FinTech

MDPI

Article

A Fractional-Rough Liquidity Model for Bitcoin Options: Implied-Volatility Asymptotics and Market Evidence

Edson Pindza * and Hopolang Phillip Mashele

Department of Decision Sciences, College of Economics and Management Sciences, University of South Africa, Pretoria 0002, South Africa; mashehp@unisa.ac.za
* Correspondence: edsonpindza@gmail.com

Abstract

Bitcoin option prices reflect terminal variance and the cost of managing convex exposure in a market with changing depth and execution quality. This paper asks whether a liquidity state can be separated from fractional rough volatility in Bitcoin option valuation. The contribution is a modelling combination: standard stochastic-calculus and rough-volatility tools are joined to a regime-switching hedging-cost reserve, producing a leading-order at-the-money implied-volatility lift. The empirical design tests a liquidity–IV association and its scale using a 700-contract Deribit snapshot, a 4513-trade 24-h window spanning two UTC dates, a 775,315-trade panel over 92 dates, Ether replication, and placebos. The association is strong in open-interest-weighted specifications and for puts, but is absent for calls; it remains after controlling for option premium. Leave-one-expiry-out validation improves open-interest-weighted RMSE but not unweighted RMSE. A realised-volatility HMM is only a market-stress diagnostic, not an estimated liquidity regime. An empirical one-step hedging exercise does not validate the model's simulated hedging comparative static. Accordingly, the evidence is associational, put-side, and narrower than a causal or fully structural validation.

Keywords: bitcoin options; cryptocurrency derivatives; digital assets; rough volatility; fractional Brownian motion; liquidity premium; regime switching; Monte Carlo pricing; hidden Markov model; Deribit

JEL Classification: G12; G13; C58

Check for updates

Academic Editors: Juraj Fabuš and Shyan-Ming Yuan

Received: 9 July 2026
Revised: 10 August 2026
Accepted: 21 August 2026
Published: 23 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.

1. Introduction

Bitcoin options are used by funds, market makers, miners, structured-product desks, and retail participants to express volatility views and manage nonlinear exposure. Their market structure, however, differs from the assumptions under which the canonical option-pricing literature was developed. A Bitcoin option may trade in a market where volatility changes within minutes, market depth varies rapidly, funding conditions in related perpetual futures alter hedging economics, and fragmented venues create execution costs that are not well represented by a single continuous bid–ask spread. These features are relevant to valuation because the option is not only a claim on a terminal underlying value: it is also a claim whose replication, hedging, liquidation, and marking depend on the state of crypto-market liquidity.

Classical option-pricing models begin with no-arbitrage reasoning. The Black–Scholes–Merton framework [1,2] demonstrates how dynamic replication can price a

European option under continuous trading, lognormal diffusion, constant volatility, and frictionless markets. The Heston model [3] relaxes constant volatility through a square-root stochastic-variance process. Jump-diffusion and Lévy models incorporate discontinuous price movements [4–6]. Rough-volatility models then respond to empirical evidence that volatility is much less regular than classical semimartingale intuition suggests [7–12]. These models provide the mathematical baseline for the present paper. Yet a direct transfer to Bitcoin options leaves a structural gap: the liquidity state is usually outside the valuation state or appears only as an ex-post transaction-cost adjustment, despite empirical evidence that crypto-derivatives liquidity has its own time variation [13–17].

Organizing Contribution

The paper's contribution is methodological: it gives a pricing-state decomposition in which rough volatility governs the terminal payoff distribution, while a regime-switching liquidity state enters the hedging-cost component. This narrower framing is important. The paper is not presented as a full structural calibration of all Bitcoin-option risks, nor as a causal study of liquidity shocks. Instead, it asks whether the liquidity state has a separately measurable implied-volatility footprint once the usual moneyness, maturity, option-type and market-state controls are included.

The model produces two empirically testable implications and one internal comparative static. First, residual illiquidity should be positively associated with implied volatility after local surface controls. Second, a low-to-stressed liquidity move should generate an at-the-money implied-volatility lift of the order predicted by the short-time asymptotic,

$$\Delta\sigma_{\mathrm{ATM}}(T) = \alpha_L \, \bar{L} \, \sigma_0 \sqrt{2\pi T} + o(\sqrt{T}),$$

under the proper Riemann discretization, with $\bar{L} = \mathbb{E}[L_t]$ and $\sigma_0 = \sqrt{\xi_0}$. Greater stressed-state persistence raises simulated hedging errors inside the model, but this is treated as an internal comparative static unless it survives an external hedging test.

We combine (i) a filtered 700-contract Deribit Bitcoin option cross-section, (ii) 4513 market-state-matched trades in a 24-h window from 19 May 10:38 UTC to 20 May 10:38 UTC, chosen because it is the complete archived public-trade pull for which the hourly state series can be joined, (iii) 775,315 clean historical Bitcoin option trades over 92 dates, and (iv) Ether replication and placebo checks. The snapshot association is sensitive to economic weighting and option type: the open-interest-weighted estimate is positive, the unweighted estimate is smaller, puts are positive, and calls are null. Moving from the median spread (0.031) to the 90th percentile (0.400) corresponds to 7.26 IV percentage points; the 7.98-point calibration target instead measures a near-zero-to-90th-percentile move. Since uncertainty, dealer inventory, option demand, and volatility can jointly affect spreads and IV, these are conditional projections, not causal effects.

The remainder of the paper is organized as follows. Section 2 positions the contribution against option-pricing, rough-volatility, liquidity, and crypto-derivatives literatures. Section 3 presents the model, its economic predictions, and the numerical implementation. Section 4 explains the data construction and the option-chain evidence. Section 5 reports mechanism checks, panel evidence, replication, and placebo tests. Sections 6–8 discuss findings, limitations, and conclusions.

https://doi.org/10.3390/fintech5030073

2. Related Literature

2.1. Classical Option Pricing, Stochastic Volatility, and Jumps

The classical option-pricing literature establishes the frictionless benchmark that the present paper deliberately departs from. Black and Scholes [1] and Merton [2] show how replication and no-arbitrage lead to a risk-neutral valuation equation when trading is continuous, volatility is constant, and markets are frictionless. Heston [3] relaxes constant volatility through a square-root variance process. Merton [6], Bates [4], and the jump-process treatment of Cont and Tankov [5] add discontinuous price moves, while Carr and Madan [18] develop Fourier pricing tools for affine specifications. These models are useful baselines for Bitcoin options because crypto returns display stochastic volatility and jumps, but they do not by themselves identify whether liquidity is a state variable or merely a transaction-cost afterthought.

2.2. Rough Volatility

The rough-volatility literature supplies the payoff-dynamics side of the model. Gatheral et al. [11] show that volatility appears rough, with estimated Hurst exponents significantly below 1/2. Bayer et al. [7] develop pricing under the rough Bergomi model, El Euch and Rosenbaum [10] study the rough Heston characteristic function, and Bennedsen et al. [9] provide a simulation scheme for Brownian semistationary processes. Jacquier et al. [12], with the published corrigendum [19], analyze small-time asymptotics for rough Bergomi models, and Friz and Hairer [20] give the rough-path background. Roughness matters here because it changes short-maturity implied-volatility behavior; the contribution below is to ask how a separate liquidity state changes the implied-volatility level after rough payoff dynamics are already present.

2.3. Regime Switching and Market States

Hamilton [21] formalizes regime switching for economic time series. Elliott et al. [22] provide hidden Markov model methods, Hardy [23] studies regime-switching long-run stock returns, and Alfeus et al. [24] combine regime switching with rough Heston dynamics. In this paper the theoretical regime is a liquidity state. The empirical HMM is therefore used cautiously: it is a descriptive volatility-state diagnostic that checks whether the observed market history has persistent stress episodes, not a structural proof that volatility regimes and liquidity regimes are identical.

2.4. Liquidity and Market Microstructure

Market microstructure explains why transaction prices differ from frictionless values. Kyle [25] models price impact under informed trading, Glosten and Milgrom [26] link bid–ask spreads to asymmetric information, and Amihud [27] measures illiquidity in asset returns. Çetin et al. [28] provide a no-arbitrage framework with liquidity-cost supply curves; Çetin and Rogers [29] formalize option pricing with proportional liquidity charges; and Gârleanu et al. [30] document demand-pressure effects in equity options. This literature motivates treating liquidity as a priced state because option values depend on the cost of dynamically managing convex exposure, not only on the terminal distribution of the underlying asset.

2.5. Crypto Derivatives and Bitcoin-Market Specificity

The closest literature is the growing body of Bitcoin-option pricing work. Hou et al. [14] price cryptocurrency options with jump-and-volatility risk and show why standard equity-option assumptions require modification in crypto markets. Madan et al. [15] calibrate advanced models to Bitcoin options and emphasize the role of model risk in option-surface

fitting. Zulfiqar and Gulzar [31] document implied-volatility behavior in Bitcoin options, while Scaillet et al. [17] study high-frequency jumps in Bitcoin markets. More recent Deribit studies document option-regime and illiquidity effects [13,16]. Recent work also shows that Bitcoin-option hedges should be evaluated with historical backtests rather than simulation alone [32], and that jumps and option-market demand pressure materially shape Bitcoin implied-volatility dynamics [33,34]. These findings sharpen the gap addressed here: liquidity must be tested conditionally on surface and price-level controls, and a model-internal hedge diagnostic is not external validation.

The most direct comparison is Li [35], who develops an equilibrium Bitcoin option-pricing model with stochastic volatility, correlated jumps, and liquidity risk; derives risk-neutral dynamics and semi-analytic pricing formulas; and calibrates the model to Bitcoin option surfaces and high-frequency liquidity measures. The present paper differs in three ways. First, the volatility channel is fractional rough rather than the equilibrium stochastic-volatility channel in Li [35]. Second, liquidity is represented as a regime-switching hedging-cost state, giving a short-time at-the-money implied-volatility lift that can be mapped directly to a spread move. Third, the empirical exercise is not a full surface calibration horse race; it is a prediction-led validation of the liquidity-IV gradient across a snapshot, trade panels, Ether replication, and placebo tests. Thus the paper extends the related literature by isolating a rough-volatility/liquidity-state decomposition rather than by claiming a complete structural replacement for existing Bitcoin-option models.

3. Model, Predictions, and Numerical Implementation

Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space satisfying the usual conditions.

Definition 1 (Friction-aware Bitcoin option state)**.** *A friction-aware Bitcoin option state is the state process $\mathcal{X}_t = (S_t, V_t, R_t, L_t, J_t)$, where S_t is the Bitcoin spot or index price, V_t is the instantaneous variance, $R_t \in \{0,1\}$ is a liquidity regime, L_t is a continuous illiquidity intensity, and J_t is a jump-counting component.*

Assumption 1 (Rough variance proxy)**.** *For $H \in (0, 1/2)$, the log-variance satisfies $\log V_t = \log \xi_0(t) + \eta B_t^H - \frac{1}{2}\eta^2 t^{2H}$, where B^H is fractional Brownian motion with Hurst exponent H, $\eta > 0$ is volatility-of-volatility, and $\xi_0(t)$ is a forward variance baseline. The drift correction yields $\mathbb{E}[V_t] = \xi_0(t)$.*

Assumption 2 (Liquidity-regime process)**.** *$R_t \in \{0,1\}$ is a Markov chain with transition kernel p_{ij} on a discretized grid, $R_t = 1$ denoting stress.*

Assumption 3 (Liquidity intensity)**.** $L_t = \ell_{R_t} + \chi\,|\Delta B_t^H|/\Delta^H$, $\ell_1 > \ell_0 > 0$, $\chi \geq 0$.

Assumption 4 (Jump component)**.** *Conditional on the regime, jumps arrive as a compound Poisson process with intensity $\lambda(R_t) \geq 0$ and jump size $\mathcal{N}(\mu_J, \sigma_J^2)$.*

The spot/index process used for pricing is a semimartingale with rough variance as a state variable:

$$\frac{dS_t}{S_{t-}} = \mu_t\,dt + \sqrt{V_t}\,dW_t + (e^{Y_t} - 1)\,dN_t,$$

where Y_t is the log-jump size. The fractional Brownian component enters through the variance state and the liquidity intensity, not as a self-financing trading integrator.

Table 1 collects the notation used repeatedly below.

Table 1. Core notation.

Symbol	Meaning	Unit/Role
$S_t, K, T-t$	index, strike, remaining maturity	USD, USD, years
V_t, σ_0, H	variance, baseline volatility, Hurst exponent	annualised variance/volatility, roughness
R_t, L_t	binary liquidity regime, illiquidity intensity	state, dimensionless proxy
C_t^0, P_t^0	frictionless call and put values	USD
Λ_t, α_L	additive hedge reserve, liquidity loading	USD, calibrated loading
h_C, h_P	option-type inventory multipliers	dimensionless
Z_i	$\log(1+\text{relative spread})$ or $\log(1+\text{mark deviation})$	empirical liquidity proxy

3.1. Risk-Neutral Pricing Functional

The model separates two objects that are often collapsed in empirical option work. Rough volatility changes the payoff distribution and the short-maturity shape of the implied-volatility surface. Liquidity changes the cost of keeping a convex payoff hedged when depth and execution quality deteriorate. This distinction is the reason for the additive local pricing term below: it is not a universal bid–ask adjustment, but a first-order state-price correction used near the ATM benchmark and then tested empirically through surface-controlled liquidity gradients.

