order stale-error law is the Rayleigh-size-biased Gaussian mixture*

$$D_{\tau_M}^{\mathrm{TWAP}} \overset{d}{\approx} \sigma\sqrt{\delta}\left(-\tfrac{1}{2}R + \tfrac{1}{\sqrt{12}}Z\right), \qquad R \sim \mathrm{Rayleigh}(1) \text{ independent of } Z \sim N(0,1), \quad (18)$$

with mean $-\frac{\sqrt{2\pi}}{4}\sigma\sqrt{\delta} \approx -0.627\,\sigma\sqrt{\delta}$: when the smoothed mark reaches the boundary, the executable price has systematically led it down. Consequently,

$$\mathbb{E}\big[O^{\mathrm{TWAP}}(\delta)\big] = c_O\,\sigma\sqrt{\delta}, \qquad P_{\mathrm{bd}}^{\mathrm{TWAP}}(\delta) = \mathbb{P}\left(\frac{R}{2} - \frac{Z}{\sqrt{12}} > \frac{-\log(1-m)}{\sigma\sqrt{\delta}}\right), \quad (19)$$

where $c_O = \mathbb{E}\big[(R/2 - Z/\sqrt{12})^+\big] = 0.637$, and the bad-debt probability is a one-dimensional Rayleigh–Gaussian integral. Both leading-order objects are, like the fixed-delay case, independent of leverage only within the constant-m, fixed-non-price-drain benchmark; tiered maintenance, funding during detection, or execution costs can restore leverage and size dependence.

Proof. Reduction to stale error at the crossing. At τ_M, $\log M_{\tau_M} = b_L$. Since

$$D_{\tau_M}^{\mathrm{TWAP}} = X_{\tau_M} - \log M_{\tau_M},$$

the executable log price is $b_L + D^{\mathrm{TWAP}}_{\tau_M}$. If $D^{\mathrm{TWAP}}_{\tau_M} < 0$, the executable price has already moved below the mark boundary. The isolated-margin calculation in Theorem 1 therefore applies with ΔX_δ replaced by $D^{\mathrm{TWAP}}_{\tau_M}$, so the leverage cancellation is unchanged under the same constant-m and fixed-non-price-drain assumptions.

Slope-biased downcrossing law. Lemma 2 gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\log M_t = \frac{X_t - X_{t-\delta}}{\delta} = \frac{\Delta_t}{\delta}.$$

A downcrossing has $\Delta_t < 0$. Kac–Rice conditioning weights a crossing by the absolute slope, so for $x < 0$,

$$f_{\Delta|\mathrm{down}}(x) = \frac{|x| f_\Delta(x)}{\int_{-\infty}^{0} |u| f_\Delta(u)\,\mathrm{d}u}, \qquad f_\Delta(x) = \frac{1}{\sigma\sqrt{\delta}}\,\varphi\!\left(\frac{x}{\sigma\sqrt{\delta}}\right).$$

With $x = -\sigma\sqrt{\delta}\,r$, this density becomes

$$f_R(r) = re^{-r^2/2}, \qquad r > 0,$$

so $\Delta_{\tau_M} \overset{d}{=} -\sigma\sqrt{\delta}\,R$ with $R \sim \mathrm{Rayleigh}(1)$.

Conditional stale error. From Lemma 2,

$$D^{\mathrm{TWAP}} \mid \Delta \sim N\!\left(\frac{\Delta}{2}, \frac{\sigma^2\delta}{12}\right).$$

Therefore, at a stationary downcrossing,

$$D^{\mathrm{TWAP}}_{\tau_M} \overset{d}{=} \sigma\sqrt{\delta}\left(-\frac{1}{2}R + \frac{1}{\sqrt{12}}Z\right), \qquad Z \sim N(0,1), \quad Z \perp R.$$

The residual level-conditioning term satisfies

$$\mathbb{E}[D^{\mathrm{TWAP}}_t \mid \log M_t = b_L] = \frac{\mathrm{Cov}(D^{\mathrm{TWAP}}_t, \log M_t)}{\mathrm{Var}(\log M_t)}\, b_L = O(\delta/t),$$

which is negligible when the crossing time is large relative to the averaging window.

Overshoot and bad debt. Since $\mathbb{E}[R] = \sqrt{\pi/2}$,

$$\mathbb{E}[D^{\mathrm{TWAP}}_{\tau_M}] = -\frac{\sigma\sqrt{\delta}}{2}\mathbb{E}[R] = -\frac{\sqrt{2\pi}}{4}\sigma\sqrt{\delta}.$$

The catch-up overshoot is

$$O^{\mathrm{TWAP}}(\delta) = \left(-D^{\mathrm{TWAP}}_{\tau_M}\right)^+ = \sigma\sqrt{\delta}\left(\frac{R}{2} - \frac{Z}{\sqrt{12}}\right)^+,$$

so

$$\mathbb{E}[O^{\mathrm{TWAP}}(\delta)] = c_O\sigma\sqrt{\delta}, \qquad c_O = \mathbb{E}\left[\left(\frac{R}{2} - \frac{Z}{\sqrt{12}}\right)^+\right] = 0.637.$$

Negative equity occurs when $D^{\mathrm{TWAP}}_{\tau_M} < \log(1-m)$, equivalently,

$$\frac{R}{2} - \frac{Z}{\sqrt{12}} > \frac{-\log(1-m)}{\sigma\sqrt{\delta}},$$

which gives (19). The first-passage condition excludes earlier mark crossings, and discrete monitoring introduces boundary overshoot. The mechanism simulation of Section 6 mea-

sures their combined effect. The displayed law is therefore an approximation, not a proved bound for the first-passage distribution. □

Remark 5 (Why the fixed-time heuristics fail)**.** *Substituting the unconditional variance* $\sigma^2\delta/3$ *into Theorem 1 as if liquidation sampled the stale error at a fixed time would predict* $\mathbb{E}[O^{\mathrm{TWAP}}] = \sigma\sqrt{\delta/6\pi} \approx 0.23\,\sigma\sqrt{\delta}$ *and a far thinner tail; the mean-lag heuristic predicts* $0.28\,\sigma\sqrt{\delta}$*. The crossing-conditioned value* $0.637\,\sigma\sqrt{\delta}$ *is* 2.3–2.8 *times larger than either, because liquidation occurs precisely when the trailing window has fallen, an event under which the error is centred near* $-0.63\,\sigma\sqrt{\delta}$ *rather than zero. Only the pure-delay (Dirac-kernel) mark escapes this conditioning through the strong Markov property.*

Corollary 3 (Stationary-crossing comparison with fixed delay)**.** *Under the approximation of Theorem 2, the same-window overshoot ratio is* $\mathbb{E}[O^{\mathrm{TWAP}}]/\mathbb{E}[O] = c_O\sqrt{2\pi} \approx 1.60$*: the TWAP overshoot equals that of a fixed delay of about* $2.6\,\delta$*, not* $\delta/3$*. Numerical solution of the bad-debt comparison gives a crossover* $\kappa^* \approx 1.12$ *in* $\kappa = -\log(1-m)/(\sigma\sqrt{\delta})$*. Within this approximation, smoothing raises the bad-debt probability when the latency–margin ratio exceeds about* 0.89 *and lowers it below that threshold.*

Proof. The overshoot ratio follows directly from Theorems 1 and 2:

$$\frac{\mathbb{E}[O^{\mathrm{TWAP}}]}{\mathbb{E}[O]} = \frac{c_O\sigma\sqrt{\delta}}{\sigma\sqrt{\delta/(2\pi)}} = c_O\sqrt{2\pi} \approx 1.60.$$

Equating this to the fixed-delay formula gives

$$\sigma\sqrt{\frac{\delta_{\mathrm{fix}}}{2\pi}} = c_O\sigma\sqrt{\delta} \quad\Longrightarrow\quad \delta_{\mathrm{fix}} = (c_O\sqrt{2\pi})^2\delta \approx 2.6\delta.$$

For the tail comparison, set

$$\kappa = \frac{-\log(1-m)}{\sigma\sqrt{\delta}} \approx \frac{m}{\sigma\sqrt{\delta}}.$$

The fixed-delay tail is

$$p_{\mathrm{fix}}(\kappa) = \Phi(-\kappa),$$

while Theorem 2 gives the TWAP tail as the one-dimensional integral

$$p_{\mathrm{TWAP}}(\kappa) = \int_0^\infty \Phi\Big(\sqrt{12}\Big(\frac{r}{2} - \kappa\Big)\Big) r e^{-r^2/2}\,\mathrm{d}r.$$

Integration by parts gives

$$p_{\mathrm{TWAP}}(\kappa) = \Phi(-\sqrt{12}\,\kappa) + \frac{\sqrt{3}}{2} e^{-3\kappa^2/2}\Phi(3\kappa).$$

The difference

$$h(\kappa) = p_{\mathrm{TWAP}}(\kappa) - p_{\mathrm{fix}}(\kappa)$$

is evaluated numerically. It changes sign at $\kappa^* = 1.1225$ on the positive parameter range: $h(\kappa) > 0$ below the crossing and $h(\kappa) < 0$ above it. Since $\kappa \approx m/(\sigma\sqrt{\delta})$, the stationary-crossing approximation predicts a larger TWAP tail when $\sigma\sqrt{\delta}/m \gtrsim 1/\kappa^* \approx 0.89$. □

5. Materials and Methods

5.1. Data Sources

The analysis uses data obtained through public endpoints. Binance USD-M futures [42] provide hourly and one-minute prices for BTCUSDT, ETHUSDT and SOLUSDT, as well as native funding histories and current depth snapshots. DeFiLlama [43] provides aggregate stablecoin circulating-supply observations, and Deribit [44] provides public BTC and ETH option summaries. The primary futures snapshot was accessed on 14 June 2026 UTC and a later sensitivity snapshot on 4 July 2026 UTC. Table 2 reports the primary futures panel and Table 3 the depth descriptors. The stablecoin and option observations are retained as contextual evidence and are not inputs to the first-passage estimates.

