# Pricing Liquidity Risk in Stablecoins

João Pedro Malim Franco*
IMECC – UNICAMP

March 22, 2026
Preliminary Version

**Abstract**

This paper proposes a no-arbitrage framework for modeling stablecoin prices using a latent liquidity risk premium driven by affine jump–diffusion dynamics. We consider two alternative specifications: a CIR–jumps model, which enforces non-negativity and state-dependent volatility, and a Vasicek–jumps model, which allows for symmetric deviations around parity. The latent liquidity factor and structural parameters are jointly estimated using a Particle Markov Chain Monte Carlo (PMCMC) algorithm within a nonlinear state-space setting. The approach captures both gradual liquidity adjustments and rare but significant de-pegging events. The affine structure ensures analytical tractability for pricing derivatives and computing forward-looking risk measures such as Value at Risk and Expected Shortfall. Overall, the framework provides a structurally consistent and empirically flexible tool for valuation and risk assessment in stablecoin markets.

## 1 Introduction

Stablecoins are digital tokens engineered to maintain a stable value relative to a reference asset, most commonly a fiat currency such as the U.S. dollar, thereby mitigating the extreme price volatility observed in unbacked cryptocurrencies like Bitcoin or Ethereum. By combining the transactional efficiency of blockchain networks with the price stability of sovereign currencies, stablecoins have become a foundational layer of the digital asset ecosystem. They serve multiple functions: as a medium of exchange for on-chain payments and remittances, a store of value within crypto markets during periods of turbulence, a unit of account in decentralized finance (DeFi) protocols, collateral in lending and derivatives markets, and a bridge asset facilitating cross-border settlements and liquidity provision across exchanges. Increasingly, they are also used in emerging markets as a hedge against local currency depreciation and capital controls.

The rise of Stablecoins represents one of the most consequential – and paradoxical – developments in modern digital finance. On the one hand, the market has experienced pronounced expansion, with aggregate capitalization increasing from approximately USD 3 billion in 2019 to over USD 150 billion by mid-2022, and subsequently attaining a combined market capitalization exceeding USD 200 billion by late 2024 (d'Avernas et al., 2022; Eichengreen et al., 2025). On the other hand, this rapid growth has intensified debates regarding reserve transparency, systemic risk, monetary sovereignty, and the evolving interface between private digital money and traditional financial institutions.

Annual transaction volumes exceeded USD 10 trillion in 2023, rivaling global payment incumbents such as Visa and PayPal, and cementing stablecoins as a systemically important component of global

*The author gratefully acknowledges financial support from the São Paulo Research Foundation (FAPESP) under grant number 2025/28660-1.

payment and settlement infrastructure (Bluhm et al., 2024). Traditional firms, including listed companies such as MAIA Biotechnology, have begun to allocate substantial portions of their treasuries to stablecoins such as USDC, using them for liquidity management and risk buffering. This adoption signals their growing integration into conventional financial systems.

Beyond their role in payments and corporate treasuries, a primary catalyst for this expansion has been their indispensable function within Decentralized Finance (DeFi) (Campbell R. Harvey et al., 2021). This ecosystem, which rebuilds traditional financial services such as lending, borrowing, and trading on-chain, requires a stable unit of account and medium of exchange to operate – a role that volatile cryptoassets like Bitcoin or Ethereum cannot fulfill. Stablecoins serve as the money and primary collateral asset for this parallel financial system. They are the most-borrowed asset on lending platforms (e.g., Aave and Compound), the dominant quote currency on decentralized exchanges (e.g., Uniswap), and the foundational asset for complex, low-slippage liquidity pools (e.g., Curve Finance). Consequently, stablecoins act as the core 'plumbing' and risk-management bedrock for the entire on-chain economy, enabling leverage and sophisticated financial engineering (Schär, 2021; Gorton and Zhang, 2023).

Stablecoins achieve the stability through distinct mechanisms, generally classified as collateralized, algorithmic, or hybrid. Each design entails specific operational dynamics but shares a common vulnerability: liquidity risk, the possibility that holders cannot redeem tokens at par value without significant delay or loss. Such risk arises from mismatches between asset supply, redemption demand, and the liquidity of reserves or collateral, and is particularly acute during stress episodes when confidence deteriorates and self-reinforcing redemption spirals emerge.

Fiat-collateralized stablecoins, such as Tether (USDT) and USD Coin (USDC), maintain reserves of fiat currency or near-cash equivalents, typically held in segregated accounts and backed one-to-one by issued tokens. Although this structure offers transparency and regulatory oversight, liquidity risk materializes when reserves include short-term instruments, such as commercial paper or certificates of deposit – that may become illiquid or face counterparty risk in crises. Rapid redemption waves, often triggered by arbitrage or panic, can force asset fire-sales, leading to peg deviations and bank-run dynamics, as evidenced when USDT briefly traded at USD 0.95 during the 2018 Bitfinex banking disruptions. Furthermore, centralized custody exposes these systems to operational and regulatory risks, including the failure of banking partners or government seizures that can freeze fiat outflows.

Crypto-collateralized stablecoins, exemplified by MakerDAO's DAI, rely on over-collateralized positions of volatile cryptoassets locked in smart contracts. Stability is maintained through automated liquidation mechanisms that restore solvency when collateral values fall. While eliminating fiat intermediaries enhances decentralization and censorship resistance, liquidity risk here arises from collateral volatility and network congestion during stress. In the "Black Thursday" event of March 2020, a sharp ETH price collapse overwhelmed MakerDAO's liquidation infrastructure, leaving positions undercollateralized and driving DAI's market price below USD 0.90. These episodes illustrate how endogenous liquidity shortages and systemic interdependencies within decentralized finance (DeFi) amplify fragility.

Algorithmic stablecoins, such as the now-defunct TerraUSD (UST), dispense with collateral altogether, relying on arbitrage incentives and seigniorage mechanisms to adjust supply. Their dependence on continuous market confidence makes them uniquely vulnerable to liquidity "death spirals." When arbitrage capacity evaporates, peg defense fails catastrophically, as seen in UST's collapse to near-zero in May 2022, highlighting the reflexive nature of liquidity risk in unbacked systems. Hybrid models, like Frax, combine fractional collateralization with algorithmic adjustments but inherit compounded vulnerabilities.