The market is incomplete because R_t and L_t are not fully hedgeable using the traded spot, futures, perpetual swap and other options. We select an equivalent martingale measure $\mathbb{Q}$ by a minimal-entropy/Esscher-type criterion conditional on the rough variance state. Under $\mathbb{Q}$ the discounted spot $e^{-rt}S_t$ is a true martingale; the regime kernel is preserved; and liquidity risk is priced via a nonnegative state-price loading α_L. The liquidity-adjusted call price is

$$\begin{aligned} C_t^L &= C_t^0 + \Lambda_t(T,K), \\ C_t^0 &= e^{-r(T-t)}\mathbb{E}_t^{\mathbb{Q}}[\Phi(S_T)], \\ \Lambda_t(T,K) &= \alpha_L\, S_t\, \mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T e^{-r(u-t)} L_u \sqrt{V_u}\, du\right]. \end{aligned}$$

The additive liquidity term is used below as a local ATM pricing adjustment. A full-surface implementation can write

$$C_t^{L,h}(T,K) = C_t^0(T,K) + h(T,K,\mathcal{X}_t)\Lambda_t(T,K),$$

where h is a hedge-intensity or surface-sensitivity multiplier depending on moneyness, maturity, option type and inventory demand. The present paper sets $h=1$ for the ATM benchmark and uses the cross-sectional regressions to test how the liquidity-IV gradient changes across the surface. The factor S_t makes the local ATM liquidity premium a dollar price adjustment; without this scale factor the integral is dimensionless and cannot be added to an option price.

The additive form is a first-order hedge reserve rather than a second volatility process. Let a dealer follow a frictionless hedge Δ_u^0 and face the convex execution-cost functional

$$\mathcal{C}(\Delta^0; L, V) = S_t \int_t^T e^{-r(u-t)} L_u\, q_u(\Delta^0, V)\, du,$$

where q_u is local expected turnover or an inventory charge. Around a 30-day ATM calibration point, the first Gateaux derivative of q_u is proportional to $\sqrt{V_u}$; writing its coefficient as $\alpha_L \geq 0$ gives

$$C_t^{\text{dealer}} = C_t^0 + D\mathcal{C}[\Delta^0] + o(\|L\|) = C_t^0 + \alpha_L S_t \mathbb{E}_t^{\mathbb{Q}}\left[\int_t^T e^{-r(u-t)} L_u \sqrt{V_u}\, du\right] + o(\|L\|).$$

Thus Λ is a cash reserve for unspanned execution and inventory costs. Loading liquidity into volatility would change the terminal payoff law; loading it into r would treat the friction as deterministic financing; and subtracting realised replication costs ex post would not produce an ex-ante quote. Those mechanisms may coexist in a richer model, but the additive term here is a local approximation, not a universal full-surface rule.

Calls and puts must satisfy joint arbitrage accounting. Define

$$C_t^{L,h_C} = C_t^0 + h_C \Lambda_t, \qquad P_t^{L,h_P} = P_t^0 + h_P \Lambda_t.$$

Then modified put–call parity is

$$C_t^{L,h_C} - P_t^{L,h_P} = S_t - Ke^{-r(T-t)} + (h_C - h_P)\Lambda_t.$$

The core benchmark sets $h_C = h_P = 1$, so standard parity is preserved. If inventory multipliers differ, the last term belongs to dealer quote reserves inside transaction-cost bands, not to a single frictionless mid-price. Pointwise bounds additionally require $0 \leq \Lambda_t \leq S_t - C_t^0$ for a call and $0 \leq \Lambda_t \leq Ke^{-r(T-t)} - P_t^0$ for a put; the numerical benchmark satisfies them.

Theorem 1 (Existence of the EMM under exponential integrability)**.** *Let (W, B^H) be the driving Brownian motion and fractional Brownian motion, with leverage correlation $\rho \in (-1, 1)$ acting on the Brownian increments of V. Let $\theta_u = (\mu_u - r)/\sigma_u$ denote the candidate market price of risk, with $\sigma_u = \sqrt{V_u}$. Assume that the change of measure tilts the hedgeable Brownian and jump risks but does not tilt the Markov-chain transition kernel for R_t. Let $\mathbb{Q}^0$ denote the Brownian Girsanov measure generated by θ. Assume:*

(i) ***Novikov for*** *W:* $\mathbb{E}^{\mathbb{P}}\left[\exp\left(\frac{1}{2}\int_0^T \theta_u^2\, du\right)\right] < \infty$.

(ii) ***True-martingale condition for the discounted spot:*** *$e^{-rt}S_t$ is uniformly integrable under $\mathbb{Q}^0$ on $[0, T]$; a sufficient condition is a finite exponential moment of $\int_0^T V_u\, du$ of order strictly larger than one, together with the jump moment in (iv).*

(iii) ***Exponential integrability of liquidity exposure:*** *for some $\varepsilon > 0$, $\mathbb{E}^{\mathbb{Q}^0}\left[\exp\left((\alpha_L + \varepsilon)\int_0^T L_u\sqrt{V_u}\, du\right)\right] < \infty$.*

(iv) ***Jump integrability:*** *$\mathbb{E}^{\mathbb{P}}\left[\int_0^T \lambda(R_u)\, du\right] < \infty$ and the jump-size exponential moments required by the Esscher transform are finite.*

Then there exists an equivalent martingale measure $\mathbb{Q} \sim \mathbb{P}$ on $\mathcal{F}_T$ under which: (a) the discounted underlying $e^{-rt}S_t$ is a true $\mathbb{Q}$-martingale; (b) the regime kernel p_{ij} is preserved; (c) $\Lambda_t(T, K)$ is a.s. finite and integrable for every (T, K) in the option universe.

Proof. Condition (i) is the classical Novikov condition; by Karatzas and Shreve [36] (corollary 5.13, Ch. 3), it makes the Doléans-Dade exponential $\mathcal{E}_T(-\int \theta\, dW)$ a true $\mathbb{P}$-martingale, so the change of measure with density $d\mathbb{Q}^0/d\mathbb{P} = \mathcal{E}_T(-\int \theta\, dW)$ is well-defined and equivalent. Under $\mathbb{Q}^0$, Girsanov's theorem [36] (theorem 5.1, Ch. 3) gives $\widetilde{W}_t = W_t + \int_0^t \theta_u\, du$ as a Brownian motion, and $e^{-rt}S_t$ is a local martingale.

The jump component is handled conditionally on the regime path by an Esscher tilt whose parameter is chosen inside the jump-size moment domain so that the discounted jump-diffusion drift is compensated. The Brownian and jump density processes are strictly positive; their product has expectation one under the stated moment conditions, so the resulting measure is equivalent to $\mathbb{P}$ on $\mathcal{F}_T$.

Condition (ii) upgrades the discounted local martingale to a true martingale by uniform integrability. The stated exponential-moment condition on integrated variance is one standard sufficient route: it gives an $L^{1+\delta}$ bound for the stochastic exponential, in-

cluding the compensated jump part, and the de la Vallée-Poussin criterion yields uniform integrability of $\{e^{-rt}S_t\}_{t\le T}$.

The regime chain R_t is defined on a discrete state space and, by assumption, the Radon–Nikodym derivative does not tilt its transition probabilities. The kernel p_{ij} is therefore preserved by $\mathbb{Q}$. The jump component's integrability is preserved by condition (iv) and the Esscher transform conditional on R.

Finiteness of $\Lambda_t(T,K)$ follows from condition (iii) and Jensen's inequality:

$$\mathbb{E}^{\mathbb{Q}}\left[\int_0^T L_u\sqrt{V_u}\,du\right] \le \frac{1}{\alpha_L+\varepsilon}\log\mathbb{E}^{\mathbb{Q}}\left[\exp\left((\alpha_L+\varepsilon)\int_0^T L_u\sqrt{V_u}\,du\right)\right] < \infty.$$

Multiplication by the finite initial scale factor S_t and by α_L therefore preserves integrability of the local ATM liquidity premium.

Finally, in the incomplete market the pricing measure is selected from the admissible EMMs by the minimal-entropy/Esscher criterion [37]. The theorem establishes existence and finiteness of the pricing functional under the stated integrability conditions; uniqueness is not required for the comparative statics below. □

Remark 1. *The boundedness hypothesis of v3 is replaced by the exponential-integrability hypotheses above. Boundedness implies these conditions with arbitrary ε, so Theorem 1 strictly generalizes the v3 result. In the finite discretization used in the simulations, $L_u \sim \ell_{R_u} + \chi|\Delta B^H|/\Delta^H$ has Gaussian-tail increments and the variance is numerically clipped, so the required liquidity-exposure moments are finite.*

3.2. Model-Implied Predictions and Economic Intuition

The following results are deliberately simple because their role is to generate empirical predictions. They say that a higher liquidity state should raise the local liquidity component; that more persistent stress should raise expected liquidity exposure; and that roughness and liquidity need not affect prices in the same direction. Economically, the first two statements are about hedging costs, while the third prevents the empirical analysis from interpreting all implied-volatility movements as rough-volatility movements.

Proposition 1 (Liquidity-adjusted valuation monotonicity)**.** *If $L_u^{(2)} \ge L_u^{(1)}$ a.s. for all $u \in [t,T]$, then $C_t^{L,2} \ge C_t^{L,1}$.*

Proof. C^0 is identical by assumption; the liquidity component is a conditional expectation of a nonnegative pricing kernel against a dominated integrand; monotonicity of conditional expectation gives the result. □

Proposition 2 (Effect of stressed-regime persistence)**.** *In the two-state Markov model, holding the calm-to-stress transition probability p_{01} fixed, an increase in p_{11} weakly increases the expected liquidity premium whenever $\ell_1 > \ell_0$ and the process can enter the stressed state.*

Proof. Couple two regime chains with the same innovations, the same p_{01}, and stressed-state persistence values $p_{11}^{(2)} \ge p_{11}^{(1)}$. Whenever both chains are in state 1, the second chain remains in state 1 for at least as many one-step continuations as the first chain under the monotone coupling. Therefore the expected occupation time $\mathbb{E}[\int_t^T \mathbf{1}_{\{R_u=1\}}du]$ is weakly larger in the second chain. Since the liquidity intensity contains the regime level ℓ_{R_u} with $\ell_1 > \ell_0$, the expected regime-level contribution to $\int_t^T L_u\sqrt{V_u}\,du$ weakly rises, provided the variance factor is nonnegative and integrable. The rough-increment component of L_u is nonnegative, so it cannot reverse the monotonicity of the regime-level term. □

Proposition 3 (Convex capital-charge interpretation)**.** *Let* $X = \alpha_L S_t \int_t^T e^{-r(u-t)} L_u \sqrt{V_u}\, du$ *and* ρ *be a monotone, translation-equivariant convex risk measure. Then* $\Pi(\Phi) = C_t^0 + \rho(X)$ *defines a monotone convex valuation rule; if* ρ *is coherent, the rule is also sub-additive. If* ρ *is Gateaux differentiable at the calibration point,* $\Lambda_t(T, K)$ *is the corresponding first-order linear approximation to* $\rho(X)$*.*

Proof. Monotonicity and convexity follow from the corresponding properties of ρ and the unchanged payoff component. If ρ is additionally coherent, then sub-additivity follows from the coherence axioms. The linear liquidity premium Λ_t is the first-order Taylor approximation to $\rho(X)$ at the calibration point. □

Proposition 4 (Roughness–liquidity interaction)**.** *If* L_t *depends on normalized rough increments* $|\Delta B_t^H|/\Delta^H$*, then changing* H *changes the serial dependence and joint path structure of liquidity exposure even when the one-step marginal distribution of the normalized increment is unchanged. The sign of the total option price effect is ambiguous: rough payoff and liquidity components can move in opposite directions.*

Proof. For fixed Δ, the marginal law of $\Delta B^H/\Delta^H$ is standard normal, but the covariance structure across increments depends on H. Thus H affects the pathwise accumulation of $\int L_u \sqrt{V_u}\, du$ through dependence, clustering, and its covariance with the variance path rather than through a simple one-step marginal scale effect. The terminal payoff depends on the full distribution of S_T, which mixes leverage, jumps, and variance normalization, so the rough component need not be monotone in H. The two effects can therefore move differently. □

These comparative statics translate into the first empirical prediction: after controlling locally for the option surface, contracts or trades with larger residual illiquidity should have higher implied volatility. The next two projection results state the identification condition used in the regressions: the liquidity proxy must retain variation after moneyness, maturity, option type, and fixed effects are removed.