Table 2. Public-data panel summary statistics; funding-rate statistics in basis points per 8-hour event. Funding Hurst and AR(1) are computed on the native 8-hourly funding series (not forward-filled).

Asset	Hourly Obs	Start	End	Mean Hourly Ret.	Hourly Vol	Funding Mean (bp)	Funding q05 (bp)	Funding q95 (bp)	Funding $\widehat{H}$	Funding AR(1)
BTC	1000	3 May 2026	14 June 2026	−0.0002083	0.004448	0.27	−0.46	1.00	0.9363	0.7126
ETH	1000	3 May 2026	14 June 2026	−0.0003419	0.005825	0.23	−0.67	0.94	0.9323	0.6992
SOL	1000	3 May 2026	14 June 2026	−0.0002243	0.006687	−0.24	−1.97	1.00	0.8197	0.5429

The hourly panel contains 1000 observations per asset, from 3 May to 14 June 2026. BTC and ETH have positive mean funding in the window, while SOL has negative mean funding, consistent with asset-specific positioning. Funding persistence is reported on the native 8-hourly series (Section 6), not the forward-filled hourly grid.

Table 3. Current Binance USD-M order-book depth descriptors by basis-point band, from the public L2 snapshot (1000 price levels per side). "Capped" marks bands that exhaust the level cap, whose depth is therefore a *lower bound* (BTC at 50 bp); the 10 bp band is fully covered for all three assets and is the band used for cross-asset comparisons. The 100 bp band is omitted because it is capped for both majors.

Asset	Band (bps)	Mid	Depth (USD m)	Imbalance	Bid Levels	Ask Levels	Capped
BTCUSDT	10	64,272	25.7	−0.06059	537	553	–
BTCUSDT	50	64,272	63.8	−0.001978	1000	1000	yes
ETHUSDT	10	1,665	20.5	−0.01446	167	167	–
ETHUSDT	50	1,665	111.6	0.1408	817	809	–
SOLUSDT	10	68	4.4	−0.151	7	7	–
SOLUSDT	50	68	19.6	−0.03336	34	34	–

The depth table is a current snapshot, not a historical panel; it is a public descriptor of liquidity scale, and it carries one mechanical limitation for cross-asset comparison: the public endpoint returns at most 1000 price levels per side, and for BTC the 50-basis-point band exhausts that cap, so BTC's 50-bp figure is a *lower bound* (Table 3 flags capped cells; the 100-bp band is omitted because it is capped for both majors). Within the 10-bp band, which the snapshot covers fully for all three assets, the majors' books ($25.7M for BTC, $20.5M for ETH) are five to six times deeper than SOL's ($4.4M). At 50 bp, SOL's $19.6M is *at most* a third of BTC's capped $63.8M. This order-of-magnitude gap, not the BTC–ETH ordering (which the cap leaves undetermined), is what drives the cross-asset impact and bad-debt differences below.

5.2. Measurement Design

Returns, realised volatility, funding distributions, native-series persistence, stale-price errors and depth are measured directly. Empirical volatility is then substituted into the fixed-delay overshoot (14), bad-debt probability (15) and TWAP stationary-crossing approximation (19). The first-passage simulation uses observed volatility and funding to evaluate the specified mechanism; it is not an estimate from historical liquidation records. Order-book depth enters separately through the impact function (6).

Stale-mark errors are computed from observed one-minute prices. Path-level simulation validates the fixed-delay formulas and evaluates the TWAP stationary-crossing approximation (Appendix B). First-passage breach probabilities are calibrated simulations because public endpoints do not provide complete user-level liquidation trajectories. The supporting data, code, generated tables and figures, diagnostics, and provenance files are provided in Supplementary File S1.

The central empirical design follows the rows of Table 1. Hourly price and native funding histories calibrate and score the boundary-breach event; one-minute prices and mechanism simulation validate conditional mark-delay severity; and the depth snapshot maps a stated liquidation size to terminal displacement and average execution concession. Stablecoin supply and option summaries are reported later under an explicitly contextual heading and do not enter any of these estimators, simulations, or forecast scores.

5.2.1. Hourly Volatility and Drift Calibration

For asset i, let

$$r_{i,t} = \log P_{i,t} - \log P_{i,t-1}, \qquad \widehat{\mu}_{h,i} = \frac{1}{n_i}\sum_{t=1}^{n_i} r_{i,t}, \qquad \widehat{\sigma}_{h,i} = \left[\frac{1}{n_i - 1}\sum_{t=1}^{n_i}(r_{i,t} - \widehat{\mu}_{h,i})^2\right]^{1/2}. \tag{20}$$

The 1000 hourly closing prices produce $n_i = 999$ non-missing log returns, so the variance denominator is 998. This Bessel-corrected sample standard deviation gives $\widehat{\sigma}_h = 0.004448$, 0.005825, and 0.006687 for BTC, ETH, and SOL, respectively, as reported in Table 2. The principal simulation draws $r_{i,t}^{\mathrm{sim}} \sim N(0, \widehat{\sigma}_{h,i}^2)$ to isolate boundary, funding, and volatility effects and because the sample means are not statistically different from zero at conventional two-sided levels ($|t| \leq 1.9$). This is a transparent benchmark, not a claim that drift is economically irrelevant. The observed-drift sensitivity instead draws $N(\widehat{\mu}_{h,i}, \widehat{\sigma}_{h,i}^2)$ and materially raises several one-week breach probabilities; it is therefore reported as a stress result and discussed as a limitation of the zero-drift benchmark.

5.2.2. Funding Persistence

We estimate the R/S Hurst exponent on the native 8-hourly funding series (200 observations per asset) and report (i) a moving-block-bootstrap 95% confidence interval, (ii) an i.i.d.-shuffle surrogate baseline that quantifies the estimator's small-sample upward bias, (iii) the DFA exponent, and (iv) the AR(1) coefficient. Estimating on a forward-filled hourly grid would repeat each 8-hourly value eight times and manufacture persistence, which is why every estimate here uses the native series.

5.2.3. Stale-Mark and Catch-Up Experiment

From one-minute prices, we compute the empirical stale error $|X_t - X_{t-\delta}|$ on the delay grid $\delta \in \{1, 5, 15, 30, 60\}$ minutes; from the minute-scale volatility σ, we evaluate the fixed-delay objects of (14) and (15) and the TWAP approximation (19). A mechanism simulation liquidates minute-scale paths on the delayed mark ($\sigma_\delta = \tau + \delta$) and on the first downcrossing of the boundary by the trailing TWAP. It validates the fixed-delay formulas

and measures the combined first-passage and discrete-monitoring error of the TWAP approximation (Appendix B).

5.2.4. First-Passage Simulation

For each asset and leverage $L \in \{3, 5, 10, 20, 50\}$, we simulate $N = 20{,}000$ one-week paths. Price increments are mean-zero Gaussian at the estimated hourly volatility (20); the observed-drift alternative is reported separately. At each hour, equity and maintenance are evaluated from the contemporaneous price ratio $\rho_t = P_t/P_0$:

$$\frac{E_t}{C_0} = 1 + L(\rho_t - 1) - L\sum_{u \le t} \rho_u f_u, \qquad \frac{M_t^{\text{maint}}}{C_0} = mL\rho_t,$$

where f_u is the hourly allocation of the 8-hour funding rate. This implements the notional funding payment in Section 3 rather than subtracting funding from log return. The AR(1), fractional-Gaussian-noise and i.i.d. funding models have matched marginal variance and share price draws. Funding is centred to isolate serial dependence; the observed mean is small relative to the weekly price move. Wilson intervals report Monte Carlo binomial error conditional on the estimated inputs and do not include calibration uncertainty.

5.3. Forecast Target and Rolling Comparison

At forecast origin t, horizon h and leverage L, the binary forecast target is the first passage of a hypothetical isolated-margin long along the subsequent path,

$$Y_{t,h,L} = \mathbf{1}\left\{\min_{1 \le j \le h}\left[1 + L(\rho_{t,j} - 1) - L\sum_{k=1}^{j} \rho_{t,k}\bar{f}_{t+k} - mL\rho_{t,j}\right] \le 0\right\}, \qquad \rho_{t,j} = \frac{P_{t+j}}{P_t}, \quad (21)$$

where $\bar{f}_{t+k}$ is the hourly allocation of the contemporaneous eight-hour funding rate. The structural first-passage simulation supplies $\pi_{t,h,L} = \mathbb{P}(Y_{t,h,L} = 1 \mid \mathcal{I}_t)$, and hence, the Bernoulli predictive distribution $\mathbb{P}(Y_{t,h,L} = y \mid \mathcal{I}_t) = \pi_{t,h,L}^{y}(1 - \pi_{t,h,L})^{1-y}$. Conditional on a mark-detected crossing, Theorems 1 and 2 add the continuous predictive distribution of catch-up severity. Thus, the full forecast is a joint event–severity distribution, although the public-data comparison can score only the event component.

Forecast origins are one day apart. At each origin, the structural model uses the preceding 336 hourly observations to estimate volatility and all native funding observations then available to estimate the funding scale and AR(1) coefficient. The benchmark is historical simulation: it computes the same margin-boundary event over every overlapping return–funding path in the preceding 336 h, with a Jeffreys correction for probabilities of zero or one. We evaluate 24-, 72- and 168-hour horizons at 10×, 20× and 50×. Brier and logarithmic scores are reported; lower values are better. The realised $Y_{t,h,L}$ is constructed from the subsequently observed public price and funding path. It is a model-consistent hypothetical-account outcome, not an exchange record of a customer liquidation. The conditional severity distribution is instead checked through the path-level mechanism simulations because public endpoints do not supply complete mark histories or account trajectories.