Across all designs, liquidity risk interacts with secondary-market dynamics. Stablecoins trade on centralized and decentralized exchanges, where order-book depth and pool balance determine redemption efficiency. While high trading volumes can provide temporary liquidity buffers, stress events expose latent fragilities, including exchange insolvencies (e.g., FTX 2022) or smart-contract exploits draining liquidity pools. High-profile de-pegging events, such as the collapse of the Terra–Luna ecosystem in May 2022 and the temporary dislocation of USD Coin (USDC) during the Silicon Valley Bank (SVB) crisis in March 2023, underscore the inherent instability of these ostensibly "stable" assets (d'Avernas et al., 2022). These disruptions reveal the structural tension between perceived safety and underlying liquidity constraints that define the stablecoin market.

This duality – systemic growth alongside recurrent fragility – has transformed stablecoin stability from a niche technical concern into a first order issue for global macro-financial stability. Regulatory bodies and international institutions, including the Financial Stability Board (FSB), the Bank for International Settlements (BIS), and the International Monetary Fund (IMF), have repeatedly warned of potential contagion channels linking stablecoins to the broader financial system, as well as implications for monetary sovereignty and policy transmission (Wen and Lau, 2025).

In frictionless markets, identical cash flows must trade at identical prices under the Law of One Price. However, when assets differ in their degree of immediacy, convertibility, or market depth, liquidity becomes an economically relevant state variable. In such environments, investors require compensation for bearing liquidity risk, and this compensation appears as a *liquidity risk premium* embedded in asset prices.

This paper develops a structural no-arbitrage framework in which liquidity risk is explicitly modeled as a stochastic state variable governing deviations of stablecoin prices from par. We treat stablecoins as short-duration monetary claims subject to redemption frictions and secondary-market liquidity constraints. Within this setting, departures from the one-dollar peg are not interpreted as anomalies, but as equilibrium price adjustments reflecting time-varying compensation for liquidity risk.

Our baseline specification models the liquidity premium as a Cox–Ingersoll–Ross (CIR) jump-diffusion process under the risk-neutral measure. The square-root diffusion component ensures non-negativity of the liquidity spread, consistent with the interpretation of liquidity risk as an additive premium to the short rate. The jump component captures abrupt liquidity shocks arising from bank failures, exchange insolvencies, regulatory announcements, or smart-contract disruptions. This structure yields an affine term-structure representation, allowing closed-form or semi-analytical solutions for stablecoin prices, forward contracts, and European-style derivatives written on the stablecoin. The exponential-affine form ensures analytical tractability while preserving economically meaningful state dependence.

Specifically, we derive pricing formulas for: (i) the spot value of a redeemable stablecoin as a discounted claim subject to stochastic liquidity spreads; (ii) futures contracts reflecting expectations of future liquidity conditions; and (iii) European options whose valuation incorporates both diffusive and jump-driven liquidity risks. The resulting formulas extend classical affine term-structure techniques to digital monetary instruments operating in hybrid on-chain/off-chain environments. Importantly, the model generates endogenous peg deviations, convexity effects, and implied volatility patterns consistent with observed de-pegging episodes.

To recover the latent liquidity factor from market data, we implement a sequential Monte Carlo (particle filtering) methodology. Because liquidity risk is not directly observable, its estimation requires filtering techniques capable of handling non-linear and non-Gaussian state dynamics induced by jump components and affine transformations. The particle filter provides a flexible framework for estimating the dynamic trajectory of liquidity risk using observed stablecoin prices, derivatives data, and trading volume information. This approach enables real-time monitoring of liquidity stress and allows decomposition of observed peg deviations into systematic and jump-driven components.

In addition to the CIR jump-diffusion specification, we propose an alternative Vasicek jump-diffusion formulation. While the CIR framework imposes non-negativity of the liquidity premium, the Vasicek specification permits negative liquidity spreads. This extension is economically relevant in periods when stablecoins trade persistently above par—reflecting scarcity premia, collateral demand in DeFi protocols, or flight-to-safety dynamics within crypto markets. By allowing the liquidity premium to become negative, the model captures episodes in which investors are willing to pay a premium for immediate on-chain dollar liquidity, effectively treating stablecoins as convenience-yield-bearing monetary assets.

Together, these specifications provide a unified asset-pricing framework capable of capturing both downside liquidity spirals and upside scarcity premia. The affine jump-diffusion structure ensures that both models remain analytically tractable and suitable for empirical implementation.

Methodologically, our contributions are fourfold. First, we introduce a no-arbitrage pricing framework that embeds liquidity risk directly into the stochastic discount factor governing stablecoin valuation. Second, we derive semi-analytical pricing formulas for stablecoins and their derivatives under affine jump-diffusion dynamics. Third, we implement particle filtering techniques to estimate latent liquidity states and jump intensities in real time. Fourth, we demonstrate how alternative diffusion specifications alter the sign and interpretation of liquidity premia, thereby accommodating both depegging crises and scarcity-driven premiums above par.

From a market perspective, the framework provides practical tools for risk management, derivative pricing, stress testing, and regulatory surveillance. Exchanges and DeFi protocols can use the model to price options and structured products referencing stablecoins. Institutional treasuries holding large stablecoin reserves can quantify liquidity-adjusted value-at-risk and jump exposure. Regulators can interpret estimated liquidity premia as forward-looking indicators of systemic fragility, analogous to credit spreads in traditional financial markets.

More broadly, the paper contributes to the literature at the intersection of monetary economics, market microstructure, and digital asset pricing. By treating stablecoins as monetary instruments subject to stochastic liquidity conditions, we bridge classical term-structure modeling with decentralized financial architecture. The resulting framework rationalizes peg deviations within a coherent equilibrium setting and provides a tractable foundation for the quantitative analysis of liquidity risk in tokenized monetary systems.

# 2 Methodology

## 2.1 No-Arbitrage Pricing under a CIR Jump-Diffusion Liquidity Premium

This section develops a no-arbitrage asset-pricing framework (Björk, 2019; Musiela and Rutkowski, 2005) to quantify and price liquidity risk in stablecoins. The use of a no-arbitrage approach is motivated by the institutional nature of stablecoins as redeemable monetary claims with well-defined payoff structures. In equilibrium, a stablecoin that promises convertibility into one U.S. dollar represents a short-maturity contingent claim whose value must equal the discounted expectation of its redemption payoff under an equivalent martingale measure. Any persistent deviation from par therefore reflects either compensation for risk or market frictions. By embedding liquidity risk directly into the stochastic discount factor, we interpret peg deviations as equilibrium outcomes arising from time-varying liquidity premia rather than pricing anomalies.