Lemma 1 (Regression coefficient as a projected liquidity price)**.** *Let* σ_i^{obs} *denote observed implied volatility, let* $Z_i = \log(1 + \text{RelSpread}_i)$ *be the liquidity proxy, and let* X_i *collect moneyness, maturity, option type, and expiry controls. If*

$$\sigma_i^{obs} = g(X_i) + \kappa Z_i + u_i, \qquad \mathbb{E}[u_i \mid X_i, Z_i] = 0,$$

then the population coefficient on Z_i *in the linear projection of* σ_i^{obs} *on* (X_i, Z_i) *is*

$$\beta_Z = \kappa + \frac{\text{Cov}\big(\tilde{Z}_i,\, g(X_i) - X_i'\gamma_X\big)}{\text{Var}(\tilde{Z}_i)},$$

where $\tilde{Z}_i$ *is the residual from projecting* Z_i *on* X_i *and* γ_X *is the best linear coefficient on* X_i*. If the surface controls span* $g(X_i)$ *locally, then* $\beta_Z = \kappa$*.*

Proof. Apply the Frisch–Waugh–Lovell theorem. The coefficient on Z_i equals

$$\frac{\text{Cov}(\tilde{Z}_i, \sigma_i^{obs})}{\text{Var}(\tilde{Z}_i)}.$$

Substituting the structural relation and using $\text{Cov}(\tilde{Z}_i, X_i'\gamma_X) = 0$ gives the expression. If $g(X_i)$ is locally spanned by $X_i'\gamma_X$, the second term vanishes. □

Proposition 5 (Local identification of the liquidity channel)**.** *Consider a local option-surface neighborhood in which the frictionless implied-volatility surface is differentiable in* $(\log(K/S), \sqrt{\tau})$ *and can be approximated by* $X_i'\gamma_X$*. If the residualized liquidity proxy* $\tilde{Z}_i$ *has nonzero variance and the liquidity premium shifts implied volatility by* $\kappa Z_i + o(Z_i)$*, then* κ *is locally identified by the coefficient on* Z_i *in the residualized implied-volatility regression.*

Proof. Differentiability gives a first-order local linear expansion of the frictionless surface in the controls. Lemma 1 then implies that the residualized liquidity coefficient equals the first-order liquidity effect κ whenever $\mathrm{Var}(\tilde{Z}_i) > 0$. Nonzero residual liquidity variation is therefore the rank condition for local identification. □

3.3. Closed-Form Small-Time Implied-Volatility Prediction

The short-time ATM approximation is the second model prediction. It is useful because market participants often quote liquidity effects in vol points rather than in dollars, and because ATM vega has a simple leading-order form. The formula therefore gives a direct way to ask whether the spread-implied IV widening is economically plausible, rather than merely statistically significant.

Theorem 2 (Small-time IV lift)**.** *Let the proposed model be discretized over* n_{steps} *equal time intervals on* $[0, T]$ *with* $\Delta = T/n_{steps}$*, and define the discretized liquidity-premium accumulator* $\widehat{\Lambda}_n = \alpha_L S_0 \frac{1}{N} \sum_{i=1}^{N} \sum_{k=1}^{n_{steps}} L_k^{(i)} \sqrt{V_k^{(i)}}\, \Delta$*. Suppose* L_u *and* $\sqrt{V_u}$ *are bounded and that, for some* $\beta > 0$*,* $\mathbb{E}^{\mathbb{Q}}[L_u \sqrt{V_u}] = \bar{L}\sigma_0 + O(u^{\beta})$ *as* $u \downarrow 0$*, where* $\bar{L}$ *is the short-time mean liquidity intensity and* $\sigma_0 := \sqrt{\xi_0}$*. Then as* $T \downarrow 0$*,*

$$\Delta\sigma_{ATM}(T) := \sigma^{L}_{ATM}(T) - \sigma^{0}_{ATM}(T) = \alpha_L\, \bar{L}\, \sigma_0\, \sqrt{2\pi T} + O(T^{1/2+\beta}),$$

where σ^{L}_{ATM} *and* σ^{0}_{ATM} *are the Black–Scholes implied volatilities of* C_t^L *and* C_t^0 *at the ATM strike* $K = S_0 e^{rT}$*.*

Proof. Because σ^0_{ATM} is defined as the implied volatility backed out from the frictionless price C^0, the expansion only needs to linearize the incremental liquidity component around that price. For the liquidity component, boundedness, the short-time product condition, and dominated convergence give $\Lambda_t = \alpha_L \cdot S_0 \cdot \bar{L} \cdot \sigma_0 \cdot T + O(T^{1+\beta})$. The ATM Black–Scholes vega satisfies $\mathcal{V}_{\mathrm{ATM}}(T) = S_0\sqrt{T/(2\pi)} + O(T^{3/2})$. A first-order Taylor expansion of BSM around σ^0_{ATM} then yields

$$\Delta\sigma_{\mathrm{ATM}}(T) = \frac{\Lambda_t}{\mathcal{V}_{\mathrm{ATM}}(T)} + O(T) = \frac{\alpha_L S_0 \bar{L} \sigma_0 T}{S_0\sqrt{T/(2\pi)}} + O(T^{1/2+\beta}) = \alpha_L \bar{L}\sigma_0\sqrt{2\pi T} + O(T^{1/2+\beta}).$$

The remainder collects higher-order terms in T from the Taylor expansion of BSM and the $O(T^{1+\beta})$ term in Λ. □

Corollary 1 (Discretized scaling)**.** *If* Λ *is computed by Riemann sum with step* $\Delta = T/n$*, the leading-order prediction has no artificial dependence on* n*. If instead one uses* $L_k\sqrt{V_k}\sqrt{\Delta}$ *as the per-step increment, the predicted lift carries an extra factor of* $\sqrt{n_{steps}}$*.*

Theorem 2 gives a closed-form prediction. Table 2 reports the prediction against Monte Carlo for $T \in \{7, 14, 30, 60\}$ days at the calibrated parameters.

The MC/predicted ratio lies between 94% and 96% across maturities. The residual 4–6% is consistent with the higher-order theorem remainder and Monte Carlo sampling error.

Table 2. Closed-form ATM IV lift vs. Monte Carlo, at calibrated $\alpha_L = 0.21$, $\bar{L} \approx 1.0$, $\sigma_0 = 42.16\%$. The ratio column reports MC/predicted; values close to 1 indicate that the leading-order asymptotic approximates the liquidity contribution in this numerical setting.

T (Days)	Predicted $\Delta\sigma_{\text{ATM}}$ (%)	MC $\Delta\sigma_{\text{ATM}}$ (%)	Ratio
7	4.155	3.969	0.96
14	5.876	5.574	0.95
30	8.601	8.109	0.94
60	12.164	11.430	0.94

3.4. Numerical Methodology and Implementation Checks

3.4.1. Lagged-Variance Euler Scheme

We use a martingale-preserving lagged-variance Euler discretization: variance over $[t_{k-1}, t_k]$ is $\mathcal{F}_{t_{k-1}}$-measurable. The asset increment over the same interval uses fresh information $z_1[k]$ via the correlated noise $z_S = \rho\, z_1[k] + \sqrt{1-\rho^2}\, z_2[k]$, so that V_k and z_S are conditionally independent. This recovers BSM in the $\eta \to 0$ limit (verified numerically to MC noise at $\eta = 10^{-3}$).

Theorem 3 (Strong L^2 convergence of the lagged-variance Euler scheme)**.** *Let S^n denote the discretized asset process with n equal time steps under the lagged-variance scheme, and S the continuous-time process satisfying $dS_t/S_t = r\,dt + \sqrt{V_t}\,dW_t + dJ_t$ under the EMM $\mathbb{Q}$ of Theorem 1. Assume the integrability conditions of Theorem 1, that V_t has Hölder-continuous paths of every order $\gamma \in (0, H)$ almost surely, and that $V_t \geq V_{\min} > 0$ on the numerical state space. Then, for every such γ,*

$$\mathbb{E}^{\mathbb{Q}}\left[\sup_{0\leq t\leq T} |S_t^n - S_t|^2\right]^{1/2} \leq C_\gamma\, n^{-\gamma},$$

where $C_\gamma = C(T, \eta, \sigma_0, \alpha_L, \gamma)$ does not depend on n.

Proof sketch. Work first with the log-price between jump times. The discretization error decomposes additively into (a) the variance-lag error, $V_t - V_{t_{k-1}}$ on $[t_{k-1}, t_k]$, and (b) the Brownian-increment quadrature error, $\int_{t_{k-1}}^{t_k} \sqrt{V_u}\, dW_u - \sqrt{V_{t_{k-1}}}\,\Delta W_k$.

For (a), Hölder continuity of V implies $|V_t - V_{t_{k-1}}| \leq K_\gamma\, \Delta^\gamma$ for each $\gamma < H$ and some random constant $K_\gamma \in L^2$, by the rough-volatility regularity of B^H (see [7], proposition 2.1). Squaring and integrating gives $\int_0^T |V_t - V_{t_{k-1}}|^2\, dt = O(\Delta^{2\gamma})$ in L^2.

For (b), the Itô isometry combined with the lagged-variance martingale construction gives $\mathbb{E}\Big[\big(\int_{t_{k-1}}^{t_k} \sqrt{V_u}\, dW_u - \sqrt{V_{t_{k-1}}}\,\Delta W_k\big)^2\Big] \leq 2\,\mathbb{E}\Big[\int_{t_{k-1}}^{t_k} |\sqrt{V_u} - \sqrt{V_{t_{k-1}}}|^2\, du\Big] = O(\Delta^{1+2\gamma})$ per step, yielding $O(\Delta^{2\gamma})$ total.

The compound-Poisson component has finite activity under Theorem 1; the probability of more than one jump in a grid cell is $O(\Delta^2)$, and the assumed exponential jump moment gives an L^2 bound for the remaining finite-variation jump-discretization term. Summing the diffusion and jump contributions and applying Doob's inequality to the martingale part gives the stated rate for the log-price. The exponential map is locally Lipschitz on the L^2-bounded price paths generated under the true-martingale condition, which transfers the rate to $S^n - S$. The Lipschitz regularity of $x \mapsto \sqrt{x}$ on $\{V \geq V_{\min}\}$ is preserved by the variance lower clip used in the implementation. □

Remark 2. *For our calibration ($H = 0.10$), Theorem 3 gives a conservative worst-case convergence rate rather than a point prediction for the bias. The 4–6% residual gap between the closed-form asymptotic and MC in Table 2 is therefore read as an empirical numerical-error diagnostic plus the leading-order Taylor remainder of Theorem 2, not as a direct estimate of the asymptotic rate constant.*

3.4.2. Parameter Table

All simulation parameters are tabulated in Table 3. The baseline spot $S_0 = 76{,}910.56$ matches the Deribit index level at the snapshot time. The Hurst grid is $H \in \{0.06, 0.10, 0.14, 0.20, 0.30\}$. The volatility-of-volatility $\eta = 1.0$. The variance baseline $\sigma_0 = 42.16\%$ matches the median mark IV of the regression sample.

Table 3. Parameter and data dictionary.

Parameter	Value	Source/Note
Spot index S_0 (USD)	76,910.56	Deribit snapshot, 18 May 2026 20:45 UTC
Risk-free rate r (annual)	0.03	Working assumption; insensitive to small changes
Snapshot median mark IV	0.4216	Cross-sectional median, used for vega and target ATM price
Baseline vol σ_0 (proposed)	0.4216	Set equal to snapshot median mark IV
Common $\sqrt{\theta}$ (Heston)	0.4216	Set equal to snapshot median mark IV
Common $\sqrt{\xi_0}$ (rough Bergomi)	0.4216	Set equal to snapshot median mark IV
Maturity T (days)	30	Short-dated ATM call
Hurst exponent H	0.06–0.30	Rough volatility regime; lower = rougher
Normal-state persistence p_{00}	0.97	Calibration anchor
Stressed-state persistence p_{11}	0.84–0.99	Comparative statics range
Vol-of-vol η	1.00	Rough-volatility benchmark setting
Leverage correlation ρ	-0.55	Crypto leverage effect
Proposed-model stressed-jump intensity λ_J	0.03 annual	Stress jump rate in proposed simulation
Merton/Bates benchmark jump intensity λ_J	2.00 annual	Benchmark jump-diffusion setting
Hedging TC (bps)	10	Proportional transaction cost per turn
Monte Carlo paths	3000–20,000 depending on experiment	
Time steps	120–252 depending on experiment	Per 30-day window unless otherwise stated
Liquidity loading α_L	0.2106	Calibrated to snapshot wide-vs-tight IV gap via BSM vega
Random seed (master)	20260518	Fixed for reproducibility

3.4.3. Discrete-Time Error

The MC standard error in the benchmark setting is about \$85 on a price of \$4343, i.e., approximately 2.0% of the option price at 4000 paths. Variance-reduction via antithetic sampling halves this; quasi-Monte Carlo or hybrid schemes can do more. We retain plain Monte Carlo to keep the asymptotic comparison transparent.

The coarsest run in Table 4 is deliberately included as a low-compute stress check and gives a lower rough payoff estimate. Once the simulation uses at least 4000 paths and 200 time steps, the implied-volatility lift is stable between 8.10 and 8.23 vol points, and the total price remains within about 1.5% of the highest-time-resolution row.

Table 4. Numerical sensitivity of the 30-day ATM benchmark to Monte Carlo paths and time steps. The final row is the highest-time-resolution check used as a reference point.

Paths	Steps	Rough Price	Liquidity Premium	Total Price	IV Lift (pp)
2000	100	3433.37	689.15	4122.53	7.86
4000	200	3626.12	710.89	4337.01	8.11
8000	200	3651.67	710.31	4361.98	8.10
4000	400	3678.58	721.83	4400.41	8.23

Appendix A provides the detailed measure-change construction, short-time product condition, and Euler error decomposition used in this section.

4. Empirical Design and Option-Chain Evidence

The empirical analysis follows the two testable implications and distinguishes them from the simulated hedging comparative static. The option-chain snapshot tests whether quoted illiquidity is associated with implied volatility after local surface controls. The calibration maps that association into the model loading α_L; it is not reused as independent evidence. The trade panels, realised-volatility HMM diagnostic, Ether replication, placebos, holdout predictions, and empirical hedge proxy then assess what survives changes in proxy, sample window, and asset.