6. Results

6.1. Prices and Funding

Figure 2 plots normalised public futures prices over the window, which includes substantial movement but no single crisis event. Figure 3 shows funding dynamics. BTC and ETH carry positive mean funding and SOL negative; fragility is therefore direction-dependent, and a risk engine should model $q_t F_t$, not F_t.

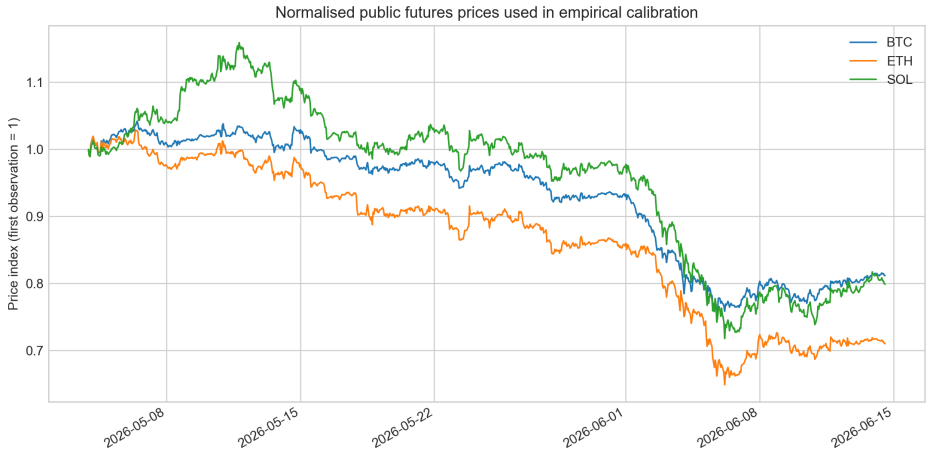


Figure 2. Normalised Binance USD-M futures prices used in the empirical calibration, from public hourly kline data.

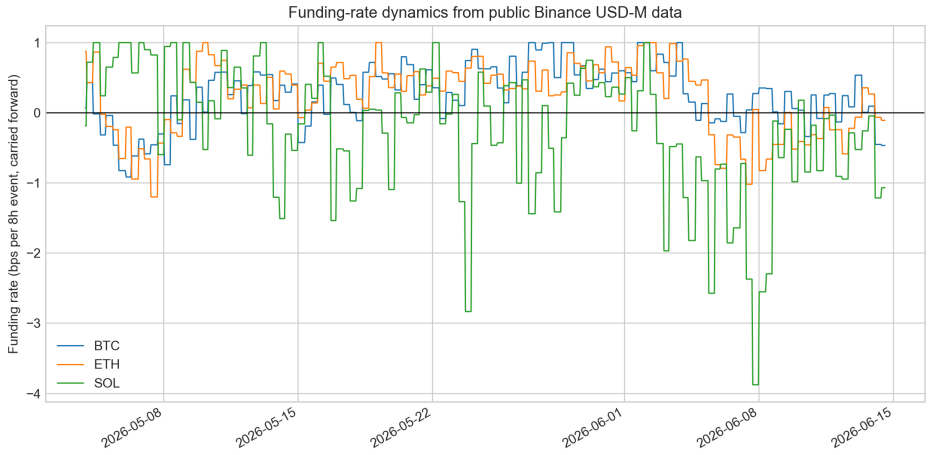


Figure 3. Funding-rate dynamics from public Binance USD-M funding data, in basis points per 8-hour event carried forward to the hourly grid for display.

6.2. Funding-Rate Persistence on Native Data

Table 4 reports persistence on the native 8-hourly funding series. The R/S Hurst exponents (0.94 for BTC, 0.93 for ETH, 0.82 for SOL) lie 4.8, 4.9 and 3.1 surrogate standard deviations above the i.i.d.-shuffle baseline; AR(1) coefficients are 0.71, 0.70 and 0.54. The block-bootstrap intervals are wide at $n = 200$, and the point estimates are method-sensitive, so only positive serial dependence relative to the shuffle benchmark is inferred. The raw R/S level is not interpreted as a fractional parameter. Runs at the venue's 0.01% default rate contribute genuine serial dependence but may reflect the funding formula's regime stickiness.

Table 4. Funding-rate persistence on native 8-hourly data: R/S Hurst with moving-block-bootstrap 95% CI, i.i.d.-shuffle surrogate baseline (mean and s.d.), DFA exponent, AR(1) coefficient, and the standardised distance of the Hurst estimate from the i.i.d. surrogate null.

Asset	Obs.	$\widehat{H}$ (R/S)	$\widehat{H}$ CI Low	$\widehat{H}$ CI High	i.i.d. Surrogate	Surrogate s.d.	DFA α	AR(1)	z vs. Null
BTC	200	0.9363	0.7522	1.01	0.6218	0.06575	1.059	0.7126	4.784
ETH	200	0.9323	0.6984	0.9665	0.6112	0.06588	1.049	0.6992	4.874
SOL	200	0.8197	0.6447	0.8967	0.6155	0.06687	0.9113	0.5429	3.054

6.3. The Cost of Oracle Delay

Table 5 reports the stale-mark and catch-up objects. The sample p95 stale error increases with delay for every asset, consistent with Lemma 1. In the fixed-delay mechanism simulation, realised mean overshoot matches (14) to within 1.3%, and resolved bad-debt frequencies match (15) to within 3%. The Gaussian benchmark rises from negligible at one minute to 0.017 (BTC), 0.041 (ETH) and 0.135 (SOL) at one hour. The cross-asset ordering follows $\sigma\sqrt{\delta}/m$, which reaches 0.91 for SOL against 0.47 for BTC. The delay grid is a stress range; the one-hour cell represents a prolonged feed impairment, not a claim about normal venue operation.

Table 5. Oracle-delay stress from one-minute futures prices. The fixed-delay bad-debt probability is set by $\sigma\sqrt{\delta}/m$ at $m = 0.5\%$; $P_{\mathrm{bd}}^{\mathrm{TWAP}}$ is the stationary-downcrossing approximation for a uniform TWAP mark. p95 stale error and overshoot are in log-return units.

Asset	Delay (min)	p95 Stale Err.	$\sigma\sqrt{\delta}/m$	$\mathbb{E}[O(\delta)]$	P_{bd}	$P_{\mathrm{bd}}^{\mathrm{TWAP}}$	Obs.
BTCUSDT	1	0.0005669	0.06107	0.0001218	$<10^{-6}$	$<10^{-6}$	999
BTCUSDT	5	0.001145	0.1366	0.0002724	$<10^{-6}$	$<10^{-6}$	995
BTCUSDT	15	0.00203	0.2365	0.0004718	1.1×10^{-5}	$<10^{-6}$	985
BTCUSDT	30	0.002901	0.3345	0.0006672	0.001	1.2×10^{-6}	970
BTCUSDT	60	0.003582	0.473	0.0009436	0.017	0.001	940
ETHUSDT	1	0.0006753	0.07457	0.0001487	$<10^{-6}$	$<10^{-6}$	999
ETHUSDT	5	0.00143	0.1667	0.0003326	$<10^{-6}$	$<10^{-6}$	995
ETHUSDT	15	0.002601	0.2888	0.0005761	2.6×10^{-4}	$<10^{-6}$	985
ETHUSDT	30	0.003783	0.4084	0.0008147	0.007	1.0×10^{-4}	970
ETHUSDT	60	0.00437	0.5776	0.001152	0.041	0.009	940
SOLUSDT	1	0.001158	0.1173	0.000234	$<10^{-6}$	$<10^{-6}$	999
SOLUSDT	5	0.002242	0.2623	0.0005232	6.6×10^{-5}	$<10^{-6}$	995
SOLUSDT	15	0.004075	0.4543	0.0009062	0.014	5.8×10^{-4}	985
SOLUSDT	30	0.006065	0.6424	0.001281	0.059	0.022	970
SOLUSDT	60	0.008306	0.9086	0.001812	0.135	0.139	940

Figure 4 plots the fixed-delay bad-debt probabilities in Table 5 across the delay grid.

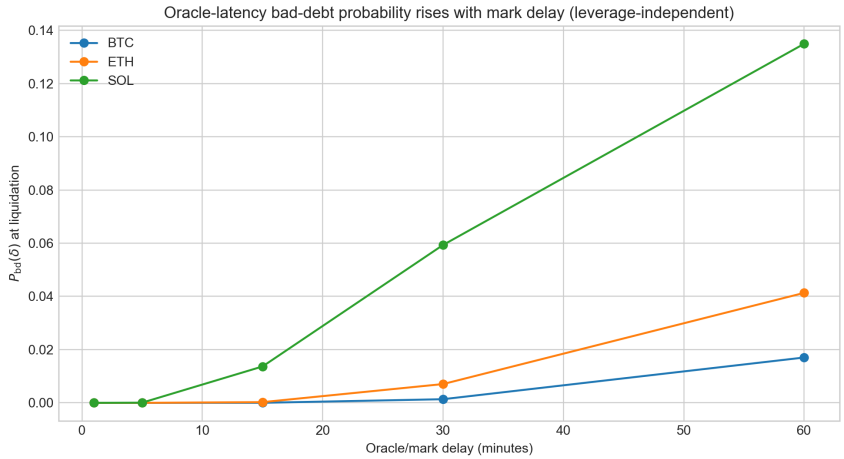


Figure 4. Oracle-latency bad-debt probability $P_{\mathrm{bd}}(\delta)$, the constant-maintenance, fixed-non-price-drain benchmark (15) evaluated at empirical one-minute volatility. Leverage independence applies only within that benchmark. The ordering tracks the latency–margin ratio $\sigma\sqrt{\delta}/m$.