A no-arbitrage framework is particularly suitable in this context for three reasons. First, stablecoins are actively traded across centralized exchanges, decentralized protocols, and derivative markets, creating cross-market restrictions that require internally consistent pricing. Second, arbitrage mechanisms – minting and redemption facilities, cross-exchange trading, and on-chain collateral adjustments – impose structural links between spot prices, forward prices, and derivative valuations. Third, liquidity shocks propagate through both primary (redemption) and secondary (trading) markets, implying that liquidity risk affects discounting in a manner analogous to short-rate spreads or convenience yields in traditional term-structure models.

Modeling liquidity risk as a stochastic state variable within an affine no-arbitrage setting therefore allows us to (i) derive semi-analytical valuation formulas for stablecoins and their derivatives, (ii) ensure consistency across maturities and instruments, and (iii) recover forward-looking liquidity premia from market data. This approach transforms observed peg deviations into economically interpretable risk prices and provides a tractable foundation for empirical estimation and policy analysis.

We first model the liquidity premium as a Cox–Ingersoll–Ross (CIR) square-root diffusion augmented with jumps (Cox et al., 1985). The choice of the CIR structure is motivated by both economic and econometric considerations. First, liquidity risk in stablecoin markets is naturally non-negative when interpreted as an additive spread over the risk-free rate. During stress episodes, such as redemption waves, exchange insolvencies, or banking disruptions, the required compensation for immediacy and convertibility risk increases, generating positive liquidity spreads and downward pressure on the peg. The square-root diffusion of the CIR process ensures non-negativity of the state variable while preserving stochastic variability. This property is economically consistent with the interpretation of liquidity risk as a premium demanded by investors for bearing redemption and market-depth uncertainty.

Second, liquidity risk exhibits strong mean-reversion. Empirically, de-pegging episodes are typically short-lived: arbitrageurs restore parity once balance sheet constraints relax and market depth recovers. The mean-reverting drift of the CIR process captures this stabilizing mechanism, reflecting the institutional design of stablecoins in which redemption facilities and collateral adjustments anchor long-run value to par. The speed-of-adjustment parameter therefore has a direct economic interpretation as the efficiency of arbitrage and reserve management.

Third, liquidity volatility is state-dependent. Periods of elevated liquidity stress are typically accompanied by heightened volatility and jump risk. The square-root diffusion term implies that the conditional variance of the liquidity premium increases with its level, generating endogenous heteroskedasticity. This feature is crucial for reproducing the clustering of volatility observed during systemic events such as banking crises or exchange failures. Linear Gaussian models, by contrast, impose constant variance and may understate tail risk during stress regimes.

Fourth, the CIR process belongs to the affine class of term-structure models, which ensures analytical tractability. Within an affine framework, bond prices, forwards, and European options admit exponential-affine representations, reducing valuation to the solution of Riccati equations. This tractability allows us to derive semi-analytical pricing formulas for stablecoins and their derivatives while preserving structural interpretability. The affine property also facilitates filtering and likelihood-based estimation, particularly when combined with jump components.

Finally, augmenting the CIR diffusion with a jump process accommodates rare but severe liquidity dislocations. Stablecoin markets are exposed to discrete events, as regulatory announcements, custodial failures, smart-contract exploits, or bank insolvencies – that generate abrupt repricing. The jump component captures these discontinuities without sacrificing the closed-form structure of the affine model.

Taken together, the CIR jump-diffusion specification provides a parsimonious yet economically coherent representation of liquidity risk: it enforces non-negativity, incorporates mean reversion consistent with arbitrage-based peg stabilization, generates state-dependent volatility, accommodates tail events, and preserves analytical tractability for derivative pricing and empirical estimation. These properties make it particularly well-suited for modeling liquidity premia in stablecoin markets operating at the intersection of monetary claims and decentralized financial infrastructure.

Under this framework, we model the liquidity risk premium $x_t$ of USDC under the risk-neutral measure $\mathbb{Q}$ as a Cox–Ingersoll–Ross (CIR) diffusion with compound Poisson jumps:

$$dx_t \;=\; \kappa(\theta - x_t)\,dt \;+\; \sigma\sqrt{x_t}\,dW_t^{\mathbb{Q}} \;+\; dJ_t, \tag{1}$$

where $W^{\mathbb{Q}}$ is a $\mathbb{Q}$-Brownian motion and the jump process is

$$J_t \;=\; \sum_{i=1}^{N_t} Y_i,$$

with $N_t \sim \text{Poisson}(\lambda_J t)$ and i.i.d. exponential jump sizes $Y_i \sim \text{Exp}(\beta_J)$. Parameters satisfy $\kappa, \theta, \sigma, \lambda_J, \beta_J > 0$. This affine jump-diffusion specification captures mean-reversion, state-dependent variance, and sudden liquidity shocks (jumps). See Cox et al. (1985) and the affine-jump literature (Duffie et al., 2000, 2003) for theoretical background.

To establish an arbitrage-free valuation, consider a contingent claim that pays one unit of the numéraire, specifically United States Dollars (USD), at the maturity time $T$. Under $\mathbb{Q}$, the arbitrage-free price is

$$P(t,T) \;=\; \mathbb{E}_t^{\mathbb{Q}}\left[\exp\Big(-\int_t^T (r + x_s)\, ds\Big)\right], \tag{2}$$

where $r$ is the (constant) USD risk-free rate. We seek an affine representation

$$P(t,T) \;=\; \exp\big[-r\tau + A(\tau) - B(\tau)\, x_t\big], \qquad \tau := T - t, \tag{3}$$

with deterministic functions $A(\tau)$ and $B(\tau)$ satisfying Riccati-type ordinary differential equations (ODEs) obtained by the Feynman–Kac argument (see Karatzas and Shreve (1991); Filipović (2009)).

## 2.2 Riccati system and Laplace transform of jumps

Recall that in time-to-maturity $\tau = T - t$, the pricing Partial Differential Equation (PDE) reads (Cont and Tankov, 2003; Filipović, 2009):

$$-\frac{\partial P}{\partial \tau} + \kappa(\theta - x)\frac{\partial P}{\partial x} + \frac{1}{2}\sigma^2 x \frac{\partial^2 P}{\partial x^2} + \lambda_J \int_0^\infty \big[P(\tau, x+y) - P(\tau, x)\big] f_Y(y)\, dy - (r+x)P = 0, \tag{4}$$

with initial condition $P(0,x) = 1$.