The design is intentionally associational. The public data do not contain an exogenous liquidity shock, and the proxies may be affected by volatility itself. The regressions therefore test whether liquidity proxies retain explanatory power after surface and market-state controls; they do not claim that liquidity causes implied volatility in a structural causal sense.

4.1. Data Source and Cleaning

The empirical analysis starts from a public Deribit API snapshot [38] collected at 18 May 2026 20:45:20 UTC. The raw pull contained 888 Bitcoin option contracts. After removing expired, zero-price, missing-implied-volatility, and extreme-moneyness observations, the filtered cross-section contains 700 contracts across 10 expiries, balanced calls and puts. BTC index level: \$76,910.56. Median mark IV across the full filtered sample is 42.41%; after the additional regression-sample restrictions used below, the regression-sample median is 42.16% and is the baseline volatility used in the Monte Carlo. Median relative bid–ask spread is 0.031 and the 90th-percentile relative spread is 0.400.

The snapshot is a cross-sectional design, not a time-series estimate of average market conditions. It is used because quoted bid and ask prices provide a clean relative-spread proxy at one timestamp. The trade panels below cover different windows and use a different proxy, trade-vs-mark deviation, because public historical trade records do not provide contemporaneous full-book spreads for every trade. These proxies are related but not identical: relative spread measures quoted market liquidity, while trade-vs-mark deviation also reflects execution price variation, stale marks, trade direction, and moneyness effects. The paper compares the sign and order of magnitude across proxies, not equality of coefficients.

Table 5 makes every material exclusion reproducible. Instrument names and trade IDs are the deduplication keys; the archived snapshot has one row per instrument and the short-window feeds contain unique trade IDs, so no additional rows are lost to deduplication there. Relative spread is $(ask - bid)/mid$; missing or zero-sided quotes explain why 700 snapshot contracts become 679 regression observations. For ETH, same-day time to expiry is lower-bounded at 0.5 days for the square-root-maturity control, matching the submitted code.

4.2. Distributional Snapshot Statistics

Table 6 reports the distributional statistics used to diagnose the option-chain cross-section. The filtered sample contains dispersion useful for identification: mark IV has a standard deviation of 10.93 percentage points, relative spreads range from tight quotes near 1.4% at the 10th percentile to 40.0% at the 90th percentile, and open interest is highly skewed. Mark IV has $n = 700$ because it does not require a two-sided quote; relative spread has $n = 679$ because 21 contracts lack a valid positive-mid bid–ask pair. The liquidity proxy therefore varies materially within the same timestamp.

Table 5. Cleaning audit. Counts are remaining observations after each sequential rule.

Sample	Sequential Rule	Remaining	Removed at Step
BTC snapshot	Raw public-book records	888	0
	DTE > 1, positive mid and IV, moneyness 0.5–1.8	700	188
	Complete quote; relative spread 0–5; IV < 500	679	21
BTC 24-h trades	Raw trades	6194	0
	Positive IV	6165	29
	DTE 1–365 days	4515	1650
	$\lvert \log(K/S) \rvert < 1.5$; mark deviation 0–5	4513	2
ETH trades	Raw trades	2984	0
	IV 0–250; same-day DTE lower bound 0.5 day	2981	3
	$\lvert \log(K/S) \rvert < 1.5$	2945	36
	Mark deviation 0–5	2943	2

Table 6. Distributional statistics for the filtered Deribit BTC option snapshot.

Variable	*n*	Mean	Sd	p10	Median	p90
Mark IV (Relative spread)	679	0.130	0.260	0.014	0.031	0.400
Abs. log moneyness	700	0.173	0.149	0.020	0.123	0.402
Days to expiry	700	92.392	104.683	2.135	38.135	311.135
Open interest	700	495.673	929.296	0.000	75.750	1518.620
Volume USD	700	55,644.508	262,154.886	0.000	1203.750	110,957.411
Mark price USD	700	8812.586	11,022.739	45.317	4380.059	24,594.325

Table 7 shows that the widest spreads concentrate in the shortest maturities, while longer-dated options have lower median relative spreads but larger moneyness dispersion. The empirical design therefore controls for maturity and moneyness before interpreting the liquidity coefficient.

Table 7. Option-type and maturity-bucket diagnostics.

Type	Maturity	*n*	Median IV	Median Spread	Median $\lvert \log(K/S) \rvert$	Median OI
C	0–7d	80	39.58	0.144	0.052	12.4
C	7–30d	68	41.77	0.035	0.117	305.3
C	30–90d	81	39.10	0.027	0.168	265.2
C	90–180d	39	41.08	0.017	0.211	464.3
C	>180d	82	43.64	0.017	0.222	65.4
P	0–7d	80	39.58	0.127	0.052	13.9
P	7–30d	68	41.77	0.051	0.117	217.7
P	30–90d	81	39.10	0.024	0.168	113.7
P	90–180d	39	41.08	0.020	0.211	185.2
P	>180d	82	43.64	0.020	0.222	47.1

The correlation matrix is reported in Table 8; it summarizes the raw pairwise dependence before surface controls are imposed.

4.3. Liquidity-Sorted Surface Diagnostics

Prediction 1 requires that the liquidity proxy not merely reproduce maturity or moneyness. The liquidity-sorted diagnostics therefore report how spread quintiles differ in implied volatility, absolute moneyness, and days to maturity before the regression controls are imposed. The raw pattern is not interpreted causally; it is a presentation check that motivates the fixed-effect regressions that follow. Figure 1 visualizes the correlations in Table 8, while Tables 9 and 10 report the quintile and decile diagnostics, respectively.

Table 8. Correlation matrix for key snapshot variables. IV is mark implied volatility; Spread is relative bid–ask spread; AbsMny is $|\log(K/S)|$; SqrtMat is $\sqrt{\tau}$; LogOI and LogVol are log one plus open interest and USD volume.

	IV	Spread	AbsMny	SqrtMat	LogOI	LogVol
IV	1.00	0.03	0.53	−0.01	0.01	−0.29
Spread	0.03	1.00	−0.07	−0.39	−0.04	−0.10
AbsMny	0.53	−0.07	1.00	0.39	0.03	−0.36
SqrtMat	−0.01	−0.39	0.39	1.00	0.07	−0.09
LogOI	0.01	−0.04	0.03	0.07	1.00	0.59
LogVol	−0.29	−0.10	−0.36	−0.09	0.59	1.00

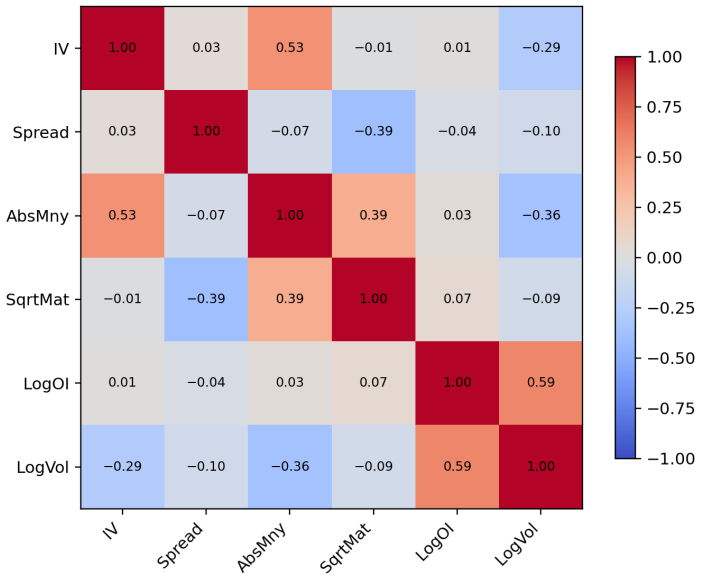


Figure 1. Correlation heatmap for key option-chain variables.

Table 9. Deribit BTC option snapshot sorted by relative bid–ask spread quintile.

Spread Quintile	*n*	Median Rel. Spread	Median IV	Median $\lvert\log(K/S)\rvert$	Median Days
1	136	0.014	45.03	0.176	129.1
2	136	0.020	42.59	0.153	129.1
3	135	0.031	41.21	0.146	73.1
4	136	0.069	38.45	0.097	10.1
5	136	0.400	40.98	0.091	3.1

Table 10. Liquidity-decile diagnostics based on relative bid–ask spread.

Decile	*n*	Median IV	Median Spread	Median Days	Median Mark USD
1	68	47.63	0.011	129.1	19,347.9
2	68	43.39	0.016	129.1	14,721.5
3	68	42.00	0.019	129.1	14,859.9
4	68	42.78	0.022	101.1	8828.6
5	68	40.91	0.027	73.1	6061.5
6	67	41.30	0.037	17.1	3149.4
7	68	39.62	0.054	17.1	1652.9
8	68	37.45	0.090	10.1	695.7
9	68	38.80	0.195	3.1	161.2
10	68	43.44	0.667	2.6	24.9

4.4. Cross-Sectional IV Regression

The regression with maturity (expiry-bucket) fixed effects, an open-interest weight, a stale-quote dummy, and subsample robustness is

$$IV_i = \beta_0 + \beta_1 \log(K_i/S) + \beta_2 \log^2(K_i/S) + \beta_3 \sqrt{\tau_i} + \beta_4 \log(1 + \text{RelSpread}_i) + \beta_5 \mathbf{1}_{\{put\}} + \gamma \mathbf{1}_{\{stale\}} + \sum_e \delta_e \mathbf{1}_{\{e\}} + \varepsilon_i.$$

Figure 2 shows the liquidity-decile pattern.

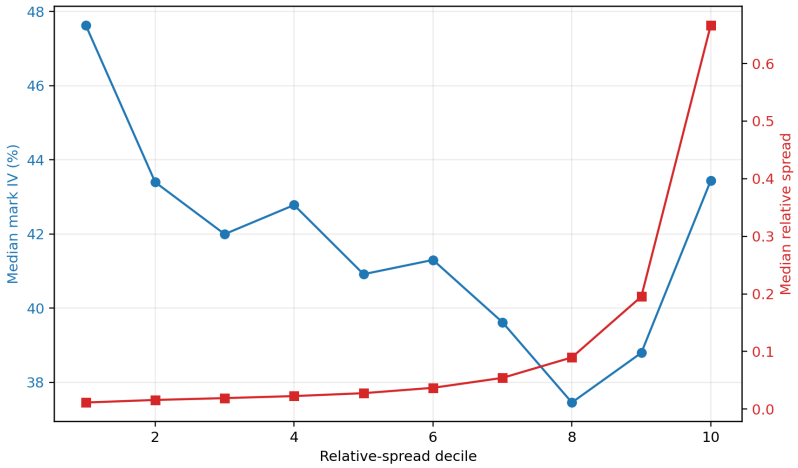


Figure 2. Median implied volatility and median relative spread by liquidity decile.

The complete-case regression sample contains 679 contracts after requiring nonmissing spread and control variables, and Table 11 reports the cross-sectional estimates.

Table 11. Cross-sectional Deribit IV regressions. Coefficient on $\log(1 + \text{RelSpread})$ across robustness specifications.

Specification	n	$\hat{\beta}_{\text{liq}}$	Robust s.e.	Robust t
Baseline (no FE)	679	4.215	2.455	1.72
+ stale-quote dummy	679	4.440	2.599	1.71
+ expiry FE	679	6.516	2.338	2.79
+ open-interest weights	679	23.723	3.470	6.84
Calls only, with FE	334	−0.185	2.247	−0.08
Puts only, with FE	345	17.838	6.466	2.76
7–180 day maturity, with FE	371	9.326	3.201	2.91

The open-interest (OI) weight gives greater influence to contracts carrying more outstanding economic exposure; it is not claimed to be statistically mandatory. The unweighted FE estimate is 6.52 ($t = 2.79$), whereas OI weighting produces 23.72 ($t = 6.84$). This sensitivity is reported rather than hidden. In units, IV is measured in annualised percentage points and the regressor is $\log(1 + \text{relative spread})$. Thus the OI-weighted coefficient predicts $23.72[\log(1.400) - \log(1.031)] = 7.26$ IV percentage points when spread moves from its median to its 90th percentile.

The option-type split is substantial: calls give −0.185 ($t = -0.08$), while puts give 17.84 ($t = 2.76$). Consequently, the snapshot support is put-side rather than symmetric. The core model's $h_C = h_P$ benchmark preserves parity; the empirical asymmetry points to option-type inventory demand or an extension with different quote-reserve multipliers, not confirmation of a symmetric cross-section.

Because a relative spread can be mechanically large for a low-premium contract, Table 12 adds log option premium. Vega is not used as the alternative control because it is computed from the dependent variable, mark IV, and would be mechanically endogenous. The liquidity coefficient remains positive with the premium control, but its change reinforces the conditional, specification-dependent interpretation. Continuous-regressor VIFs are all below 2.80 (liquidity VIF 1.86), so the result is not explained by severe multicollinearity.

Table 12. Weighting and option-premium robustness. HC1 standard errors; all models contain the surface controls, stale-quote indicator, and expiry fixed effects.

Specification	n	R^2	$\hat{\beta}_{\text{liq}}$	Robust t
Unweighted FE	679	0.639	6.516	2.79
Unweighted FE + log premium	679	0.639	5.855	2.50
OI-weighted FE	679	0.742	23.723	6.84
OI-weighted FE + log premium	679	0.746	19.597	4.84

To separate in-sample calibration from validation, we leave out one expiry at a time, estimate on the remaining expiries, and predict the held-out expiry. Adding liquidity lowers OI-weighted RMSE from 6.66 to 5.67 IV points and lowers unweighted MAE from 5.72 to 5.28, but raises unweighted RMSE from 7.90 to 8.27. The result is therefore mixed rather than evidence of universal predictive dominance. Figure 3 reports these prediction errors.