Fixed Delay Versus TWAP Mark

The last column of Table 5 reports the stationary-crossing TWAP approximation (19), and Figure 5 compares it with the naive unconditional $\delta/3$ value and the mechanism Monte

Carlo. At the SOL one-hour cell, the realised bad-debt frequency at the TWAP liquidation time is 0.164, against 0.028 from the unconditional substitution. The stationary-crossing approximation tracks mean overshoot: realised TWAP overshoot is within 5% of $c_O\sigma\sqrt{\delta}$ in every cell and is 1.6–1.7 times the fixed-delay overshoot. Realised bad-debt frequency is 15–40% above the approximation in populated cells. This gap combines the no-earlier-crossing condition with discrete boundary overshoot, so the approximation and mechanism estimate are reported separately. At one hour, the mechanism comparison lowers bad-debt probability for BTC and ETH but raises it for SOL: 0.164 under TWAP against 0.136 under a same-window fixed delay. The stationary approximation predicts the same sign change near the latency–margin ratio 0.89.

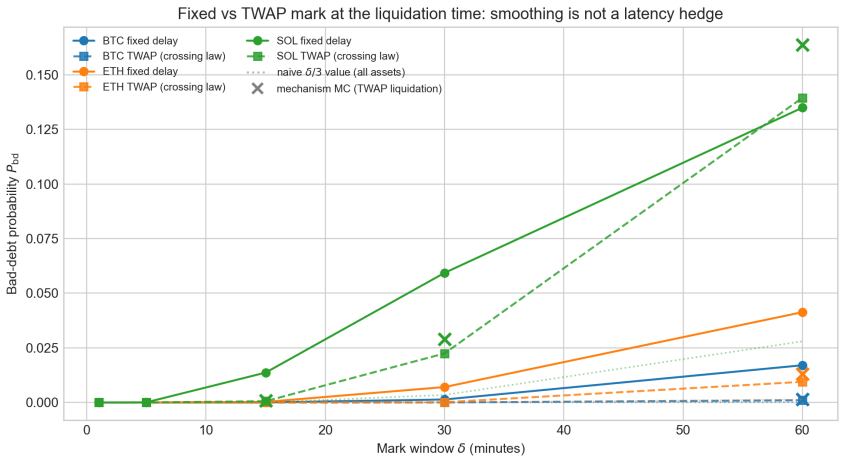


Figure 5. Fixed-delay (solid) versus uniform-TWAP bad-debt probability against mark window δ at empirical one-minute volatility: the stationary-crossing approximation (19) (dashed), the naive unconditional $\delta/3$ value (dotted), and the path-level mechanism Monte Carlo at the TWAP liquidation time (crosses, where resolvable).

6.4. First-Passage Breach Probabilities

Table 6 reports the calibrated one-week breach probabilities with Wilson intervals. They vary across the leverage grid: essentially zero at 3–5×, 0.07–0.22 at 10×, 0.39–0.56 at 20×, and 0.76–0.83 at 50×, increasing monotonically in leverage and in volatility (SOL > ETH > BTC). Mean time-to-liquidation falls from roughly 80 h at 20× to 30 at 50×. Figure 6 shows $\widehat{\Pi}(L)$ with confidence bands. The boundary algebra is material: a mis-specified boundary—for instance, cross-margin algebra applied to an isolated-margin account—saturates these probabilities at 0.87–1.0 across the grid and removes the comparative statics, whereas the isolated-margin boundary (8) evaluated on the empirical hourly volatility over the one-week horizon preserves leverage and volatility variation. We hold $m = 0.5\%$ flat across leverages; because the boundary distance $d_L = (1/L - m)/(1 - m)$ shrinks as m rises, a real tiered schedule (which raises m at high leverage) would liquidate earlier and push breach probabilities up—at 20×, moving m from 0.5% to 2% raises the BTC breach probability from 0.39 to 0.56 in the margin-robustness output. The flat-m figures are a conservative lower bound on breach frequency, while the higher residual equity at a larger m simultaneously lowers bad-debt severity: the two risks again move oppositely (Corollary 2).

Table 6. Calibrated one-week first-passage liquidation probability $\widehat{\Pi}_i(L)$ with Wilson 95% intervals, by asset and leverage. Boundaries use isolated-margin algebra $d_L = (1/L - m)/(1 - m)$ at $m = 0.5\%$, showing one-sided long liquidation. $\widehat{\sigma}_h$ is the Bessel-corrected sample standard deviation of 999 hourly log returns (variance denominator 998); the principal simulation sets drift to zero and reports observed-drift sensitivity separately.

Asset	Leverage	σ_h	$\widehat{\Pi}$	CI Low	CI High	Mean Hrs to Liq.
BTCUSDT	3	0.004448	0	0	1.9×10^{-4}	–
BTCUSDT	5	0.004448	1.0×10^{-4}	2.7×10^{-5}	3.6×10^{-4}	133.5
BTCUSDT	10	0.004448	0.071	0.068	0.075	118.9
BTCUSDT	20	0.004448	0.392	0.385	0.398	79.4
BTCUSDT	50	0.004448	0.757	0.751	0.763	35.4
ETHUSDT	3	0.005825	0	0	1.9×10^{-4}	–
ETHUSDT	5	0.005825	0.004	0.003	0.005	140.6
ETHUSDT	10	0.005825	0.166	0.161	0.171	105.4
ETHUSDT	20	0.005825	0.506	0.499	0.512	66.1
ETHUSDT	50	0.005825	0.807	0.801	0.812	28.9
SOLUSDT	3	0.006687	0	0	1.9×10^{-4}	–
SOLUSDT	5	0.006687	0.012	0.010	0.013	137.3
SOLUSDT	10	0.006687	0.224	0.219	0.230	98.0
SOLUSDT	20	0.006687	0.559	0.552	0.566	59.9
SOLUSDT	50	0.006687	0.826	0.821	0.832	26.1

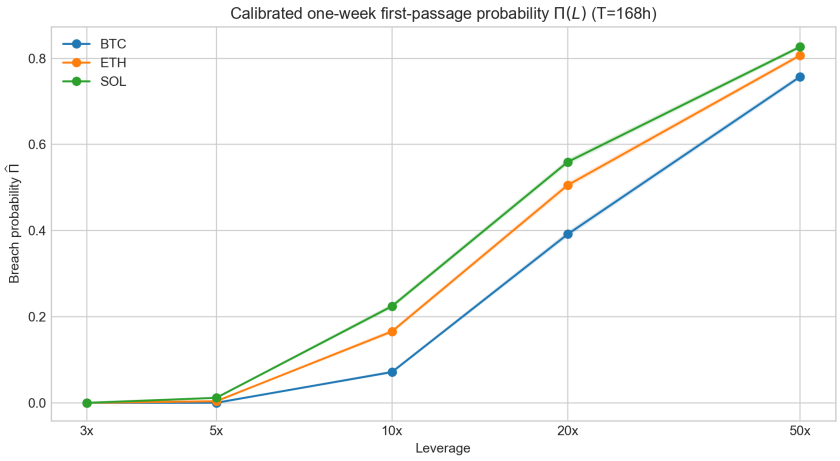


Figure 6. Calibrated one-week first-passage probability $\widehat{\Pi}(L)$ by asset, with Wilson 95% bands, using empirical hourly volatility and the isolated-margin boundary (8).

6.5. Out-of-Sample First-Passage Forecasts

Table 7 reports the rolling comparison with historical simulation. The structural forecast has a lower Brier score at all three horizons for 10× positions, including 0.377 versus 0.457 over one week, and at every leverage over the 24-hour horizon. Its log score is also lower at 24 h for 50×, at 72 h for 10×, and at one week for 10×. These results show that the first-passage mapping can add forecast value when the margin boundary is not close to saturation.

The comparison is not uniformly favourable. At 20× and 50× over one week, the historical path benchmark has lower Brier and log scores. In this sample, those configurations realise breaches in 51 and 60 of 63 overlapping forecasts, respectively; the short rolling history responds more directly to the sustained drawdown than the stationary Gaussian price specification. The result is useful diagnostic evidence rather than a claim of universal superiority. It identifies a clear extension: a production implementation should replace the

stationary innovation law with a conditionally heteroskedastic or regime-sensitive price density while retaining the same margin, funding and mark-severity map.

Table 7. Rolling out-of-sample forecasts of the model-consistent isolated-margin first-passage event. FP is the structural first-passage forecast, and HS is historical simulation of trailing joint return–funding paths; lower Brier and log scores are better. The forecast origin uses the prior 336 h, origins are 24 h apart, and the historical probability has a Jeffreys correction. Negative FP–HS values favour the structural forecast. Outcomes are realised hypothetical-account breaches on subsequent public paths, not observed account liquidations.

Horizon (h)	Leverage	Forecasts	Events	FP Brier	HS Brier	FP–HS	FP Log Score	HS Log Score
24	10×	81	1	0.012	0.013	−0.000	0.091	0.088
24	20×	81	11	0.125	0.126	−0.001	0.473	0.452
24	50×	81	41	0.270	0.283	−0.014	0.734	0.763
72	10×	75	15	0.195	0.224	−0.029	1.115	1.132
72	20×	75	30	0.271	0.270	+0.001	0.768	0.737
72	50×	75	59	0.192	0.169	+0.023	0.575	0.519
168	10×	63	28	0.377	0.457	−0.080	1.248	2.154
168	20×	63	51	0.325	0.183	+0.142	0.852	0.644
168	50×	63	60	0.091	0.060	+0.031	0.353	0.327

6.6. Funding Memory Is Regime-Dependent

Table 8 compares persistent funding with a matched-variance i.i.d. counterfactual using paired paths. At observed funding, every difference is within two paired standard errors of zero. At the ×20 stress, the AR(1) differences are +0.002 to +0.005 (1.3–2.7 standard errors); at ×60, they are +0.010 to +0.020 (3.2–6.9 standard errors). The fractional specification gives +0.023 for BTC at ×60 but only +0.005 and +0.004 for ETH and SOL, both within 1.5 standard errors of zero. Serial dependence is therefore material in some stressed calibrations, but its magnitude and statistical resolution are specification-dependent.