Under exponential–affine form the derivatives are

$$\frac{\partial P}{\partial \tau} = \big(-r + A'(\tau) - B'(\tau)x\big)P. \tag{5}$$

$$\frac{\partial P}{\partial x} = -B(\tau)P. \tag{6}$$

$$\frac{\partial^2 P}{\partial x^2} = B(\tau)^2 P. \tag{7}$$

For the contribution arising from the jump term, we compute

$$\begin{aligned} P(\tau, x+y) &= \exp\big[-r\tau + A(\tau) - B(\tau)(x+y)\big] && (8)\\ &= P(\tau, x)\, e^{-B(\tau)y}. && (9) \end{aligned}$$

Therefore,

$$\int_0^\infty \big[P(\tau, x+y) - P(\tau, x)\big] f_Y(y)\, dy = P(\tau, x)\left(\mathbb{E}[e^{-B(\tau)Y}] - 1\right). \tag{10}$$

Substituting all terms into (4):

$$-\big(-r+A'-B'x\big)P+\kappa(\theta-x)(-BP)+\frac{1}{2}\sigma^2xB^2P+\lambda_JP\big(\mathbb{E}[e^{-BY}]-1\big)-(r+x)P=0. \tag{11}$$

Factor out $P(\tau,x)>0$:

$$-(-r+A'-B'x)-\kappa(\theta-x)B+\frac{1}{2}\sigma^2xB^2+\lambda_J\big(\mathbb{E}[e^{-BY}]-1\big)-(r+x)=0. \tag{12}$$

Simplify the first term:

$$-(-r+A'-B'x)=r-A'+B'x.$$

Hence,

$$r-A'+B'x-\kappa\theta B+\kappa xB+\frac{1}{2}\sigma^2xB^2+\lambda_J(\mathbb{E}[e^{-BY}]-1)-r-x=0. \tag{13}$$

The constant $r$ cancels. Rearranging:

$$\big(-A'-\kappa\theta B+\lambda_J(\mathbb{E}[e^{-BY}]-1)\big)+x\big(B'+\kappa B+\frac{1}{2}\sigma^2B^2-1\big)=0. \tag{14}$$

Since the equality must hold for all $x$, the coefficients of 1 and $x$ must vanish separately:

$$B'(\tau)=1-\kappa B(\tau)-\frac{1}{2}\sigma^2B(\tau)^2, \tag{15}$$

with initial condition $B(0)=0$ and

$$A'(\tau)=-\kappa\theta B(\tau)+\lambda_J\big(\mathbb{E}[e^{-B(\tau)Y}]-1\big), \tag{16}$$

with $A(0)=0$.

Therefore, plugging (3) into the pricing PDE yields the coupled system

$$\frac{dB}{d\tau}(\tau)=1-\kappa B(\tau)-\tfrac{1}{2}\sigma^2B(\tau)^2, \qquad B(0)=0, \tag{17}$$

$$\frac{dA}{d\tau}(\tau)=-\kappa\theta B(\tau)+\lambda_J\big(\mathbb{E}[e^{-B(\tau)Y}]-1\big), \qquad A(0)=0. \tag{18}$$

Since $Y\sim\mathrm{Exp}(\beta_J)$, its Laplace transform is

$$\mathbb{E}[e^{-uY}] \;=\; \frac{\beta_J}{\beta_J+u}, \qquad u\geq 0, \tag{19}$$

and therefore

$$\frac{dA}{d\tau}(\tau) \;=\; -\kappa\theta B(\tau)+\lambda_J\left(\frac{\beta_J}{\beta_J+B(\tau)}-1\right). \tag{20}$$

The Riccati ODE (17) admits the standard CIR solution (see Cox et al. (1985); Duffie et al. (2000)). Define

$$\gamma \;:=\; \sqrt{\kappa^2 + 2\sigma^2}\,. \tag{21}$$

Then

$$B(\tau) \;=\; \frac{2\big(e^{\gamma\tau}-1\big)}{(\gamma+\kappa)\big(e^{\gamma\tau}-1\big)+2\gamma}\,. \tag{22}$$

One checks that $B(0) = 0$ and that $B(\tau) > 0$ for $\tau > 0$[1]. Thus, for this term we have an analytical formula.

For the diffusion-only CIR (no jumps) the classical expression for $A(\tau)$ is

$$A_{\mathrm{CIR}}(\tau) \;=\; \frac{2\kappa\theta}{\sigma^2}\,\ln\!\left(\frac{2\gamma e^{(\kappa+\gamma)\tau/2}}{(\gamma+\kappa)(e^{\gamma\tau}-1)+2\gamma}\right), \tag{23}$$

which satisfies $\frac{dA_{\mathrm{CIR}}}{d\tau} = -\kappa\theta B(\tau)$. With jumps, using (20), $A(\tau)$ decomposes as

$$A(\tau) \;=\; A_{\mathrm{CIR}}(\tau) \;+\; \lambda_J \int_0^{\tau} \left(\frac{\beta_J}{\beta_J + B(s)} - 1\right) ds. \tag{24}$$

The first term is given in closed form by (23); the second term captures the cumulative effect of exponential jumps and generally must be evaluated numerically because $B(s)$ is a nontrivial function of $s$. Note that the jump integral can also be written as

$$\lambda_J \int_0^{\tau} \left(\frac{\beta_J}{\beta_J + B(s)} - 1\right) ds = -\lambda_J\tau + \lambda_J\beta_J \int_0^{\tau} \frac{ds}{\beta_J + B(s)}.$$

A practical and stable Numerical scheme for computing $A(\tau)$ is:

1. Evaluate $B(s)$ for $s \in [0,\tau]$ using the closed-form expression (22) on a fine grid $0 = s_0 < s_1 < \cdots < s_N = \tau$.
2. Compute the integral in (24) using an adaptive quadrature rule (e.g., Gauss–Kronrod) or composite Simpson's rule:
$$I(\tau) \;=\; \int_0^{\tau} \left(\frac{\beta_J}{\beta_J + B(s)} - 1\right) ds \approx \sum_{n=0}^{N-1} w_n \left(\frac{\beta_J}{\beta_J + B(s_n)} - 1\right),$$
where $w_n$ are the quadrature weights.
3. Finally set $A(\tau) = A_{\mathrm{CIR}}(\tau) + \lambda_J I(\tau)$.

This procedure is numerically inexpensive because evaluating $B(s)$ by (22) is direct and smooth; the integrand $\frac{\beta_J}{\beta_J + B(s)} - 1$ is well-behaved for $B(s) \geq 0$ (no singularity as $\beta_J > 0$). For high accuracy use adaptive quadrature that refines near $s \approx 0$ if needed.

[1] See Cox et al. (1985).