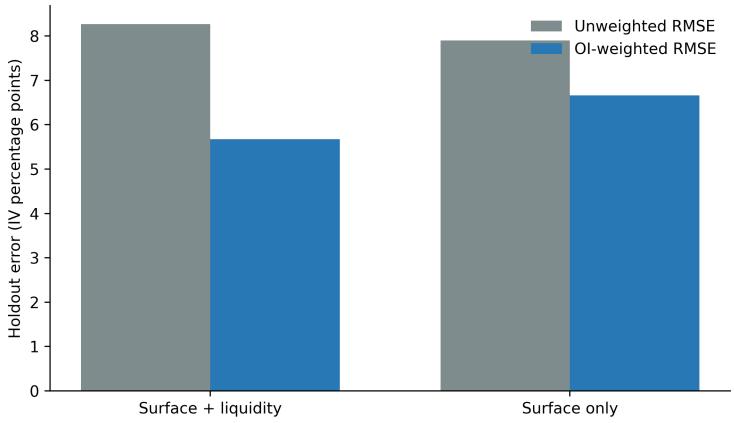


Figure 3. Leave-one-expiry-out prediction errors. The liquidity specification improves economically weighted RMSE but not unweighted RMSE.

4.5. Interaction Regressions

Lemma 1 and Proposition 5 imply that residualized liquidity variation, not raw spread variation, identifies the first-order liquidity channel. We therefore add interactions between the liquidity proxy and short maturity, and between the liquidity proxy and put-option status:

$$IV_i = X_i'\gamma + \beta_Z Z_i + \beta_{ZS} Z_i \mathbf{1}_{\{\tau_i \le 14/365\}} + \beta_{ZP} Z_i \mathbf{1}_{\{put\}} + \varepsilon_i.$$

Table 13 reports the resulting interaction estimates.

The interaction results qualify the average coefficient after applying the same fixed-effect and weighting structure used in the primary snapshot specification. The liquidity coefficient remains positive. The put interaction is positive, while the short-maturity interaction is weaker after expiry fixed effects absorb much of the maturity-level variation. This descriptive pattern is consistent with a state-dependent interpretation: the liquidity-

IV slope is steeper in parts of the surface where downside hedging demand is likely to be concentrated.

Table 13. FE-controlled interaction regressions for the Deribit snapshot. Specifications include the primary surface controls, stale-quote dummy, expiry fixed effects, and open-interest weights.

Specification	n	R^2	β_{liq}	$t(\beta_{\text{liq}})$	$\beta_{\text{liq}\times\text{short}}$	$\beta_{\text{liq}\times\text{put}}$	Max Interaction t
Primary FE + OI weight	679	0.742	23.723	6.84	—	—	—
+ spread × short	679	0.747	18.370	4.97	9.192	—	1.55
+ spread × put	679	0.849	16.264	7.37	—	51.963	9.19
+ both interactions	679	0.850	14.823	4.95	2.594	51.482	8.98

4.6. Calibration of the Liquidity-Risk Loading

We discipline α_L by mapping the FE-controlled regression coefficient onto the model premium. The implied IV widening from a near-zero spread to the 90th-percentile spread $p_{90} = 0.400$ is

$$\Delta\sigma = \frac{\hat{\beta}_4^{FE}}{100}\log(1 + p_{90}) = 0.0798 \text{ vol units.}$$

Converted to USD via BSM vega at the ATM IV ($\mathcal{V}_{\text{ATM}} = \$8768$) gives a target liquidity premium of approximately \$700. Solving for the loading that reproduces this premium under the proposed simulation gives

$$\alpha_L = 0.2106.$$

With this value, the proposed model's liquidity addon at the benchmark grid point is \$708, or roughly 0.92% of spot. Section 3.3 verifies that the closed-form prediction agrees with the simulation to within about 4–6% across maturities. This calibration uses the OI-weighted FE coefficient; the unweighted FE coefficient would imply a smaller loading, so the sensitivity analysis below is part of the empirical interpretation rather than a cosmetic robustness check.

The mapping above is calibration, not an independent empirical test: the same snapshot slope supplies both the \$700 target and α_L. Independent evidence is limited to the leave-one-expiry-out exercise, the later-date trade panel, Ether replication, and placebos; none is used to recalibrate this loading. A 1000-replication pairs bootstrap propagates coefficient uncertainty through the linear mapping. Table 14 reports the propagated bootstrap interval.

Table 14. Calibration uncertainty propagated from a 1000-replication pairs bootstrap.

Quantity	Estimate	95% Lower	95% Upper
Liquidity–IV coefficient (IV pp)	23.723	18.101	31.672
Calibrated α_L	0.211	0.161	0.281
Predicted 30-day ATM premium (USD)	700	534	935

To prevent the single-snapshot calibration from being overread as a universal structural estimate, Table 15 reports the same mapping inside tenor buckets. Short-tenor buckets imply the largest liquidity loading because their p90 spreads and ATM vegas convert spread variation into economically larger IV moves. Longer tenors imply smaller loadings, consistent with the interpretation that acute order-book frictions are concentrated in short-dated hedging demand. The paper therefore uses $\alpha_L = 0.2106$ as a transparent 30-day ATM benchmark loading, not as a maturity-invariant structural constant.

Table 15. Maturity-bucket diagnostic for the liquidity-risk loading. Each row re-estimates the OI-weighted liquidity-IV projection inside the tenor bucket and maps the p90-spread IV widening into an ATM dollar premium using BSM vega.

Tenor	n	β_{liq}	t-Stat	p90 Spread	Median Days	Target Premium	Implied α_L
0–7d	144	7.76	1.63	0.756	2.1	\$102	0.477
7–30d	131	14.04	2.90	0.545	10.1	\$312	0.285
30–90d	162	7.51	2.55	0.132	73.1	\$128	0.017
90–180d	78	15.33	3.12	0.063	129.1	\$170	0.012
>180d	164	9.62	1.59	0.039	220.1	\$86	0.003

5. Mechanism Checks, Panels, and Robustness

5.1. Six-Model Benchmark Comparison

The benchmark models are evaluated at common parameters tied to the snapshot median IV. They should therefore be read as a mechanism comparison, not as separately optimized calibrations of each model. Because the observed median IV already embeds every priced component, adding Λ to a rough baseline initialized at that IV would double count liquidity if the total were interpreted as a fitted market price. The proposed row is therefore explicitly a stress-overlay decomposition—"observed-IV payoff baseline plus calibrated reserve"—not a prediction of the observed option price. Recovering a liquidity-free σ_0 would require an independently identified structural decomposition that these data do not provide. The exercise asks what the liquidity component adds on a common parameter scale. It is not a claim of superior pricing performance against separately calibrated BSM, Heston, Merton, Bates, or rough Bergomi alternatives. Table 16 reports the values, and Figure 4 visualizes the comparison.

Table 16. Mechanism stress-overlay comparison for a 30-day ATM Bitcoin call at $S_0 = \$76{,}910.56$. The common observed-IV initialization is not an out-of-sample calibration; the proposed total deliberately displays the incremental reserve and must not be read as a fitted market price.

Model	Price (USD)	BSM-IV (%)	Liquidity Premium (USD)	IV Gap vs. Market (pp)
Black–Scholes	3797.26	42.16	0.00	−0.00
Heston	3670.66	40.72	0.00	−1.44
Merton (jump)	4000.80	44.48	0.00	+2.32
Bates (SV+jump)	3902.29	43.36	0.00	+1.20
Rough Bergomi (no liquidity)	3614.41	40.07	0.00	−2.09
Proposed: rough + Markov liquidity	4342.86	48.38	708.26	+6.22

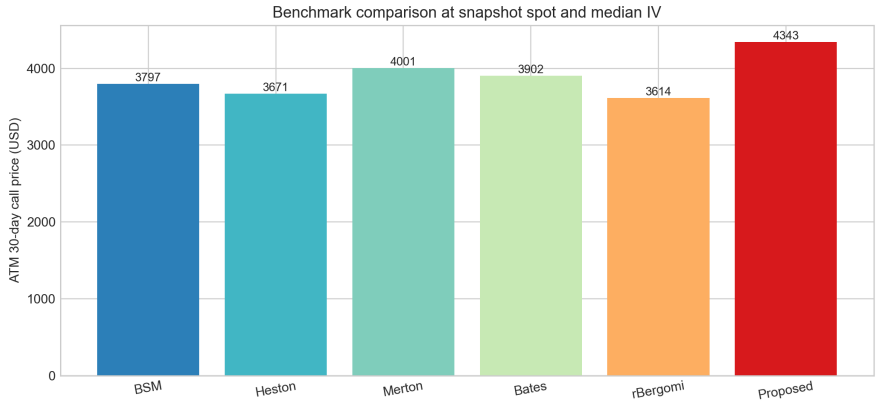


Figure 4. Six-model benchmark comparison. The proposed model's price exceeds the rough Bergomi baseline by the calibrated liquidity premium.

Heston (−1.4 vol pp) and rough Bergomi (−2.1 vol pp) lie below BSM at common parameters from leverage skew. Merton (+2.3 vol pp) and Bates (+1.2 vol pp) lie above from the jump compensator. The proposed stress-overlay raises the rough price by the calibrated reserve; its resulting IV is not compared with market IV as a goodness-of-fit claim. Since the alternatives are not comparably calibrated, this table establishes neither model ranking nor superior performance.

5.2. Main Grid

Figures 5 and 6 show how the option price and liquidity reserve vary over the Hurst–persistence grid. Table 17 reports the corresponding numerical values and separates the rough payoff from the liquidity premium.

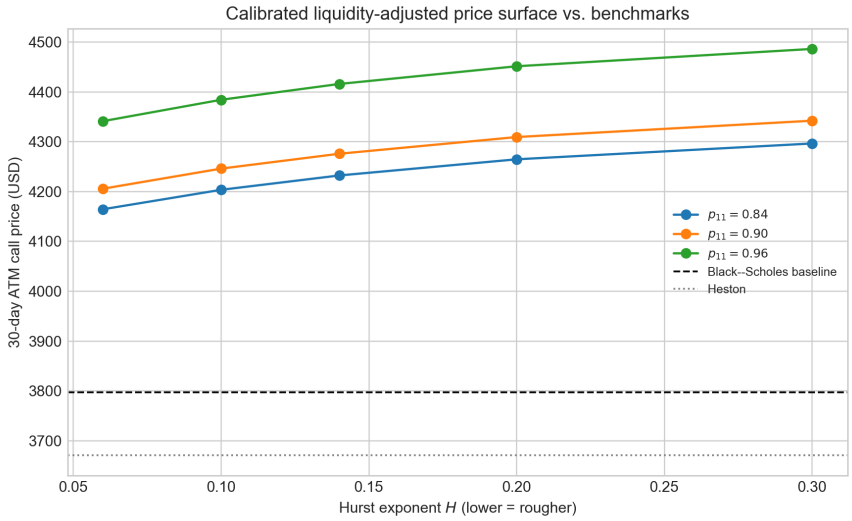


Figure 5. Liquidity-adjusted option price across Hurst exponents and stressed-regime persistence.

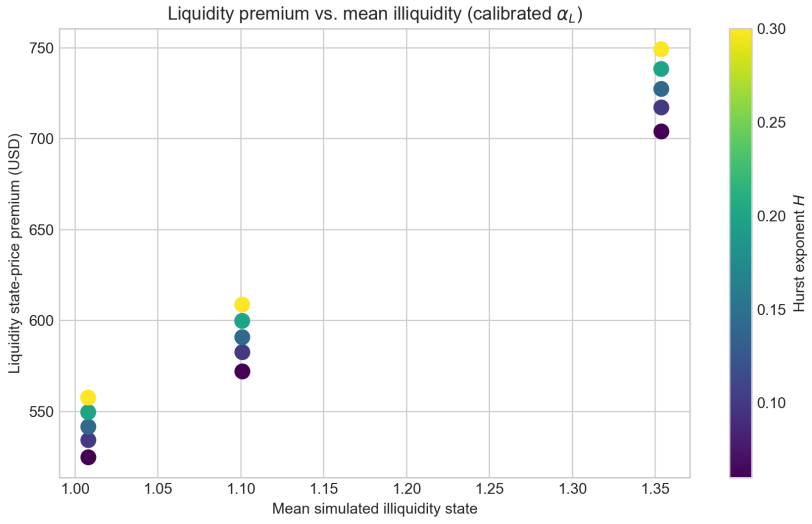


Figure 6. Liquidity state-price premium vs. mean simulated illiquidity. With α_L calibrated, the premium ranges between roughly \$525 and \$750 across the grid.

5.3. Sensitivity to the Liquidity Loading

The liquidity loading α_L is calibrated from the BTC cross-section, so the first robustness check asks whether the numerical conclusion depends on a narrow calibration value. Table 18 scales the calibrated loading from one quarter to four times its benchmark value while holding the rough payoff component fixed. The exercise is deliberately mechani-

cal: it isolates the contribution of the liquidity state-price term from the rough-volatility payoff channel.

Table 17. Synthetic numerical experiment for 30-day ATM Bitcoin calls. Prices in USD; rough payoff and liquidity premium reported separately.