Table 8. Funding-memory premium at 20×: incremental one-week first-passage probability under persistent funding relative to the matched-variance i.i.d. null, across funding-stress regimes. "AR(1) premium" uses the short-memory AR(1) drain actually simulated; "fBm premium" uses fractional Gaussian noise at the native Hurst index (genuine long memory). All drains share the price draws (common random numbers), so each premium is a paired difference; paired standard errors in parentheses.

Asset	Funding Regime	$\widehat{\phi}$	Π i.i.d.	AR(1) Premium (s.e.)	fBm Premium (s.e.)
BTCUSDT	calm (observed)	0.7126	0.3907	+0.0009 (0.0005)	+0.0011 (0.0005)
BTCUSDT	crowded ×20	0.7126	0.3926	+0.0048 (0.0018)	+0.0056 (0.0021)
BTCUSDT	extreme ×60	0.7126	0.4037	+0.0204 (0.0030)	+0.0226 (0.0034)
ETHUSDT	calm (observed)	0.6992	0.506	−0.0005 (0.0004)	−0.0004 (0.0005)
ETHUSDT	crowded ×20	0.6992	0.5044	+0.0022 (0.0017)	+0.0032 (0.0020)
ETHUSDT	extreme ×60	0.6992	0.511	+0.0131 (0.0029)	+0.0045 (0.0032)
SOLUSDT	calm (observed)	0.5429	0.5591	+0.0003 (0.0004)	+0.0003 (0.0004)
SOLUSDT	crowded ×20	0.5429	0.5585	+0.0034 (0.0018)	+0.0022 (0.0020)
SOLUSDT	extreme ×60	0.5429	0.5669	+0.0095 (0.0029)	+0.0036 (0.0032)

Two variance statements are distinct. The fractional *scaling exponent* of Proposition 2, $2H + 2$, is validated directly in Remark 2 and Figure 1 (fitted 2.99/3.36/3.79 versus 3.0/3.4/3.8). Distinct from this exponent, the AR(1) drain used in the simulation inflates

the *level* of one-week cumulative-drain variance by the factor $(1+\widehat{\phi})/(1-\widehat{\phi}) = 5.9$ at $\widehat{\phi} = 0.71$, which is a level multiplier and not a change in horizon scaling. Funding memory therefore inflates the variance of collateral depletion several-fold and, under long memory, steepens its horizon scaling, but it translates into breach probability only when funding itself is large, as in crowded trades.

6.7. Convex Depth Amplification

Figure 7 plots the terminal book displacement (5) against depth utilisation z/D for linear ($\gamma = 1$) and convex ($\gamma = 0.5$) books. For a single account liquidating a small fraction of the book ($z/D \leq 0.1$), displacement is modest, and the more convex book reaches a nearer terminal level. At $z/D = 2$, terminal displacement is 200 bp for $\gamma = 0.5$ versus 100 bp for $\gamma = 1$. The corresponding *average* execution concessions from (6) are 66.7 bp and 50 bp, respectively; the terminal and average objects must not be interchanged. Because SOL's 50-bp book is at most a third of BTC's (Table 3, the BTC figure is a level-capped lower bound, so the true gap is at least this large), the same dollar liquidation reaches a higher z/D on SOL. This raises terminal displacement and, in sufficiently stressed utilisation ranges, the average and log-equivalent concessions that enter Theorem 1.

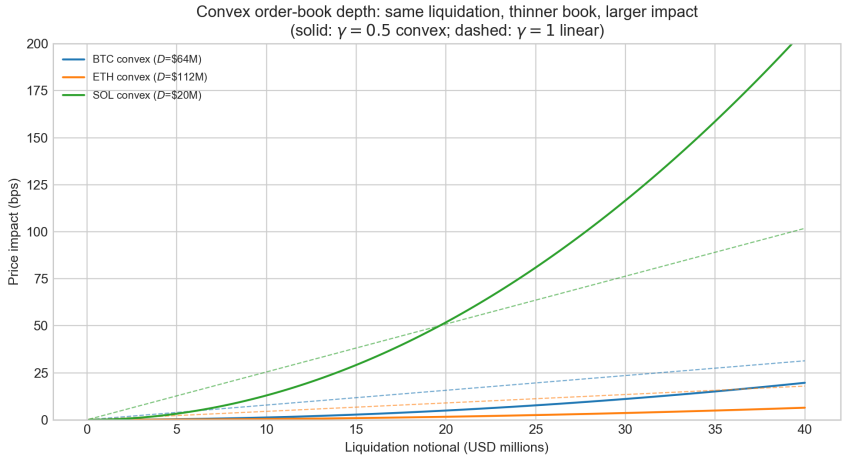


Figure 7. Convex order-book depth amplification: terminal book displacement $y_\gamma(z;D) = b(z/D)^{1/\gamma}$ for linear and convex books. The average execution concession is $[\gamma/(1+\gamma)]y_\gamma$ and is therefore smaller than the plotted endpoint.

6.8. Contextual Evidence Outside the Central Calibration

This subsection is deliberately separated from the tests of the liquidation forecast. Neither stablecoin supply nor the option cross-section enters the boundary, the funding process, the latency formulas, the execution mapping, or the rolling scores. Retaining them records the wider public information set and motivates future state variables, but no headline result depends on them.

The liquidation model is calibrated from futures prices, funding and depth, but those variables operate within a wider digital-asset funding environment. Figure 8 reports aggregate stablecoin supply from the public liquidity series. Stablecoins facilitate settlement and often serve as collateral or quote assets, so their aggregate availability provides a market-state descriptor relevant to funding liquidity and the capacity to meet margin calls [31,32]. The series cannot reveal which assets are held by individual futures traders, whether they are eligible collateral at a particular venue, or whether a supply change causes liquidation risk. It is consequently not included in the boundary, persistence estimates or breach simulations.

The public option cross-section contains 1684 BTC and ETH contract summaries. These observations preserve a forward-looking volatility layer that is conceptually relevant to the latency–margin ratio $\sigma\sqrt{\delta}/m$. The present calibration deliberately uses realised minute volatility so that the stale-price measurements and analytical inputs refer to the same observed price process. Replacing or combining that input with maturity- and strike-specific implied volatility would require a separate term-structure treatment and is left as an extension. Keeping the option snapshot in the empirical record clarifies both the available information set and the boundary of the present claims.

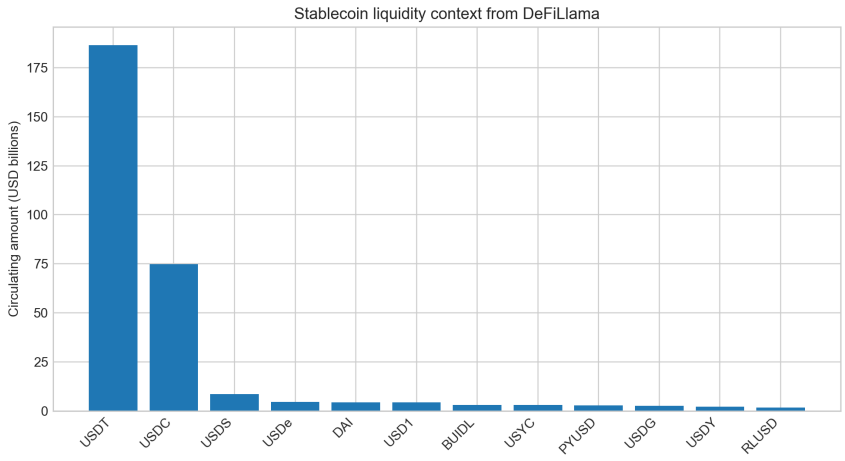


Figure 8. Aggregate stablecoin liquidity from public circulating-supply observations. The series describes market-wide conditions and is not interpreted as user-level collateral or liquidation evidence.

7. Discussion

The theory and evidence separate boundary crossing from conditional loss severity. *Leverage* moves the breach boundary (Corollary 2). *Oracle delay* governs conditional bad-debt severity through the $\sqrt{\delta}$ overshoot and the ratio $\sigma\sqrt{\delta}/m$ (Theorem 1); leverage still changes how often liquidation occurs. *Funding dependence* changes cumulative-drain variance (Proposition 2) and affects breach probability mainly when funding is large. *Depth and book shape* determine terminal displacement and the average execution concession; the ordering in book convexity is utilisation-dependent (Proposition 3).

Taken together, these mechanisms answer the single research question rather than four separate ones. The event forecast begins with a leverage-specific margin boundary and updates its crossing probability using price and funding inputs. Conditional on a crossing, the observation rule supplies a hidden-price-loss distribution. Conditional on execution size and observed depth, the average execution concession shifts that severity distribution. The public evidence is matched to this order: hourly paths score the event, minute paths validate the latency channel, and the book snapshot scales the execution channel. The conclusions below are limited to those linked objects; the contextual stablecoin and option series do not widen the causal claim.

Economic Interpretation of the Stopping Mechanism

Detection probability and loss severity respond differently to delay. Under a pure lag, a continuous executable-price crossing is recorded after the lag has elapsed. Equation (12) therefore permits the probability of a *recorded* liquidation by a fixed calendar horizon to fall as δ increases. Conditional on subsequent detection, however, the executable price contains an additional increment over the hidden interval. Its expected adverse part is proportional to $\sqrt{\delta}$, and the corresponding bad-debt probability depends on the maintenance cushion

through $\sigma\sqrt{\delta/m}$. Counts truncated at a fixed horizon can thus fall because detection is postponed, even when conditional loss increases.

For a TWAP mark, averaging reduces stale-error variance at a fixed time, but liquidation samples the error when the smoothed process downcrosses its boundary. This conditioning changes the distribution of the executable-minus-mark error. The Rayleigh mixture in Theorem 2 is the stationary-downcrossing law induced by that conditioning. It is not an exact law for a discretely monitored venue. It supplies a benchmark for the direction and scale of crossing selection, while the path-level experiment measures the error created by first passage and discrete monitoring. The comparison shows why fixed-time variance can understate liquidation-time error.