## 2.3 Implied liquidity premium and inversion

Using (3) and (24), the model-implied price is

$$P(t,T) \;=\; \exp\bigl[-r\tau + A_{\mathrm{CIR}}(\tau) - B(\tau)x_t \;+\; \lambda_J \int_0^{\tau} \bigl(\tfrac{\beta_J}{\beta_J+B(s)} - 1\bigr)\, ds\bigr]. \tag{25}$$

Given observed market prices $P^{\mathrm{obs}}(t,T)$, the latent liquidity state $x_t$ can be inverted (when other parameters are known) from

$$x_t \;=\; \frac{A(\tau) - r\tau - \ln P^{\mathrm{obs}}(t,T)}{B(\tau)}. \tag{26}$$

This yields a time series of implied liquidity premia that can be studied empirically. When parameters are unknown, expression (25) provides the model likelihood for calibration. We discuss a particle-MCMC approach to estimate parameters and liquidity risk in Section 4.

# 3 Derivative Pricing under the CIR–Jump Liquidity Model

Having characterized the dynamics of the liquidity risk factor under the risk-neutral measure, we now turn to the pricing of derivatives written on the stablecoin. In our framework, deviations from par are driven by a latent liquidity state variable $x_t$ that follows a Cox–Ingersoll–Ross (CIR) process augmented with positive jumps. This specification captures both the persistent component of liquidity stress and the abrupt dislocations frequently observed during redemption waves, exchange outages, or collateral reallocation events.

Because the spot value of the stablecoin with settlement delay $\tau_s$ admits an exponential–affine representation in the state variable, the model belongs to the class of affine term structure models. This structure is crucial for tractability: conditional expectations of exponentially affine functions of $x_t$ remain exponentially affine. Consequently, futures and option prices can be expressed semi-analytically through solutions to Riccati differential equations.

We first derive a closed-form expression for the futures price, defined as the risk-neutral expectation of the future settlement value. The affine property implies that the pricing problem reduces to evaluating the conditional Laplace transform of the jump–CIR state variable. Despite the presence of Poisson jumps, this transform remains exponentially affine, with coefficients solving a modified Riccati system.

We then extend the analysis to European options written on the stablecoin. Since the terminal log-price is affine in the state variable, its characteristic function is available in closed form. This permits the use of Fourier inversion methods, such as the Carr–Madan (Carr and Madan, 1999) representation, to obtain numerically stable option prices. The resulting formulas highlight how persistent liquidity risk (through the CIR diffusion) and sudden liquidity shocks (through jumps) jointly determine the shape of the implied volatility surface and the pricing of tail risk.

Together, these results provide a coherent no-arbitrage framework for pricing stablecoin derivatives under endogenous liquidity risk, linking microstructural redemption frictions to observable derivative premia.

## 3.1 Semi-Analytical Futures Pricing Formula

The time-$t$ futures price for settlement at time $T$ is defined as the risk-neutral expectation of the future spot value of the liquidity-adjusted asset:

$$F(t,T) = \mathbb{E}_t^{\mathbb{Q}}[P(T, T+\tau_s)], \tag{27}$$

where $\tau_s$ denotes the spot settlement delay (e.g., one day).

Using the affine form of $P(\cdot,\cdot)$, this expectation can be written as

$$F(t,T) = \exp(-r\tau_s + A(\tau_s))\, \mathbb{E}_t^{\mathbb{Q}}[\exp(-B(\tau_s)\, x_T)]\,. \tag{28}$$

The conditional Laplace transform

$$M(u,\Delta) = \mathbb{E}_t^{\mathbb{Q}}[\exp(-u\, x_{t+\Delta}) \mid x_t]$$

is also affine and satisfies its own system of Riccati equations with initial condition $B(0) = u$:

$$\frac{dB}{d\tau} = -\kappa B(\tau) - \tfrac{1}{2}\sigma^2 B(\tau)^2, \quad B(0) = u, \tag{29}$$

$$\frac{dA}{d\tau} = -\kappa\theta B(\tau) + \lambda_J\left(\frac{\beta_J}{\beta_J + B(\tau)} - 1\right), \quad A(0) = 0. \tag{30}$$

Solving this system over $[0,\Delta]$, with $\Delta = T - t$, yields $A(\Delta; u)$ and $B(\Delta; u)$, and therefore

$$M(u,\Delta) = \exp\big(A(\Delta; u) - B(\Delta; u)\, x_t\big). \tag{31}$$

Combining these expressions, the semi-analytical futures price is:

$$F(t,T) = \exp\Big( - r\tau_s + A(\tau_s) + A(\Delta; B(\tau_s)) - B(\Delta; B(\tau_s))\, x_t\Big). \tag{32}$$

In practice, $A(\tau)$ is obtained via numerical quadrature (trapezoidal integration), and the Riccati systems are integrated using a Runge–Kutta scheme (as implemented in `scipy.integrate.solve_ivp`). This semi-analytical procedure preserves numerical stability and efficiency, while the affine structure guarantees positivity of $x_t$ and closed-form tractability of conditional expectations.

### 3.2 European option pricing

We derive a no-arbitrage price for a European call option written on the stablecoin, maturing at $T$ with strike $K$. Let $t$ denote the pricing time and $\Delta = T - t$. As in Section 2.1 the spot (settlement) value of the stablecoin at time $s$ with settlement delay $\tau_s$ is given by

$$S_s = P(s, s+\tau_s) = \exp\big( - r\tau_s + A(\tau_s) - B(\tau_s)x_s\big) = \mathcal{M}\, e^{-B_s x_s}, \tag{33}$$

where $\mathcal{M} := \exp(-r\tau_s + A(\tau_s))$ and $B_s := B(\tau_s) > 0$. Hence,

$$S_T = \mathcal{M}e^{-B_s x_T}.$$

The arbitrage-free price of the call is

$$C(t) = e^{-r\Delta}\, \mathbb{E}_t^{\mathbb{Q}}\big[(\mathcal{M}e^{-B_s x_T} - K)^+\big]\,. \tag{34}$$

Define

$$x^* = \frac{1}{B_s} \ln\left(\frac{\mathcal{M}}{K}\right).$$

Then

$$C(t) = e^{-r\Delta} \int_0^{x^*} (\mathcal{M}e^{-B_s x} - K) f_{x_T|t}(x)\, dx.$$