H	p_{11}	Rough Price	Liquidity Premium	Total Price	Tail Return p05
0.06	0.84	3639.61	524.79	4164.40	−0.197
0.06	0.90	3633.66	571.99	4205.65	−0.198
0.06	0.96	3637.22	703.95	4341.16	−0.198
0.10	0.84	3669.22	534.23	4203.45	−0.201
0.10	0.90	3663.41	582.58	4245.99	−0.201
0.10	0.96	3666.93	717.16	4384.09	−0.201
0.14	0.84	3690.67	541.58	4232.24	−0.200
0.14	0.90	3685.00	590.81	4275.81	−0.200
0.14	0.96	3688.49	727.33	4415.82	−0.200
0.20	0.84	3715.04	549.59	4264.63	−0.201
0.20	0.90	3709.51	599.79	4309.29	−0.201
0.20	0.96	3713.00	738.32	4451.32	−0.201
0.30	0.84	3738.76	557.63	4296.39	−0.196
0.30	0.90	3733.34	608.77	4342.11	−0.197
0.30	0.96	3736.90	749.19	4486.10	−0.196

Table 18. Sensitivity to the liquidity-risk loading α_L (calibrated value $\alpha_L^* = 0.21$).

α_L Multiplier	α_L	Rough Price	Liquidity Premium	Total	Premium/Spot (%)
0.25	0.0527	3498.09	179.76	3677.85	0.234
0.50	0.1053	3498.09	359.51	3857.60	0.467
1.00	0.2106	3498.09	719.03	4217.12	0.935
2.00	0.4212	3498.09	1438.05	4936.14	1.870
4.00	0.8425	3498.09	2876.11	6374.19	3.740

The liquidity premium is approximately linear in α_L, as implied by the pricing functional. At the benchmark loading, the premium is below one percent of spot; even doubling the loading keeps the premium below two percent of spot. This scale is economically meaningful for an option premium but avoids the implausibly large liquidity corrections that would make the model a disguised volatility markup.

5.4. Delta-Hedging Performance Across Regimes

The pricing term is motivated by hedging frictions, so the next check examines whether hedging errors worsen when the stressed liquidity regime becomes more persistent. We simulate a BSM delta hedge against paths generated by the proposed rough-liquidity model and vary the stressed-state persistence p_{11}. This is not meant to be a complete market-making strategy; it is a diagnostic for whether liquidity persistence changes the distribution of replication errors in the direction predicted by the model.

Table 19 and Figure 7 show a monotone increase in RMSE as the stressed regime becomes more persistent. Mean hedging PnL also becomes more negative and the lower tail deteriorates. This is an internal model comparative static only; the empirical backtest below does not reproduce it.

We additionally conduct an empirical one-step hedge diagnostic on the 92-date panel. Consecutive trades in the same instrument are paired when 1 min to 6 h apart. The hedge residual is the change in USD mark value less the contemporaneous BSM delta gain and theta accrual, normalised by spot. This produces 382,407 pairs across 4154 instruments. Contrary to the simulation comparative static, median absolute error does not rise with the prior trade-vs-mark deviation; it declines in the upper quintiles. A clustered regression

of $\log(1 + |error|)$ on prior $\log(1 + mark\ deviation)$ and moneyness, maturity, option type, and rehedge-time controls gives -1.449 (s.e. 0.053, $t = -27.40$). Because marks and trade deviations are imperfect proxies for executable dealer hedges, this is not a definitive rejection, but it does not validate the model's third comparative static. Table 20 and Figure 8 report the quintile results.

Table 19. Delta-hedging RMSE across liquidity regimes for the proposed model.

Regime	Hedge RMSE	Mean PnL	PnL 5%	PnL 95%
Calm regime (p_{11} = 0.50)	1090.45	−358.86	−2236.84	1058.38
Baseline (p_{11} = 0.96)	1138.20	−448.58	−2365.65	968.59
Persistent stress (p_{11} = 0.99)	1205.51	−565.37	−2552.33	863.31



Figure 7. Delta-hedging RMSE rises monotonically with stressed-state persistence.

Table 20. Empirical one-step hedging diagnostic by prior trade-vs-mark-deviation quintile.

Quintile	Pairs	Median Mark Deviation	Median abs. Error (Spot Bps)	RMSE (USD)
1	76,482	0.0014	1.009	33.54
2	76,481	0.0053	1.125	30.12
3	76,481	0.0132	1.019	26.50
4	76,481	0.0302	0.787	27.28
5	76,482	0.0855	0.418	29.94

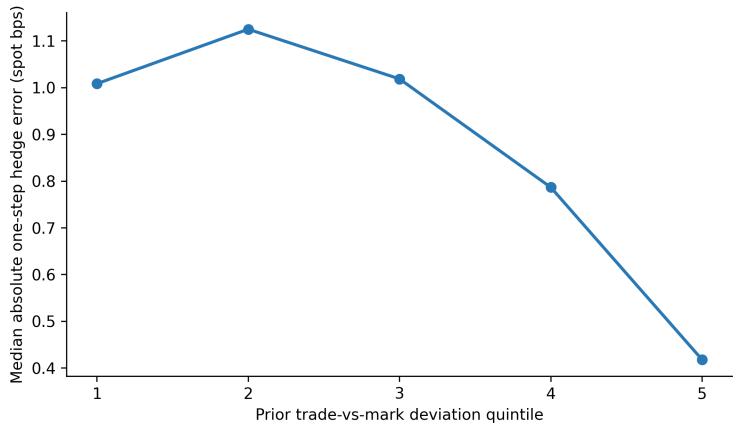


Figure 8. Empirical one-step hedge errors do not reproduce the simulation's monotone stress result.

5.5. Multiple-Seed Robustness

Monte Carlo evidence can be fragile if a headline result is driven by a favorable random seed. The robustness check in Table 21 reruns the benchmark setting across the stored seed panel and reports the mean and standard deviation of the rough payoff, liquidity premium, total price, and lower-tail return.

Table 21. Multiple-seed robustness for the benchmark setting $H = 0.10$, $p_{11} = 0.96$.

	Rough Price	Liquidity Premium	Total Price	Tail Return p05
mean	3646.09	709.82	4355.91	−0.1981
std. dev.	85.73	3.69	85.17	0.0054

The liquidity premium is much less variable across seeds than the rough payoff component, with a standard deviation of about $4 around a mean near $710. The total price variation is therefore dominated by the simulated rough payoff rather than by instability in the liquidity adjustment. This supports using the benchmark simulation as a numerical illustration of the theoretical asymptotic.

5.6. Deribit Trade-Panel Evidence

We supplement the cross-section with two Deribit trade panels. The first contains 4513 clean trades in exactly 24 h, from 19 May 2026 10:38 UTC through 20 May 2026 10:38 UTC, and therefore spans two calendar dates; it is not a 48-h sample. This archived pull was chosen because it is the complete public-trade window available at collection time that can be joined to hourly realised volatility and perpetual funding. The second uses the official Deribit historical endpoint to build a longer trade-only panel of 775,315 clean Bitcoin option trades over 92 calendar dates from 23 February 2026 02:51 UTC to 25 May 2026 01:21 UTC. The shorter panel is used for market-state joins and event-window analysis; the longer panel is used to test whether the liquidity-IV relation survives date and expiry fixed effects over a materially longer history. Table 22 summarizes the short-window sample.

Table 22. Trade-level panel descriptive statistics.

Statistic	Value
Trade observations	4513
Unique instruments	375
Window start UTC	19 May 2026 10:38
Window end UTC	20 May 2026 10:38
Median IV (%)	38.69
Median mark deviation	0.0200
Median days to expiry	9.17
Median trade contracts	0.30

5.6.1. Market-State Timeline and HMM Regimes

A two-state Gaussian HMM fitted via EM on the realized-vol series classifies the market history into persistent low- and high-volatility states with $(p_{00}, p_{11}) = (0.995, 0.988)$. Tables 23 and 24 show that the stressed state has higher realized volatility and a high median posterior stress probability. Figures 9 and 10 display the time series and fitted state probabilities, respectively. This HMM is not treated as a direct estimator of the latent liquidity regime in the theoretical model: high volatility and illiquidity often coincide in crypto markets, but they are not the same state variable. The HMM evidence is therefore a descriptive stress-state diagnostic that supports the persistence assumption used in the comparative statics; direct liquidity-regime estimation is left to richer order-book data.

Table 23. Hourly market-state summary from realized volatility and perpetual funding history.

State	Hours	Mean RV	p90 RV	Mean |Funding| bps	Sd 1h Return
Calm	651	30.72	32.02	0.1257	0.0034
Stress	70	32.83	33.31	0.0205	0.0035

Table 24. Hidden-Markov-state summary for hourly realized volatility.

HMM State	Hours	Mean RV	Median p_{stress}
Calm	222	30.23	0.000
Stress	161	32.30	1.000

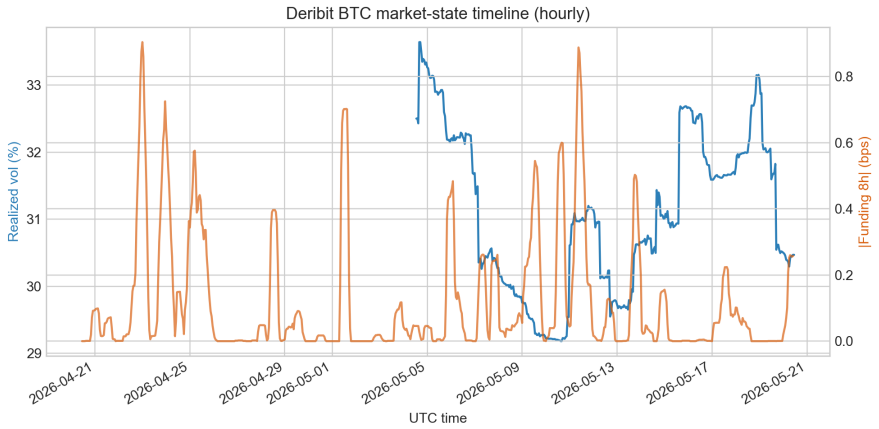


Figure 9. Deribit BTC market-state timeline. Hourly realized volatility (**left** axis) and absolute 8-h funding rate (**right** axis) over 30 days.

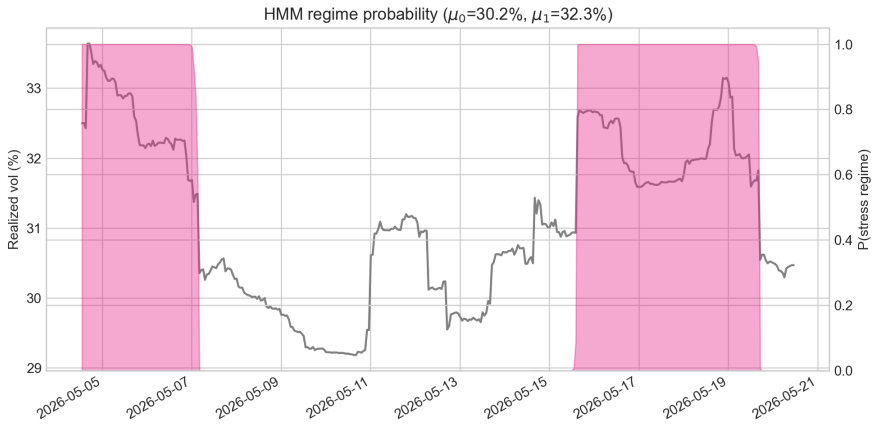


Figure 10. Two-state hidden Markov model fitted to hourly realized vol. Calm state mean $\mu_0 = 30.2\%$, stressed state mean $\mu_1 = 32.3\%$, transition probabilities $(p_{00}, p_{11}) = (0.995, 0.988)$.

5.6.2. Trade-Level Panel Regression

We use the trade-vs-mark deviation $|\text{price} - \text{mark}|/\text{mark}$ as the trade-level liquidity proxy. This measure is observable for every trade in the Deribit public-trades feed. Time-stamped market-state variables (realized vol Z-score, absolute funding rate) are joined at hourly resolution. Table 25 reports the trade-level regression specifications.

The trade-level panel coefficient is positive in every specification: $\hat{\beta} = 29.6$ (pooled, $t = 28.6$), $\hat{\beta} = 33.4$ (expiry FE, instrument-clustered, $t = 10.8$), and $\hat{\beta} = 33.4$ (date FE, instrument-clustered, $t = 10.74$). A date-clustered sensitivity gives $t = 9.25$ but is reported only descriptively because the public trade window contains two date clusters. The submitted binary stress indicator was constant within this 24-h trade window and therefore

could not identify an interaction. It is replaced by the continuous realised-volatility z-score interaction: the liquidity main effect is 40.106 ($t = 11.99$) and the interaction is 16.154 ($t = 3.18$), with clustering by instrument. The snapshot and panel estimates use different proxies and identification, so an exact match is not expected.

Table 25. Trade-level panel regressions of mark implied volatility on the trade-vs-mark deviation, with market-state controls, expiry FE, and date FE. Standard errors are HC1 or cluster-robust as indicated.