Serially dependent funding changes cumulative collateral drain before the boundary is reached. Its effect depends on persistence, funding scale, position direction and horizon; persistence alone does not imply a large change in breach probability. Depth affects execution loss once liquidation consumes available liquidity. Terminal displacement and average concession are distinct; for the average concession, a smaller γ is not automatically more costly at $z/D = 1$, and the relevant ordering must be checked over the actual utilisation range using (6). The persistence estimate, depth snapshot and delay formulas therefore measure different mechanisms and should not be combined as proxies for an unobserved common fragility index.

Aggregate stablecoin supply describes the scale of a settlement and collateral medium but cannot be mapped to an account's surplus. The option cross-section contains forward-looking price information, although its use in $\sigma\sqrt{\delta/m}$ would require a maturity- and strike-consistent volatility model. Accordingly, futures prices and funding determine the calibrated state variables, the depth snapshot describes execution capacity at its observation time, and simulation evaluates the stated mechanism. The stablecoin and option series provide market context but no additional causal identification.

The rolling comparison further separates forecast content from claim scope. The model produces a probability distribution for the stated margin mechanism and a conditional severity distribution for delayed detection. Its score is informative at moderate leverage and short horizons, but the one-week high-leverage result shows that a historical path distribution can be more accurate in a persistent drawdown. This is a limitation of the current price-density specification, not evidence that public data identify historical customer liquidations.

8. Limitations and Robustness

8.1. Observability

Public endpoints do not report trader leverage, collateral balances, liquidation paths, insurance-fund transfers or historical mark construction. The first-passage probabilities are therefore calibrated mechanism estimates, not venue-level liquidation frequencies. Stablecoin supply and option summaries provide market context but do not identify account collateral or liquidation events.

8.2. Sampling

Depth is observed once and may decline under stress. The 1000-level endpoint cap makes BTC's 50-bp depth a lower bound and leaves the BTC–ETH ordering at that band unidentified; cross-asset comparisons therefore use the fully covered 10-bp band in Table 3. Funding histories are also short, so persistence is interpreted from the R/S gap relative to the shuffle baseline, supported by DFA and AR(1), rather than as a precisely estimated fractional parameter.

8.3. Model Assumptions

A pure delay omits venue-specific index constituents, clamps and fallback rules. The Gaussian benchmark omits jumps, skewness and tail dependence, while the breach simulation uses zero drift because the sample drift is not statistically different from zero ($|t| \leq 1.9$). As a stress check, imposing the observed negative drift raises the 10× breach probabilities from 0.07/0.17/0.22 to 0.19/0.40/0.36 for BTC/ETH/SOL.

The exact detection shift (12) and fixed-delay tail (15) additionally freeze the funding boundary, fees, position size, and maintenance tier during the detection interval. If those quantities evolve, liquidation is governed by the time-varying boundary (11); incremental funding changes the bad-debt threshold according to (16). These formulas identify the isolated price-latency channel and should not be read as an exact venue rule under continuously evolving funding.

8.4. Forecast Evaluation

The rolling score comparison uses only 63–81 forecast origins per leverage–horizon cell and realised horizons overlap. It is therefore a descriptive out-of-sample test, not a large-sample forecast-encompassing test. The observed path is sufficient to score the hypothetical isolated-margin event in (21), but not to score venue-specific mark detection, user collateral, liquidation fees or insurance-fund losses. The high-leverage weekly comparison also rejects any claim of uniform dominance over historical simulation.

8.5. Later-Snapshot Sensitivity

The later hourly window overlaps the primary window by about 21 days, and its longer funding history largely contains the primary funding period. Table 9 is therefore a rolling-window sensitivity check, not an independent holdout. In the later snapshot, 10× breach probabilities rise to 0.12–0.30, the cross-asset ordering is unchanged, and the ETH and SOL TWAP bad-debt frequencies exceed their fixed-delay values.

Table 9. Later-snapshot sensitivity. W1 is the primary window (3 May to 14 June 2026, accessed 14 June 2026); W2 is the later hourly window (24 May to 4 July 2026, accessed 4 July 2026). The hourly windows overlap by about 21 days, and W2 funding uses a longer 500-event history, so this is a rolling-window and expanded-history sensitivity check rather than an independent holdout. Breach probabilities and bad-debt objects use each snapshot's volatility calibration.

Object	BTC W1	BTC W2	ETH W1	ETH W2	SOL W1	SOL W2
Hourly volatility σ_h	0.0044	0.0052	0.0058	0.0068	0.0067	0.0079
Minute volatility $\sigma_{\min}$	0.00031	0.00035	0.00037	0.00065	0.00059	0.00072
Funding Hurst $\widehat{H}$ (R/S)	0.94	0.92	0.93	0.81	0.82	0.80
z vs. i.i.d. surrogate	4.8	7.4	4.9	4.7	3.1	5.0
Funding AR(1)	0.71	0.63	0.70	0.52	0.54	0.59
Share at default rate	0.035	0.024	0.040	0.024	0.105	0.066
$\widehat{\Pi}$ (one week, 10×)	0.071	0.123	0.166	0.231	0.224	0.303
$\widehat{\Pi}$ (one week, 20×)	0.392	0.464	0.506	0.568	0.559	0.620
$\sigma\sqrt{60}/m$	0.47	0.54	0.58	1.01	0.91	1.11
P_{bd} (fixed, 60 min)	0.017	0.032	0.041	0.160	0.135	0.184
$P_{\text{bd}}^{\text{TWAP}}$ (60 min)	0.001	0.005	0.009	0.196	0.139	0.257
10-bp depth (USD m)	25.7	29.6	20.5	13.9	4.4	4.0

The reported objects remain separable: stale error, persistence, overshoot, bad-debt probability, breach probability and depth impact.

8.6. Additional Stress Scenarios

A joint volatility–depth stress sets $\sigma_s = s\sigma$ and $d_s = d/s^{\rho}$, $s > 1$. Theorem 1 then raises P_{bd} through σ, while (6) raises execution loss through lower depth. Directional funding stress uses $q_t F_t$ to distinguish long-crowded and short-crowded positions. Maintenance-margin sensitivity replaces $m = 0.5\%$ with venue tiers; the latency–margin rule is then evaluated tier by tier rather than called leverage-independent. The remaining checks vary $\gamma \in [0.25, 1]$ and set $T \in \{72, 168, 336\}$ hours.

9. Risk-Management Implications

Risk controls can be stated in terms of the model's three measured quantities. The stale displacement $D_t(\delta)$ and ratio $\sigma\sqrt{\delta}/m$ determine the fixed-delay benchmark for conditional bad debt. Funding dependence determines whether cumulative drain differs materially from an independent-funding benchmark. Liquidation notional and depth determine the terminal displacement, average concession, and log-equivalent execution loss through (5)–(7); the book-shape ordering is utilization-dependent, as stated in Proposition 3.

The operational recommendations are straightforward:

1. Faster price detection or a larger maintenance cushion reduces the loss that can be hidden before liquidation;
2. Leverage limits should be used to control how often the boundary is reached, while latency controls address loss after the crossing;
3. Funding stress should be assessed in the direction of the position and over the forecast horizon, rather than summarized by an unconditional average rate;
4. Liquidation size should be compared with executable depth using the average execution price, not only the last book level consumed.

These recommendations require exchanges and protocols to disclose the mark rule, maximum fallback delay, funding rule, maintenance schedule, liquidation fees, and usable depth assumptions.

The model motivates state-contingent margin schedules through a dynamic maintenance fraction

$$m_t = m_0 + a_1\widehat{\sigma}_t + a_2|\widehat{F}_t| + a_3\,\widehat{\sigma}_t\sqrt{\delta} + a_4\,\widehat{\phi}'_t(N_t), \tag{22}$$

in which the delay term enters through $\sigma\sqrt{\delta}$, the quantity Theorem 1 identifies. For decentralised protocols, the oracle design is often posed as a two-risk trade-off: a slower, smoother oracle reduces manipulation [3,35] but raises catch-up risk. Theorem 2 formalises that trade-off because the relevant object is not the smoothing window's unconditional dispersion but the stale error *at the mark's own downcrossing of the boundary*: a TWAP of window δ carries about 1.6 times the expected catch-up overshoot of a same-window fixed delay, and in the regime $\sigma\sqrt{\delta}/m \gtrsim 0.89$, it enlarges the bad-debt tail as well (Corollary 3). The design lever is the *crossing-conditioned* overshoot of the kernel, so the smoothing window trades manipulation cost against $c_O\sigma\sqrt{\delta}$ rather than against the naive $\delta/3$ equivalence, which understates the latency cost of smoothing by nearly a factor of three. For regulators, the relevant disclosure set is the mark formula, maximum delay, smoothing window, funding caps, liquidation fees and insurance-fund rules, without which users cannot evaluate Π or P_{bd} even approximately.

10. Conclusions

This study shows that liquidation risk in crypto perpetual futures depends not only on leverage and volatility but also on how the exchange observes prices. A delayed mark records liquidation after the executable price has crossed the margin boundary. Under the fixed-delay benchmark, the expected catch-up overshoot is $\sigma\sqrt{\delta/2\pi}$, while conditional

bad-debt probability is governed by the latency–margin ratio $\sigma\sqrt{\delta}/m$. Leverage changes the frequency of boundary crossings but, under the model assumptions, not the conditional bad-debt probability once liquidation occurs.

Price smoothing does not remove this latency cost. Under a TWAP mark, liquidation selects adverse stale-price errors at boundary downcrossings. The stationary-downcrossing approximation implies a mean overshoot about 1.6 times the same-window fixed-delay value. Path-level simulation confirms the overshoot scale but also shows that the approximate TWAP bad-debt probability should not be treated as an exact first-passage result.