The conditional Moment Generating Function (MGF) of $x_T$ has affine form

$$\mathbb{E}_t^{\mathbb{Q}}[e^{ux_T}] = \exp\Big(\widetilde{A}(\Delta; u) + \widetilde{B}(\Delta; u)x_t\Big),$$

where $(\widetilde{A}, \widetilde{B})$ solve

$$\frac{d\widetilde{B}}{ds} = \frac{1}{2}\sigma^2 \widetilde{B}^2 - \kappa \widetilde{B}, \qquad \widetilde{B}(0; u) = u, \tag{35}$$

$$\frac{d\widetilde{A}}{ds} = \kappa\theta\, \widetilde{B} + \lambda_J \left(\mathbb{E}[e^{\widetilde{B}J}] - 1\right), \qquad \widetilde{A}(0; u) = 0. \tag{36}$$

$$\mathbb{E}[e^{\widetilde{B}J}] = \frac{\beta_J}{\beta_J - \widetilde{B}}, \quad \text{for } \Re(\widetilde{B}) < \beta_J,$$

$$\frac{d\widetilde{A}}{ds} = \kappa\theta\, \widetilde{B} + \lambda_J \left(\frac{\beta_J}{\beta_J - \widetilde{B}} - 1\right).$$

Let

$$L_T = \ln S_T = \ln \mathcal{M} - B_s x_T.$$

Then, the conditional characteristic function is

$$\varphi_{L_T|t}(v) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{ivL_T}\right] = e^{iv \ln \mathcal{M}} \exp\Big(\widetilde{A}(\Delta; -ivB_s) + \widetilde{B}(\Delta; -ivB_s)x_t\Big). \tag{37}$$

### 3.2.1 Fourier pricing

The pricing formula follows from the standard Fourier-transform approach of Carr and Madan (1999), which exploits the availability of the conditional characteristic function of the log-payoff under the risk-neutral measure. Since the call payoff is not square-integrable, we introduce an exponential damping factor $e^{\alpha k}$ with $\alpha > 0$ to ensure integrability in Fourier space.

Denoting by $k = \ln K$ the log-strike and by $\varphi_{L_T|t}(u)$ the conditional characteristic function of the log-price (or log-liquidity-adjusted payoff) at maturity, conditional on the state at time $t$, the damped call price admits a Fourier representation obtained via inversion of the characteristic function and application of Fubini's theorem. The denominator $(\alpha + iv)(\alpha + iv + 1)$ arises from the Fourier transform of the damped payoff function. For damping parameter $\alpha > 0$,

$$C(t; K) = \frac{e^{-\alpha k}}{\pi} \int_0^\infty \mathrm{Re}\left\{ e^{-ivk} \frac{\varphi_{L_T|t}(v - i(\alpha + 1))}{(\alpha + iv)(\alpha + iv + 1)} \right\} dv. \tag{38}$$

In the present CIR jump-diffusion setting, the characteristic function is exponentially affine in the state variable due to the affine structure of the model. Specifically, using the conditional transform

$$\varphi_{L_T|t}(u) = \exp\big(\widetilde{A}(\Delta; uB_s) + \widetilde{B}(\Delta; uB_s)\, x_t + iu \ln \mathcal{M}\big),$$

where $\widetilde{A}$ and $\widetilde{B}$ solve the associated Riccati system and $\Delta = T - t$, substitution into the Fourier representation yields the semi-analytical pricing formula

$$C(t;K) = \frac{e^{-\alpha k}}{\pi} \int_0^\infty \mathrm{Re}\Bigg\{ e^{-ivk+i(v-i(\alpha+1))\ln \mathcal{M}} \tag{39}$$

$$\times \exp\Big( \widetilde{A}\big(\Delta; -i(v-i(\alpha+1))B_s\big) + \widetilde{B}\big(\Delta; -i(v-i(\alpha+1))B_s\big) x_t \Big) \tag{40}$$

$$\times \big[ (\alpha+iv)(\alpha+iv+1) \big]^{-1} \Bigg\} dv. \tag{41}$$

This representation reduces option pricing to a one-dimensional Fourier integral, which can be evaluated efficiently either by direct numerical integration or by Fast Fourier Transform (FFT) techniques.

# 4 Estimation Process for Parameters and Latent Liquidity Risk

## 4.1 Estimation via Particle Markov Chain Monte Carlo

We estimate jointly the structural parameters $\Theta = (\kappa, \theta, \sigma, \lambda_J, \beta_J)$ and the latent state process $\{x_t\}_{t=1}^T$ using a Particle Markov Chain Monte Carlo (PMCMC) methodology. Because the model features a nonlinear and non-Gaussian state-space structure induced by the square-root diffusion and jump component, the likelihood function is not available in closed form.

We therefore rely on a Sequential Monte Carlo (SMC) particle filter to construct an unbiased estimator of the likelihood, which is embedded within a Markov Chain Monte Carlo scheme following Andrieu et al. (2010). In particular, we implement a Particle Marginal Metropolis–Hastings (PMMH) algorithm in which each proposed parameter vector $\Theta^{(m)}$ is evaluated using a bootstrap particle filter to approximate the filtering distribution $p(x_{1:T} \mid \Theta^{(m)}, y_{1:T})$ and the marginal likelihood $p(y_{1:T} \mid \Theta^{(m)})$. This construction ensures that the Markov chain targets the exact joint posterior distribution $p(\Theta, x_{1:T} \mid y_{1:T})$ despite relying on Monte Carlo likelihood estimates. The resulting procedure allows for full Bayesian inference on both static parameters and the latent factor path, while properly accounting for parameter uncertainty and filtering uncertainty in a coherent probabilistic framework.

In order to account for pricing frictions and market microstructure effects, we augment the state-space system with an explicit measurement error component. In practice, observed stablecoin prices may exceed parity (i.e., trade above one dollar) due to temporary liquidity imbalances, exchange-specific demand pressures, order book depth limitations, or settlement frictions. Such deviations are not necessarily fully captured by the latent liquidity factor $x_t$, which is designed to describe the underlying structural liquidity conditions rather than transitory trading noise.

Accordingly, we specify the observation equation as

$$y_t = \mathcal{H}(x_t; \Theta) + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), \tag{42}$$

where $\mathcal{H}(x_t; \Theta)$ denotes the model-implied price (or spread) consistent with the risk-neutral valuation framework, and $\varepsilon_t$ captures measurement error. This additive disturbance absorbs short-lived price dislocations, bid–ask bounce, exchange heterogeneity, asynchronous trading, and other microstructure effects that are not explicitly modeled within the structural dynamics.

The inclusion of measurement error is particularly important in the presence of over-par observations. Without $\varepsilon_t$, the filtering step would attempt to rationalize every upward price deviation through extreme realizations of the latent state $x_t$, potentially inducing upward bias in the estimated volatility parameter $\sigma$ or in the jump intensity $\lambda_J$. By allowing for observation noise, the model separates persistent liquidity risk dynamics from purely transitory pricing noise, thereby stabilizing parameter inference.