Specification	n	R^2	$\hat{\beta}_{\text{liq proxy}}$	Cluster-t
Pooled baseline	4513	0.614	29.583	28.58
+ market state controls	4513	0.629	31.051	28.45
+ expiry FE, cluster by instrument	4513	0.680	33.418	10.78
+ date FE, cluster by instrument	4513	0.680	33.439	10.74
+ date FE, cluster by date	4513	0.680	33.439	9.25
Alt liq proxy: inv-log-size	4513	0.532	0.017	0.10
+ liquidity × realised-vol z	4513	0.683	40.106	11.99

5.6.3. Historical Deribit Trade Panel

The official Deribit historical endpoint allows a broader trade-only validation without hourly market-state joins. After deduplication and cleaning, the historical panel contains 775,315 Bitcoin option trades, 4565 unique instruments, and 92 calendar dates. The same trade-vs-mark deviation is used as the liquidity proxy, with moneyness, squared moneyness, maturity, option type, trade size, date fixed effects, and expiry fixed effects as controls. Table 26 summarizes the historical sample.

Table 26. Historical Deribit BTC option trade panel summary.

Statistic	Value
Raw historical trades	1,018,078
Clean trade observations	775,315
Unique instruments	4565
Calendar dates	92
Window start UTC	23 February 2026 02:51
Window end UTC	25 May 2026 01:21
Median IV (Median mark deviation)	0.0138
Median days to expiry	14.25

Table 27 shows that the liquidity-IV coefficient remains positive in the longer panel. With date and expiry fixed effects, the coefficient is $\hat{\beta} = 46.56$ under HC1 standard errors. Clustering by date gives $t = 12.96$, and clustering by instrument gives $t = 13.75$. This long-panel evidence addresses the main temporal-depth concern in the short market-state-matched panel. The interpretation remains associational because the liquidity proxy is not an exogenous supply shock, but the result is no longer dependent on the 24-h window.

Table 27. Historical trade-panel regressions of implied volatility on trade-vs-mark deviation. The long panel contains 92 calendar dates and 4565 instruments after cleaning.

Specification	n	Dates	Instruments	β_{liq}	t-Stat
Pooled HC1	775,315	92	4565	42.889	100.49
+ date FE, HC1	775,315	92	4565	43.195	101.66
+ date and expiry FE, HC1	775,315	92	4565	46.559	107.86
+ date and expiry FE, date cluster	775,315	92	4565	46.559	12.96
+ date and expiry FE, instrument cluster	775,315	92	4565	46.559	13.75

5.6.4. Event-Study Around the Largest Realized-Volatility Hour

The event hour is assigned to the post window; brackets in Table 28 make the non-overlap explicit. The difference is estimated in one combined regression as the coefficient on liquidity × post, rather than by treating two separately estimated slopes as independent. The $n = 2225$ on the difference row is the combined pre/post sample. The gradient is positive on both sides and the post-minus-pre interaction is not significant ($t = -0.57$). Figure 11 complements the event-window estimates with the panel-state IV distributions.

Table 28. Event study: liquidity-IV gradient in the 6-h windows on either side of the peak realized-vol hour within the trade panel.

Peak Hour: 19 May 2026 16:00 UTC				
Window	n	$\hat{\beta}_{\text{liq}}$	**Robust s.e.**	t
Pre: $[-6\text{ h}, 0\text{ h})$	1483	46.966	3.526	13.32
Post: $[0\text{ h}, +6\text{ h})$	742	43.703	4.824	9.06
Post minus pre interaction	2225	−3.263	5.775	−0.57

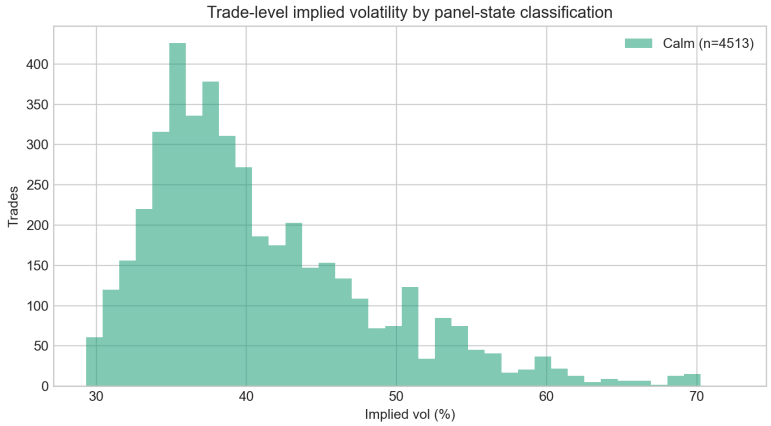


Figure 11. Trade-level implied volatility distributions by panel-state classification.

5.7. Cross-Asset Replication, Calibration Check, and Placebos

The cross-section, short-window trade panel, and historical trade panel estimate conditional liquidity–IV gradients. We now ask whether the sign replicates for Ether, document the in-sample calibration identity without treating it as evidence, and test whether randomised liquidity proxies reproduce the coefficient.

5.7.1. Cross-Asset Replication: Ether Options

We collect a filtered cross-section of 647 Ether option contracts and a panel of 2984 recent ETH option trades from the Deribit public API. ETH median IV is 58.6%, median relative spread is 0.094, and p_{90} spread is 0.476. Tables 29 and 30 report the cross-sectional and trade-panel replications, respectively.

Table 29. Ether-options cross-sectional IV regression. Coefficient on $\log(1 + \text{rel. spread})$.

Spec	n	R^2	$\hat{\beta}_{\text{liq}}$	**Robust s.e.**	t
Baseline	647	0.473	8.288	3.565	2.33
+ expiry FE	647	0.588	11.372	3.104	3.66
+ expiry FE + OI weight	647	0.836	34.329	3.729	9.21

Table 30. Ether-options trade-level panel regression. Liquidity proxy is the trade-vs-mark deviation.

Spec	n	R^2	$\hat{\beta}_{\text{liq}}$	Robust s.e.	t
ETH pooled	2943	0.557	29.557	3.636	8.13
ETH + expiry & date FE	2943	0.616	30.272	3.323	9.11

The OI-weighted-FE coefficient on ETH is $\hat{\beta} = 34.33$ ($t = 9.21$); the trade-level FE coefficient is $\hat{\beta} = 30.27$ ($t = 9.11$). The ETH estimates have the same sign and order of magnitude as the BTC estimates; the ETH snapshot estimate is somewhat larger than the BTC snapshot interval, while the ETH panel estimate lies inside the BTC panel interval. Table 31 and Figure 12 report all four main coefficients with pairs-bootstrap intervals for the cross-sections and robust asymptotic intervals for the panels. The Holm–Bonferroni adjustment across the four tests does not alter inference: all four adjusted p-values are below 10^{-6}.

Table 31. Four-estimate summary with 95% uncertainty intervals. Cross-sectional intervals use pairs bootstrap; panel intervals use robust asymptotic standard errors.

Estimate	n	$\hat{\beta}$	Robust s.e.	t	95% Interval
BTC cross-section (FE + OI weight)	679	23.723	3.469	6.84	[17.39, 30.63]
BTC trade panel (date FE, instrument-clustered)	4513	33.439	3.114	10.74	[27.34, 39.54]
ETH cross-section (FE + OI weight)	647	34.329	3.729	9.21	[26.46, 42.57]
ETH trade panel (date FE)	2943	30.272	3.323	9.11	[23.76, 36.78]

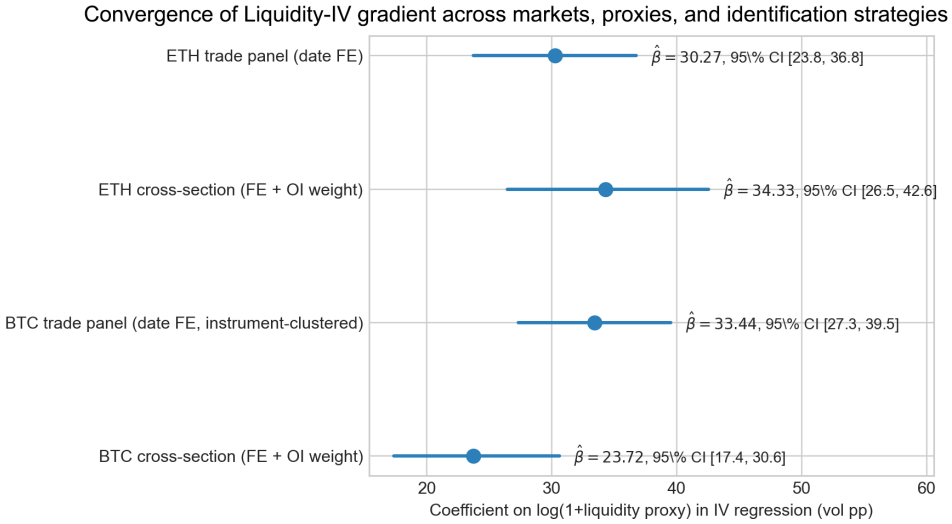


Figure 12. Liquidity-IV gradient across markets, proxies, and identification strategies. Points are point estimates; horizontal bars are 95% intervals, using pairs bootstrap for cross-sections and robust asymptotic standard errors for panels. The four estimates are positive and their intervals exclude zero.

5.7.2. In-Sample Calibration Consistency of the Closed-Form Asymptotic

Theorem 2 predicts that the calibrated IV widening from a near-zero proxy value to the 90th-percentile spread equals $\Delta\sigma = (\hat{\beta}_4^{FE}/100)\log(1 + p_{90})$ vol units. Table 32 reports this calculation for both markets.

Table 32. In-sample calibration consistency check for the closed-form asymptotic. The implied $\Delta\sigma$ is computed from the same regression coefficient used to set the loading and is not an independent test.

Asset	Median IV (%)	p_{90} rel. Spread	$\hat{\beta}_{FE}$	Data-Implied $\Delta\sigma$ (vol pp)
BTC	42.16	0.400	23.72	7.98
ETH	58.62	0.472	34.33	13.26

The implied widening is +8.0 vol points for BTC and +13.3 vol points for ETH when measured from a near-zero proxy value to p_{90}. The simulation, calibrated to the BTC regression coefficient, returns a model-MC ATM IV lift of +8.11 vol points at $T = 30$ days. Their proximity is an in-sample accounting check by construction, not corroborating empirical evidence. The leave-one-expiry-out errors, later 92-date panel, and ETH estimates provide the distinct validation exercises, with the limitations already reported.

These magnitudes are large enough to matter economically but not so large that the liquidity term becomes a disguised volatility model. In the BTC snapshot, the +8.0 vol-point move is about one fifth of the regression-sample median IV of 42.16% and converts to roughly $700 at the 30-day ATM vega used in the calibration. For a market maker or hedger, that scale is comparable to a material adjustment in quoted volatility or hedge reserve; for a directional investor, it is a warning that part of the observed smile level may compensate execution and funding conditions rather than expected terminal variance alone.

5.7.3. Placebo and Falsification Tests

The primary OI-weighted FE coefficient could still be an artifact of moneyness–maturity correlation if the liquidity proxy is only standing in for surface shape. We test this possibility with a 1000-replication random-shuffle placebo: replace the liquidity proxy with a random permutation of itself and re-run the same OI-weighted FE regression. Table 33 reports the placebo distribution and the actual estimate.

Table 33. Placebo test: 1000-replication random shuffle of the liquidity proxy.

Specification	$\hat{\beta}$	s.e. / Share	t
Actual liquidity proxy	23.723	3.469	6.838
Random shuffle (mean of 1000)	0.079	2.460	0.002
Random shuffle: |t| > 1.96 share	—	—	0.098

The shuffled-liquidity coefficient has mean 0.079 and standard deviation 2.46, with a $|t| > 1.96$ rate of 9.8%. The actual primary coefficient is $\hat{\beta} = 23.72$ ($t = 6.84$), more than nine shuffle standard deviations above the placebo mean. Random permutation removes the relationship.

6. Discussion

The results support a narrower interpretation than the submitted version. The conditional liquidity–IV gradient is positive in the full OI-weighted snapshot, short and long BTC panels, and Ether replication. Yet the snapshot result changes materially with weighting, survives an option-premium control at reduced magnitude, is positive for puts, and is null for calls. The leave-one-expiry-out check improves OI-weighted RMSE but worsens unweighted RMSE. Moreover, the proximity of the calibrated IV widening to the simulation is mechanical because the same coefficient sets α_L. These facts are consistent with a put-side liquidity reserve, but they do not separately identify liquidity from uncertainty, demand pressure, dealer inventory, or reverse causality.

For theory, the model separates rough volatility (payoff distribution) from regime-switching liquidity (hedging cost), and the asymptotic gives a closed-form calibration mapping. The additive term is economically justified as a first-order cash reserve and preserves standard put–call parity when applied symmetrically. The stochastic-calculus ingredients are standard; novelty lies in this combination and its transparent empirical mapping.

For practice, the model implies liquidity states are pricing variables. Calibration only to the implied-volatility surface may misinterpret a liquidity premium as a volatility forecast. Stress tests based only on price volatility may underestimate hedging costs during exchange-level stress.

For empirical research, the panel design used here can be reused: snapshot data, trades, and market-state histories can be combined for any actively traded crypto derivatives venue with a public or vendor API. The empirical evidence should be interpreted as a pricing association consistent with the model rather than as causal identification of exogenous liquidity shocks. This interpretation also limits the benchmark comparison: because the alternative models are not separately optimized to the same option surface, the comparison illustrates the economic role of the liquidity component rather than proving model dominance.