The public-data calibration finds positive funding dependence relative to a shuffle benchmark, leverage-sensitive breach probabilities and higher conditional bad-debt risk as delay increases. The associated forecast is a distribution over a margin-boundary event and its conditional catch-up severity. In rolling public-path evaluation, the structural probability forecast improves on historical simulation at moderate leverage and short horizons but not uniformly at one-week high leverage. Funding dependence becomes economically material mainly under stressed funding conditions. These quantities are calibrated stress measures rather than estimates of historical liquidation frequencies.

The practical implication is that margin design, mark construction and executable depth should be assessed jointly. Exchanges should disclose mark formulas, smoothing windows and fallback delays, and should evaluate maintenance margins against crossing-conditioned stale-price risk rather than fixed-time mark variance alone.

Future Research

The next empirical step is to replace the stationary Gaussian innovation law with a conditionally heteroskedastic, jump-aware, or regime-sensitive predictive density and to evaluate it on longer, non-overlapping forecast samples. Access to timestamped historical marks and anonymised account-level liquidation paths would permit direct scoring of the detection and severity components that public endpoints cannot currently identify.

A priority for mark design is a volume-weighted-average-price (VWAP) extension. Because volume and depth are endogenous to stress, a VWAP mark requires a joint price–volume crossing law rather than substituting a deterministic set of weights into the TWAP formula. This extension would connect the observation kernel directly to available liquidity while preserving the distinction between terminal book displacement and average execution price.

Further theory should also allow the funding boundary $b_L(g_t)$, tiered maintenance $m(L)$, and execution concession to evolve during the detection interval, so that the benchmark correction in (16) can be evaluated without freezing non-price drains.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/forecast8050076/s1, Supplementary File S1, a reproducibility archive containing public-data extracts, the analysis scripts `reproduce_analysis.py` and `second_window_robustness.py`, generated tables and figures, diagnostics, and provenance files.

Author Contributions: Conceptualisation, E.P.; methodology, E.P.; software, E.P.; formal analysis, E.P.; investigation, E.P.; data curation, E.P.; visualisation, E.P.; validation, E.P. and H.P.M.; writing—original draft preparation, E.P.; writing—review and editing, E.P. and H.P.M.; supervision, H.P.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The public source endpoints are Binance USD-M Futures Market Data (https://developers.binance.com/docs/derivatives/usds-margined-futures/market-data/rest-api, accessed on 14 June 2026), DeFiLlama Stablecoins API (https://stablecoins.llama.fi/stablecoins?includePrices=true, accessed on 14 June 2026) and the Deribit Public API (https://docs.deribit.com/, accessed on 14 June 2026). No private account, order or liquidation feed is used. The primary snapshot was accessed on 14 June 2026 UTC and a later snapshot on 4 July 2026 UTC. The supporting computational routines, processed datasets, figures and tables are supplied as Supplementary Materials.

Acknowledgments: Grammarly (https://www.grammarly.com) was used for language editing.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Extended Mathematical Notes

Appendix A.1. General Mark Operators

Beyond a fixed delay, let the log-mark be a kernel average of the recent log price, $\log M_t = \int_0^\infty K(a)\, X_{t-a}\, \mathrm{d}a + \eta_t$ with $\int K = 1$; a fixed delay is a Dirac mass at δ, a moving-average mark a uniform or exponential kernel. Two objects must be kept apart. The *unconditional* stale-error dispersion defines a variance-equivalent delay $\delta_{\mathrm{eff}} = \mathrm{Var}(D_t)/\sigma^2$ ($\delta/3$ for the uniform TWAP by Lemma 2, $\tau/2$ for an exponential kernel of mean lag τ), and it is the right scale for fixed-time questions such as manipulation cost. The *liquidation-time* cost is a different object: any absolutely continuous kernel makes $\log M_t$ differentiable, so its boundary downcrossings sample the local slope with Kac–Rice size-biasing, and the stale error at the crossing is biased adverse exactly as in Theorem 2; only the pure-delay Dirac kernel escapes, through the strong Markov property. For the exponential kernel, $\mathrm{d}\log M_t/\mathrm{d}t = (X_t - \log M_t)/\tau = D_t/\tau$, so the crossing slope is proportional to the stale error itself, the size-biased crossing law is the pure Rayleigh $D_{\tau_M} \overset{d}{=} -\sigma\sqrt{\tau/2}\, R$ with no Gaussian residual, and

$$\mathbb{E}[O^{\mathrm{exp}}] = \sqrt{\tfrac{\pi}{2}}\,\sigma\sqrt{\tau/2} \approx 0.886\,\sigma\sqrt{\tau}, \qquad P_{\mathrm{bd}}^{\mathrm{exp}} = \mathbb{P}\big(R > \kappa\sqrt{2}\big) = \exp\Big(-\frac{\log^2(1-m)}{\sigma^2\tau}\Big),$$

about 2.2 times the overshoot of a fixed delay at the same mean lag τ under the same stationary-crossing approximation. The uniform and exponential examples show why fixed-time dispersion need not describe liquidation-time error; they do not establish the ordering for every smoothing kernel.

Appendix A.2. VWAP Extension

A trailing VWAP would take

$$\log M_t^{\mathrm{VWAP}} = \frac{\int_{t-\delta}^{t} V_u X_u\, \mathrm{d}u}{\int_{t-\delta}^{t} V_u\, \mathrm{d}u},$$

with transaction-volume or executable-depth weight V_u. Unlike the deterministic TWAP kernel, these weights are stochastic and can covary with both returns and the crossing slope. A valid liquidation-time law must therefore condition jointly on price, volume, and the VWAP downcrossing; replacing K by average historical volume weights would miss the liquidity selection that motivates the extension. Estimating that joint crossing law is left for future work.

Appendix A.3. Depth and Impact

For $L(y) = ay^\gamma$, $0 < \gamma \leq 1$, executing notional z reaches terminal displacement $y(z) = (z/a)^{1/\gamma}$, and the total concession integrated across the consumed book is

$$\Phi(z) = \int_0^z (u/a)^{1/\gamma} \, \mathrm{d}u = \frac{\gamma}{1+\gamma} a^{-1/\gamma} z^{1+1/\gamma}, \tag{A1}$$

strictly convex for $\gamma < 1$. Dividing by z gives the average fractional concession $\Phi(z)/z$. Anchoring a to band depth $D = ab^\gamma$ yields terminal displacement (5), average execution concession (6), and log-equivalent concession (7). The bad-debt calculation uses the last of these, not the terminal book level.

Appendix B. Computational Validation Protocol

The validation uses deterministic simulation settings and public input data. Monte Carlo objects carry Wilson intervals. Minute-scale paths are liquidated on delayed and TWAP marks ($\approx 10^5$ liquidations per asset–window cell; see Table A1). Fixed-delay overshoot is within 1.3% of the formula, and bad-debt frequency is within 3% in resolved cells. For TWAP marks, mean overshoot is within 5% of the stationary-crossing approximation, while bad-debt frequency is 15–40% higher in stressed cells; this discrepancy combines first-passage selection and discrete monitoring. Separate distributional checks cover the half-normal overshoot mean, discrete-TWAP dispersion, and the T^{2H+2} scaling in Proposition 2.

Table A1. Path-level mechanism check at liquidation: minute-scale driftless paths at empirical minute volatility are liquidated on a delayed mark and on the first discrete downcrossing by a trailing uniform TWAP. Realised outcomes are compared with the fixed-delay formula, the stationary-crossing TWAP approximation, and the naive unconditional $\delta/3$ value. About 10^5 liquidations per cell.

Asset	δ (min)	$\mathbb{E}[O]$ MC	$\mathbb{E}[O]$ Thm	P_{bd} MC	P_{bd} Thm	$\mathbb{E}[O^{\mathrm{TW}}]$ MC	$\mathbb{E}[O^{\mathrm{TW}}]$ Thm	$P_{\mathrm{bd}}^{\mathrm{TW}}$ MC	$P_{\mathrm{bd}}^{\mathrm{TW}}$ Thm	$P_{\mathrm{bd}}^{\mathrm{TW}}$ Naive
BTCUSDT	15	0.000471	0.0004718	0	1.1×10^{-5}	0.0007818	0.0007538	0	$<10^{-6}$	$<10^{-6}$
BTCUSDT	30	0.0006622	0.0006672	0.001	0.001	0.00111	0.001066	0	1.2×10^{-6}	$<10^{-6}$
BTCUSDT	60	0.0009317	0.0009436	0.017	0.017	0.001569	0.001508	0.001	0.001	1.2×10^{-4}
ETHUSDT	15	0.0005783	0.0005761	4.0×10^{-4}	2.6×10^{-4}	0.000955	0.0009204	0	$<10^{-6}$	$<10^{-6}$
ETHUSDT	30	0.0008157	0.0008147	0.007	0.007	0.001354	0.001302	2.1×10^{-4}	1.0×10^{-4}	1.1×10^{-5}
ETHUSDT	60	0.001153	0.001152	0.042	0.041	0.001919	0.001841	0.013	0.009	0.001
SOLUSDT	15	0.0009054	0.0009062	0.014	0.014	0.001506	0.001448	8.4×10^{-4}	5.8×10^{-4}	6.6×10^{-5}
SOLUSDT	30	0.001285	0.001281	0.060	0.059	0.002129	0.002047	0.029	0.022	0.003
SOLUSDT	60	0.001827	0.001812	0.136	0.135	0.00303	0.002895	0.164	0.139	0.028

Appendix C. Variable Dictionary

Table A2. Variable dictionary for the model notation.

Symbol	Meaning
X_t, M_t	Executable log price; delayed mark/oracle price.
$D_t(\delta)$	Stale-mark error induced by delay δ.
F_t	Funding state; positive values drain longs, benefit shorts.
$y_\gamma(z;D)$, $\bar{\phi}_\gamma(z;D)$, $\ell_\gamma(z;D)$	Terminal book displacement; average fractional execution concession; log-equivalent concession.
b_L, d_L, $\xi(L)$	Liquidation log boundary; simple-return drop; residual equity (8) and (9).
$O(\delta)$, P_{bd}	Catch-up overshoot and bad-debt probability (14) and (15).
δ_{eff}	Variance-equivalent delay of a smoothing kernel (unconditional dispersion) (17).