From a statistical perspective, the measurement error ensures robustness of the likelihood evaluation within the particle filter. In nonlinear state-space models with sharp observation equations, even small model misspecifications can lead to particle degeneracy. The Gaussian disturbance smooths the likelihood surface and improves the numerical stability of the Sequential Monte Carlo approximation.

Economically, this specification recognizes that stablecoin prices reflect both structural liquidity conditions and short-run trading frictions. The latent CIR jump-diffusion factor captures systematic liquidity risk and de-pegging pressure, while the measurement error isolates transient deviations from parity. This decomposition is essential for obtaining structurally interpretable parameter estimates and for preventing the jump component from spuriously absorbing high-frequency market noise.

#### 4.1.1 Bayesian Estimation via Particle–Metropolis–Hastings

The latent liquidity risk premium $x_t$ in our model follows a Cox–Ingersoll–Ross (CIR) diffusion with Poisson jumps, as defined in Section 2.1. The transition density of such a process is not analytically tractable, since the convolution between the noncentral chi-square law (from the CIR diffusion) and the compound Poisson–exponential law (from the jump component) lacks a closed form. Furthermore, the latent state $x_t$ is not directly observable. Instead, we observe a sequence of prices given by

$$P_t = \exp(-r\tau + A(\tau) - B(\tau)x_t) + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2_{\text{obs}}), \tag{43}$$

which introduces an additional layer of measurement uncertainty.

This setting leads to a nonlinear, non-Gaussian state-space model, for which classical maximum likelihood or Kalman-based inference methods are not applicable. To perform likelihood-based Bayesian estimation in this context, we employ the Particle Metropolis–Hastings (PMMH) algorithm, originally proposed by Andrieu et al. (2010), which combines a Sequential Monte Carlo (SMC) particle filter with a Metropolis–Hastings (MH) parameter update. This approach is particularly suitable for jump–diffusion models such as ours, as it allows unbiased estimation of the likelihood function, ensuring asymptotically exact Bayesian inference despite the use of Monte Carlo approximation.

#### 4.1.2 Particle Filter for the CIR–Jump Process model

The particle filter approximates the filtering distribution of the latent state $x_t$ conditional on observed prices $P_{1:t}$ by a set of weighted particles $\{x_t^{(i)}, w_t^{(i)}\}_{i=1}^N$. At each time step, the algorithm performs three main operations:

1. **Prediction:** Given particles $\{x_{t-1}^{(i)}\}$ at time $t-1$, new particles are propagated forward according to the transition dynamics of the CIR–Jump process:

$$x_t^{(i)} = f(x_{t-1}^{(i)}) + J_t^{(i)}, \qquad J_t^{(i)} = \sum_{k=1}^{N_t^{(i)}} Y_k^{(i)}, \tag{44}$$

   where $N_t^{(i)} \sim \text{Poisson}(\lambda_J \Delta t)$ and $Y_k^{(i)} \sim \text{Exp}(\beta_J)$. The diffusion term is simulated using the exact transition of the CIR process or a moment-matched normal approximation:

$$x_t^{(i)} \approx \max\Big(\mathcal{N}\big(\mu_{t|t-1}, \sigma^2_{t|t-1}\big),\, 0\Big),$$

   where $\mu_{t|t-1}$ and $\sigma^2_{t|t-1}$ depend on $(\kappa, \theta, \sigma)$.

2. **Weighting:** Each particle receives a weight proportional to the observation likelihood:

$$w_t^{(i)} \propto p(P_t \mid x_t^{(i)}) = \frac{1}{\sqrt{2\pi\sigma_{\text{obs}}^2}} \exp\left[ -\frac{(P_t - \hat{P}_t^{(i)})^2}{2\sigma_{\text{obs}}^2} \right], \tag{45}$$

   with $\hat{P}_t^{(i)} = \exp(-r\tau + A(\tau) - B(\tau)x_t^{(i)})$.

3. **Resampling:** To avoid weight degeneracy, the particles are resampled according to their normalized weights.

The resulting estimate of the log-likelihood is given by

$$\log \hat{p}(P_{1:T} \mid \theta) = \sum_{t=1}^{T} \log \left( \frac{1}{N} \sum_{i=1}^{N} w_t^{(i)} \right), \tag{46}$$

where $\hat{p}(P_{1:T} \mid \theta)$ is an unbiased estimator of the likelihood $p(P_{1:T} \mid \theta)$ produced by the particle filter (Pitt and Shephard, 1999; Doucet et al., 2000).

### 4.1.3 Metropolis–Hastings Parameter Update

Given the unbiased likelihood estimate from the particle filter, the PMMH algorithm proceeds as a standard Metropolis–Hastings sampler over the parameter vector

$$\boldsymbol{\Theta} = (\phi, \theta, \sigma, \lambda_J, \beta_J),$$

where $\phi = e^{-\kappa \Delta t}$ is a reparameterization ensuring stationarity. At each iteration $m$, a proposal $\boldsymbol{\Theta}'$ is drawn from a mixed random-walk kernel:

$$\phi' = \operatorname{logit}^{-1}\Big( \operatorname{logit}(\phi^{(m-1)}) + \eta_\phi \Big), \tag{47}$$

$$\log \Theta_j' = \log \Theta_j^{(m-1)} + \eta_j, \quad j \in \{\theta, \sigma, \lambda_J, \beta_J\}, \tag{48}$$

with $\eta_\phi \sim \mathcal{N}(0, s_\phi^2)$ and $\eta_j \sim \mathcal{N}(0, s_j^2)$. The acceptance probability is computed as

$$\alpha = \min \left\{ 1, \ \frac{\hat{p}(P_{1:T} \mid \boldsymbol{\Theta}')\, p(\boldsymbol{\Theta}')\, q(\boldsymbol{\Theta}^{(m-1)} \mid \boldsymbol{\Theta}')}{\hat{p}(P_{1:T} \mid \boldsymbol{\Theta}^{(m-1)})\, p(\boldsymbol{\Theta}^{(m-1)})\, q(\boldsymbol{\Theta}' \mid \boldsymbol{\Theta}^{(m-1)})} \right\}, \tag{49}$$

where $\hat{p}(\cdot)$ denote the likelihood, $p(\cdot)$ denote the prior distribution and $q(\cdot)$ represent the proposal distribution. The vectors $\boldsymbol{\Theta}'$ and $\boldsymbol{\Theta}^{(m-1)}$ denote, respectively, the proposed parameter values and the current values of the Markov chain. Equation 49 ensure that the Markov chain has the correct posterior distribution $p(\boldsymbol{\Theta} \mid P_{1:T})$ as stationary.