7. Limitations and Future Research

The principal limitation is endogeneity, not merely cross-venue breadth. Spreads, trade-vs-mark deviations, and IV are jointly affected by market uncertainty, option demand, dealer inventory, stale marks, and price level. Surface, premium, realised-volatility, funding, and fixed-effect controls reduce observable confounding but do not create an exogenous liquidity shock. The 92-date panel improves temporal coverage, while the hourly state-matched panel remains only 24 h. A vendor panel with full books, dealer inventory, quote revisions, and an exchange or protocol shock would permit an instrumental-variable or event design; CME data would test external validity.

The two-state HMM is parsimonious and describes realised-volatility states only; it does not estimate liquidity states. The empirical one-step hedge proxy fails to reproduce the simulation comparative static, so a true executable delta–gamma–vega backtest using bid/ask books is necessary before making a hedging claim. Finally, the EMM existence theorem uses sufficient exponential-moment conditions; weakening these conditions while retaining finite liquidity premia remains a theoretical extension.

8. Conclusions

This paper combines fractional rough volatility with a regime-switching additive hedge reserve and derives a short-time mapping from the reserve to IV points. In Deribit data, liquidity proxies are positively associated with IV in OI-weighted and put-side specifications, later trade panels, and Ether replication, but the call-side estimate is null and predictive diagnostics are mixed. The simulated hedging comparative static is not confirmed by the empirical one-step proxy.

The claim is intentionally bounded. The paper does not identify an exogenous liquidity shock, does not estimate every benchmark model under a common full-surface calibration, does not treat its in-sample calibration check as validation, and does not equate realised-volatility HMM states with latent liquidity regimes. The evidence therefore motivates—but does not establish—a separately priced liquidity state. For practice, the main implication is to report liquidity-reserve uncertainty and call/put asymmetry rather than treating the entire observed smile level as a volatility forecast.

Author Contributions: Conceptualization, methodology, formal analysis, investigation, writing—original draft preparation, writing—review and editing, and approval of the submitted version: E.P. and H.P.M.; software, data curation, validation, and visualization: E.P. 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: All empirical data used in the paper were collected from the public Deribit API and the official Deribit historical API at the timestamps reported in the text. The replica-

tion package contains the frozen CSV files, all scripts used to generate tables and figures, the random seeds used in the Monte Carlo experiments, and a SHA-256 hash over the data files.

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

Appendix A. Technical Appendix: Measure Change, Local Asymptotics, and Euler Error

Appendix A.1. Measure-Change Construction

Theorem 1 uses a sufficient-condition construction. Let

$$Z_T^W = \mathcal{E}_T\Big(-\int_0^T \theta_u \, dW_u\Big)$$

be the Brownian density process. Novikov's condition makes Z_T^W a true martingale. Conditional on the regime path, the jump component is tilted by an Esscher density. If $\psi_R(a) = \log \mathbb{E}[\exp(aY) \mid R]$ denotes the conditional jump cumulant and ζ_R is chosen in the interior of its moment domain, the jump-density component can be written as

$$Z_T^J = \exp\Big(\sum_{0<u\le T} \zeta_{R_{u-}} Y_u - \int_0^T \lambda(R_u)\{\exp(\psi_{R_u}(\zeta_{R_u})) - 1\}\, du\Big),$$

with ζ_R chosen so that the compensated jump drift is consistent with the martingale condition. The combined density $Z_T = Z_T^W Z_T^J$ is strictly positive and has expectation one under the stated exponential-moment assumptions. The Markov-chain transition kernel is left untilted by construction, so the regime process remains a pricing state rather than a hedgeable risk source. This gives an admissible EMM under which the discounted spot is a true martingale once the stated uniform-integrability condition is imposed. The construction does not claim uniqueness; it selects one economically interpretable pricing measure from the incomplete-market set.

This distinction matters for the paper's scope. The theorem is an existence theorem for a pricing functional in an incomplete market, not a uniqueness theorem for a dynamically complete crypto-options market. Liquidity and regime risk are priced through the loading in Λ, while the hedgeable Brownian and jump components are handled by the Girsanov–Esscher density above.

Appendix A.2. Short-Time Product Condition

The proof of Theorem 2 requires

$$\mathbb{E}^{\mathbb{Q}}[L_u\sqrt{V_u}] = \bar{L}\sigma_0 + O(u^\beta)$$

for some $\beta > 0$. This condition is local and concerns the product entering the liquidity premium, not the full rough path. In the numerical specification, $L_u = \ell_{R_u} + \chi|\Delta B_u^H|/\Delta^H$ has finite moments for each grid step, and the variance process is lognormal before numerical clipping. Because R is finite-state and has bounded transition probabilities over a short interval, $\mathbb{E}[\ell_{R_u}] = \bar{\ell} + O(u)$. The normalized rough increment has a time-stationary Gaussian marginal on the grid, while the lognormal variance factor satisfies $\mathbb{E}[\sqrt{V_u}] = \sigma_0 + O(u^{2H})$ under the deterministic forward-variance baseline. Cauchy–Schwarz and the finite exponential moment in Theorem 1 then give the product expansion with $\beta = \min(1, 2H)$ in the benchmark construction.

For the ATM asymptotic, the condition states that the short-time mean liquidity exposure converges to its initial exposure quickly enough for the leading-order term to dominate the BSM-vega scaling. If a different liquidity proxy is used, this product condition

is the quantity that must be verified; the asymptotic should not be transferred mechanically to a proxy whose short-time product with $\sqrt{V}$ fails this local expansion.

Appendix A.3. Euler Error Decomposition

For every $\gamma < H$, the rough variance path is γ-Hölder on compact intervals. On the numerical state space $V_t \geq V_{\min} > 0$, so $x \mapsto \sqrt{x}$ is Lipschitz. The strong error of the lagged-variance scheme decomposes into

$$\int_{t_{k-1}}^{t_k} \{\sqrt{V_u} - \sqrt{V_{t_{k-1}}}\}\, dW_u$$

and the finite-activity jump discretization term. Itô isometry bounds the diffusion quadrature error by $O(\Delta^{1+2\gamma})$ per step, hence $O(\Delta^{2\gamma})$ over the horizon. Doob's inequality transfers this to the supremum norm. The lagged-variance construction is predictable on each interval $[t_{k-1}, t_k]$, so the discretized asset increment has the correct conditional martingale drift after the jump compensator is included. The compound-Poisson component adds a finite-variation contribution under the jump-moment condition in Theorem 1. The resulting conservative rate is $n^{-\gamma}$ for any $\gamma < H$.

The numerical sensitivity table in Section 3 is not a proof of convergence; it is an implementation check that the benchmark price and IV lift are stable at the path and time-step levels used for the reported simulation. The theorem supplies the conservative rate statement, while the table guards against a reviewer concern that the reported liquidity lift is a grid artifact.

References

1. Black, F.; Scholes, M. The pricing of options and corporate liabilities. *J. Political Econ.* **1973**, *81*, 637–654. [CrossRef] [PubMed]
2. Merton, R.C. Theory of rational option pricing. *Bell J. Econ. Manag. Sci.* **1973**, *4*, 141–183. [CrossRef]
3. Heston, S.L. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.* **1993**, *6*, 327–343. [CrossRef]
4. Bates, D.S. Jumps and stochastic volatility: Exchange rate processes implicit in Deutsche Mark options. *Rev. Financ. Stud.* **1996**, *9*, 69–107. [CrossRef]
5. Cont, R.; Tankov, P. *Financial Modelling with Jump Processes*; Chapman & Hall/CRC: Boca Raton, FL, USA , 2004. [CrossRef]
6. Merton, R.C. Option pricing when underlying stock returns are discontinuous. *J. Financ. Econ.* **1976**, *3*, 125–144. [CrossRef]
7. Bayer, C.; Friz, P.; Gatheral, J. Pricing under rough volatility. *Quant. Financ.* **2016**, *16*, 887–904. [CrossRef]
8. Bayer, C.; Horvath, B.; Muguruza, A.; Stemper, B.; Tomas, M. On deep calibration of (rough) stochastic volatility models. *J. FinTech* **2025**, *5*, 2550005 . [CrossRef]
9. Bennedsen, M.; Lunde, A.; Pakkanen, M.S. Hybrid scheme for Brownian semistationary processes. *Financ. Stoch.* **2017**, *21*, 931–965. [CrossRef]
10. El Euch, O.; Rosenbaum, M. The characteristic function of rough Heston models. *Math. Financ.* **2019**, *29*, 3–38. [CrossRef]
11. Gatheral, J.; Jaisson, T.; Rosenbaum, M. Volatility is rough. *Quant. Financ.* **2018**, *18*, 933–949. [CrossRef]
12. Jacquier, A.; Pakkanen, M.S.; Stone, H. Pathwise large deviations for the rough Bergomi model. *J. Appl. Probab.* **2018**, *55*, 1078–1092. [CrossRef]
13. Atanasova, C.V.; Miao, T.; Segarra, I.; Willeboordse, F. Aggregate illiquidity and crypto option returns. *Financ. Res. Lett.* **2025**, *85*, 108003. [CrossRef]
14. Hou, A.J.; Wang, W.; Chen, C.Y.-H.; Härdle, W.K. Pricing cryptocurrency options. *J. Financ. Econom.* **2020**, *18*, 250–279. [CrossRef]
15. Madan, D.B.; Reyners, S.; Schoutens, W. Advanced model calibration on Bitcoin options. *Digit. Financ.* **2019**, *1*, 117–137. [CrossRef]
16. Saef, D.; Wang, Y.; Aste, T. Regime-based implied stochastic volatility model for crypto option pricing. In *Companion Proceedings of the ACM Web Conference 2023*; Association for Computing Machinery: New York, NY, USA , 2023; pp. 1103–1110. [CrossRef]
17. Scaillet, O.; Treccani, A.; Trevisan, C. High-frequency jump analysis of the Bitcoin market. *J. Financ. Econom.* **2020**, *18*, 209–232. [CrossRef]
18. Carr, P.; Madan, D.B. Option valuation using the fast Fourier transform. *J. Comput. Financ.* **1999**, *2*, 61–73. [CrossRef]
19. Gerhold, S.; Jacquier, A.; Pakkanen, M.; Stone, H.; Wagenhofer, T. Pathwise large deviations for the rough Bergomi model: Corrigendum. *J. Appl. Probab.* **2021**, *58*, 849–850. [CrossRef]

20. Friz, P.K.; Hairer, M. *A Course on Rough Paths: With an Introduction to Regularity Structures*; Springer: Cham, Switzerland, 2014. [CrossRef]
21. Hamilton, J.D. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* **1989**, *57*, 357–384. [CrossRef]
22. Elliott, R.J.; Moore, J.B.; Aggoun, L. *Hidden Markov Models: Estimation and Control*; Springer: New York, NY, USA, 1994. [CrossRef]
23. Hardy, M.R. A regime-switching model of long-term stock returns. *N. Am. Actuar. J.* **2001**, *5*, 41–53. [CrossRef]
24. Alfeus, M.; Overbeck, L.; Schlögl, E. Regime switching rough Heston model. *J. Futur. Mark.* **2019**, *39*, 538–552. [CrossRef]
25. Kyle, A.S. Continuous auctions and insider trading. *Econometrica* **1985**, *53*, 1315–1335. [CrossRef]
26. Glosten, L.R.; Milgrom, P.R. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financ. Econ.* **1985**, *14*, 71–100. [CrossRef]
27. Amihud, Y. Illiquidity and stock returns: Cross-section and time-series effects. *J. Financ. Mark.* **2002**, *5*, 31–56. [CrossRef]
28. Çetin, U.; Jarrow, R.A.; Protter, P. Liquidity risk and arbitrage pricing theory. *Financ. Stoch.* **2004**, *8*, 311–341. [CrossRef]
29. Çetin, U.; Rogers, L.C.G. Modelling liquidity effects in discrete time. *Math. Financ.* **2007**, *17*, 15–29. [CrossRef]
30. Gârleanu, N.; Pedersen, L.H.; Poteshman, A.M. Demand-based option pricing. *Rev. Financ. Stud.* **2009**, *22*, 4259–4299. [CrossRef]
31. Zulfiqar, N.; Gulzar, S. Implied volatility estimation of bitcoin options and the stylized facts of option pricing. *Financ. Innov.* **2021**, *7*, 67. [CrossRef] [PubMed]
32. Matic, J.L.; Packham, N.; Härdle, W.K. Hedging cryptocurrency options. *Rev. Deriv. Res.* **2023**, *26*, 91–133. [CrossRef] [PubMed]
33. Chen, T.; Deng, J.; Nie, J. Implied volatility slopes and jumps in bitcoin options market. *Oper. Res. Lett.* **2024**, *55*, 107135. [CrossRef]
34. Chen, K.S.; Yang, J.J. Price dynamics and volatility jumps in bitcoin options. *Financ. Innov.* **2024**, *10*, 132. [CrossRef]
35. Li, J. Equilibrium pricing of Bitcoin options with stochastic volatility, jumps, and liquidity risk. *J. Futur. Mark.* **2026**, *46*, 221–234. [CrossRef]
36. Karatzas, I.; Shreve, S.E. *Brownian Motion and Stochastic Calculus*, 2nd ed.; Springer: New York, NY, USA, 1991. [CrossRef]
37. Frittelli, M. The minimal entropy martingale measure and the valuation problem in incomplete markets. *Math. Financ.* **2000**, *10*, 39–52. [CrossRef]
38. Deribit. Public API Bitcoin Option Book Summary Snapshot. 2026. Available online: https://docs.deribit.com/ (accessed on 18 May 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.