Table A2. *Cont.*

Symbol	Meaning
R, c_O, $P_{\text{bd}}^{\text{TWAP}}$	Rayleigh crossing variable, TWAP overshoot constant $c_O = 0.637$, and crossing-conditioned TWAP bad-debt probability (18) and (19).
$\Pi(T; m, \delta, H, \Phi)$	First-passage (breach) probability over horizon T.
$\mathbb{P}\{\tau_m^{\text{price}}(\delta) \le T\}$	Price-channel recorded-liquidation probability, shifted by delay only under the fixed-boundary conditions of (12).

References

1. Ackerer, D.; Hugonnier, J.; Jermann, U. *Perpetual Futures Pricing*; Mathematical Finance, Forthcoming (NBER Working Paper 32936); National Bureau of Economic Research: Cambridge, MA, USA, 2025. [CrossRef]
2. He, S.; Manela, A.; Ross, O.; von Wachter, V. Fundamentals of perpetual futures. *arXiv* **2024**, arXiv:2212.06888. [CrossRef]
3. Mackinga, T.; Nadahalli, T.; Wattenhofer, R. TWAP oracle attacks: Easier done than said? In Proceedings of the 2022 IEEE International Conference on Blockchain and Cryptocurrency (ICBC), Shanghai, China, 2–5 May 2022. [CrossRef]
4. Donier, J.; Bonart, J. A million metaorder analysis of market impact on the Bitcoin. *Mark. Microstruct. Liq.* **2015**, *1*, 1550008. [CrossRef]
5. Merton, R.C. On the pricing of corporate debt: The risk structure of interest rates. *J. Financ.* **1974**, *29*, 449–470. [CrossRef]
6. Black, F.; Cox, J.C. Valuing corporate securities: Some effects of bond indenture provisions. *J. Financ.* **1976**, *31*, 351–367. [CrossRef]
7. Leland, H.E. Corporate debt value, bond covenants, and optimal capital structure. *J. Financ.* **1994**, *49*, 1213–1252. [CrossRef]
8. Longstaff, F.A.; Schwartz, E.S. A simple approach to valuing risky fixed and floating rate debt. *J. Financ.* **1995**, *50*, 789–819. [CrossRef]
9. Almgren, R.; Chriss, N. Optimal execution of portfolio transactions. *J. Risk* **2001**, *3*, 5–39. [CrossRef]
10. Avellaneda, M.; Stoikov, S. High-frequency trading in a limit order book. *Quant. Financ.* **2008**, *8*, 217–224. [CrossRef]
11. Cartea, A.; Jaimungal, S.; Penalva, J. *Algorithmic and High-Frequency Trading*; Cambridge University Press: Cambridge, UK, 2015. Available online: https://www.scirp.org/reference/referencespapers?referenceid=4197823 (accessed on 4 July 2026).
12. Cont, R. Empirical properties of asset returns: Stylized facts and statistical issues. *Quant. Financ.* **2001**, *1*, 223–236. [CrossRef]
13. Hasbrouck, J. Measuring the information content of stock trades. *J. Financ.* **1991**, *46*, 179–207. [CrossRef]
14. Kyle, A.S. Continuous auctions and insider trading. *Econometrica* **1985**, *53*, 1315–1335. [CrossRef]
15. Makarov, I.; Schoar, A. Trading and arbitrage in cryptocurrency markets. *J. Financ. Econ.* **2020**, *135*, 293–319. [CrossRef]
16. Brauneis, A.; Mestel, R.; Riordan, R.; Theissen, E. How to measure the liquidity of cryptocurrency markets? *J. Bank. Financ.* **2021**, *124*, 106041. [CrossRef]
17. Mandelbrot, B.B.; Van Ness, J.W. Fractional Brownian motions, fractional noises and applications. *SIAM Rev.* **1968**, *10*, 422–437. [CrossRef]
18. Cont, R. Long range dependence in financial markets. In *Fractals in Engineering*; Springer: London, UK, 2005; pp. 159–179. [CrossRef]
19. Gatheral, J.; Jaisson, T.; Rosenbaum, M. Volatility is rough. *Quant. Financ.* **2018**, *18*, 933–949. [CrossRef]
20. Takaishi, T. Rough volatility of Bitcoin. *Financ. Res. Lett.* **2020**, *32*, 101379. [CrossRef]
21. Cheridito, P.; Kawaguchi, H.; Maejima, M. Fractional Ornstein–Uhlenbeck processes. *Electron. J. Probab.* **2003**, *8*, 1–14. [CrossRef]
22. Hu, Y.; Nualart, D. Parameter estimation for fractional Ornstein–Uhlenbeck processes. *Stat. Probab. Lett.* **2010**, *80*, 1030–1038. [CrossRef]
23. Black, F. The pricing of commodity contracts. *J. Financ. Econ.* **1976**, *3*, 167–179. [CrossRef]
24. Brunnermeier, M.K.; Pedersen, L.H. Market liquidity and funding liquidity. *Rev. Financ. Stud.* **2009**, *22*, 2201–2238. [CrossRef]
25. Qin, K.; Zhou, L.; Gamito, P.; Jovanovic, P.; Gervais, A. An empirical study of DeFi liquidations: Incentives, risks, and instabilities. In *IMC '21: Proceedings of the 21st ACM Internet Measurement Conference (IMC)*; Association for Computing Machinery: New York, NY, USA, 2021; pp. 336–350. [CrossRef]
26. Werner, S.; Perez, D.; Gudgeon, L.; Klages-Mundt, A.; Harz, D.; Knottenbelt, W.J. SoK: Decentralized finance (DeFi). In *AFT '22: Proceedings of the 4th ACM Conference on Advances in Financial Technologies*; Association for Computing Machinery: New York, NY, USA, 2022; pp. 30–46. [CrossRef]
27. Lehar, A.; Parlour, C.A. *Systemic Fragility in Decentralized Markets*; BIS Working Paper 1062; Bank for International Settlements: Basel, Switzerland, 2022. Available online: https://www.bis.org/publ/work1062.htm (accessed on 4 July 2026).
28. Klages-Mundt, A.; Minca, A. While stability lasts: A stochastic model of noncustodial stablecoins. *Math. Financ.* **2022**, *32*, 943–981. [CrossRef]

29. Gudgeon, L.; Werner, S.; Perez, D.; Knottenbelt, W.J. DeFi protocols for loanable funds: Interest rates, liquidity and market efficiency. In *AFT '20: Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*; Association for Computing Machinery: New York, NY, USA, 2020; pp. 92–112. [CrossRef]
30. Schär, F. Decentralized finance: On blockchain- and smart contract-based financial markets. *Fed. Reserve Bank St. Louis Rev.* **2021**, *103*, 153–174. [CrossRef]
31. Anadu, K.; Azar, P.; Cipriani, M.; Eisenbach, T.M.; Huang, C.; Landoni, M.; La Spada, G.; Macchiavelli, M.; Malfroy-Camine, A.; Wang, J.C. *Runs and Flights to Safety: Are Stablecoins the New Money Market Funds?* Federal Reserve Bank of New York Staff Report 1073; Federal Reserve Bank of New York: New York, NY, USA, 2023. [CrossRef]
32. Gorton, G.B.; Klee, E.C.; Ross, C.P.; Ross, S.Y.; Vardoulakis, A.P. *Leverage and Stablecoin Pegs*; NBER Working Paper 30796; National Bureau of Economic Research: Cambridge, MA, USA, 2022. [CrossRef]
33. Liu, Y.; Tsyvinski, A. Risks and returns of cryptocurrency. *Rev. Financ. Stud.* **2021**, *34*, 2689–2727. [CrossRef]
34. Griffin, J.M.; Shams, A. Is Bitcoin really untethered? *J. Financ.* **2020**, *75*, 1913–1964. [CrossRef]
35. Qin, K.; Zhou, L.; Livshits, B.; Gervais, A. Attacking the DeFi ecosystem with flash loans for fun and profit. *Financ. Cryptogr. Data Secur.* **2021**, *12674*, 3–32. [CrossRef]
36. Gneiting, T.; Raftery, A.E. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378. [CrossRef]
37. Diebold, F.X.; Gunther, T.A.; Tay, A.S. Evaluating density forecasts with applications to financial risk management. *Int. Econ. Rev.* **1998**, *39*, 863–883. [CrossRef]
38. Cheng, Z.; Deng, J.; Wang, T.; Yu, M. Liquidation, leverage and optimal margin in bitcoin futures markets. *Appl. Econ.* **2021**, *53*, 5415–5428. [CrossRef]
39. Kim, J.; Park, H. Designing funding rates for perpetual futures in cryptocurrency markets. *arXiv* **2025**, arXiv:2506.08573. [CrossRef]
40. Donnelly, R.; Lin, J.; Lorig, M. Optimal liquidation of perpetual contracts. *arXiv* **2026**, arXiv:2601.10812. [CrossRef]
41. Burnashev, M.V.; Tchamkerten, A. Estimating a random walk first-passage time from noisy or delayed observations. *IEEE Trans. Inf. Theory* **2012**, *58*, 4230–4243. [CrossRef]
42. Binance. Binance Public USD-M Futures API Documentation and Market-Data Endpoints. 2026. Available online: https://developers.binance.com/docs/derivatives/usds-margined-futures/market-data/rest-api (accessed on 14 June 2026).
43. DeFiLlama. Stablecoins Public API. 2026. Available online: https://stablecoins.llama.fi/stablecoins?includePrices=true (accessed on 14 June 2026).
44. Deribit. Public Option Book Summary API. 2026. Available online: https://docs.deribit.com/#public-get_book_summary_by_currency (accessed on 14 June 2026).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.