### 4.1.4 Implementation Details

The algorithm is implemented in Python using `Numba` for just-in-time compilation of the particle filter steps, enabling efficient vectorized simulation of diffusion and jump components. We use $N = 1,500$ particles per likelihood evaluation and $30,000$ Metropolis–Hastings iterations, with log-normal and logit proposals tuned to target an acceptance rate around 25–30%. The priors for the parameters are weakly informative: Beta for $\phi$ (stationarity) and Gamma for the scale parameters $(\theta, \sigma, \lambda_J, \beta_J)$.

# 5 Results for the CIR-Jumps

## 5.1 Dataset

This study uses daily observations of the USD Coin (USDC) price spanning the period from January 1, 2021 to October 18, 2025. The sample comprises 1,751 observations, with the first timestamp recorded at 2021-01-01T23:59:59.999 and the last at 2025-10-18T23:59:59.999. The data consist of end-of-day closing prices expressed in U.S. dollars. Figure 1 illustrates the temporal trajectory of USDC over the analyzed period.

USDC is designed to maintain parity with the U.S. dollar; therefore, its price series fluctuates tightly around one. Nevertheless, the data reveal occasional deviations from the peg, reflecting periods of liquidity stress, redemption pressure, and market-wide dislocations. These deviations are central to our identification of the liquidity risk premium embedded in the price.

Table 1 reports descriptive statistics for the full sample. The mean and median prices are both essentially equal to one, confirming the strong anchoring of USDC to its peg. However, the minimum price of 0.9715 indicates a substantial temporary de-peg episode. The maximum price reaches 1.0105, reflecting short-lived demand-driven premia.

The distribution exhibits pronounced non-normality. The large negative skewness (-23.91) and extreme kurtosis (888.15) indicate that the series is dominated by rare but severe downside events. This heavy-tailed behavior is consistent with episodic liquidity shocks and motivates the inclusion of jump components in the state dynamics.

Log returns display near-zero mean, as expected for a pegged asset, but their standard deviation (0.000983) and extreme values (minimum -2.84%, maximum 2.10%) reveal abrupt adjustments. These stylized facts justify modeling liquidity risk with an affine jump–diffusion specification, which captures both persistent deviations and sudden dislocations.

Figure 1: UDSC Price Series

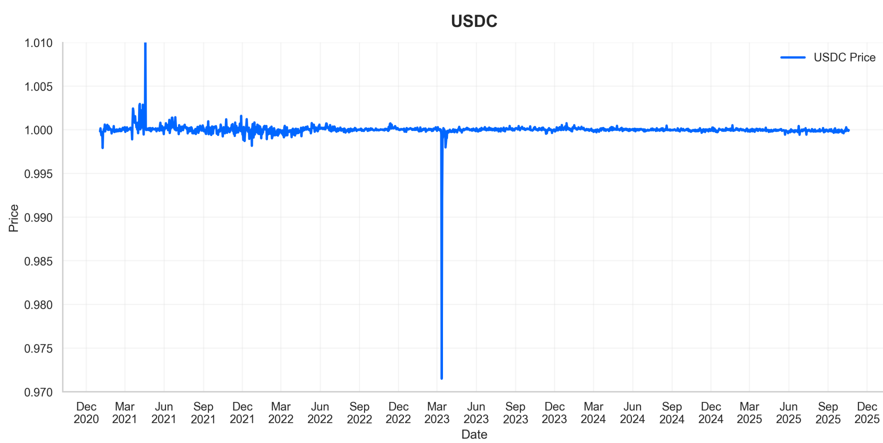


## 5.2 Particle-MCMC results

The model is estimated using 1,751 daily USDC price observations. We approximate the observed stablecoin price in empirical analysis as having a time to maturity of 5 minutes in the analytical pricing formulas, approximating the short arbitrage time between the asset and the DEFI protocols. Likelihood evaluation is performed with 1,500 particles and a fixed measurement-error standard deviation of

Table 1: Descriptive statistics of the USDC price series

| | |
|---|---|
| Obs | 1751 |
| Mean | 1.000004 |
| Median | 1.000003 |
| Standard Deviation | 0.000812 |
| Minimum | 0.971500 |
| Maximum | 1.010496 |
| Skewness | -23.905002 |
| Kurtosis | 888.145815 |
| Mean Return (log) | 0.000000 |
| Standard Deviation (log ret.) | 0.000983 |
| Minimum Return (log) | -0.028392 |
| Maximum Return (log) | 0.020952 |

$\sigma_{\text{obs}} = 0.00495$. The Particle Marginal Metropolis–Hastings algorithm is run for 30,000 iterations. Parameters are initialized at

$$(\phi, \theta, \sigma, \lambda_J, \beta_J) = \left(e^{-0.2\Delta t},\ 0.6,\ 0.4,\ 0.5,\ 3.0\right),$$

where $\phi = e^{-\kappa\Delta t}$. The first 30% of the chain is discarded as burn-in. Posterior means are computed from the remaining draws, and trace plots are inspected to assess convergence. The full chain, log-likelihood values, and acceptance rate are stored for posterior analysis.

The posterior estimates of parameters resulting of the particle-MCMC method are reported in Table 2, and provide several insights into the nature of liquidity risk embedded in USDC prices and its broader financial implications.

Table 2: Posterior Summary of Model Parameters

| Parameter | Mean | Median | Std. Dev. | 2.5% | 97.5% |
|---|---|---|---|---|---|
| $\kappa$ | 0.3267 | 0.3247 | 0.1102 | 0.1192 | 0.5273 |
| $\theta$ | 0.0055 | 0.0038 | 0.0049 | 0.0007 | 0.0195 |
| $\sigma$ | 0.1647 | 0.1123 | 0.1574 | 0.0222 | 0.6104 |
| $\lambda_J$ | 0.0018 | 0.0013 | 0.0013 | 0.0005 | 0.0053 |
| $\beta_J$ | 0.1730 | 0.1427 | 0.1467 | 0.0166 | 0.5345 |

First, the mean-reversion parameter $\kappa$ is economically significant and precisely estimated away from zero. This finding implies that liquidity shocks are transitory and decay over time, consistent with the institutional structure of fully collateralized stablecoins. In equilibrium, arbitrage mechanisms, redemption facilities, and secondary-market liquidity provision appear to restore parity following deviations. The estimated persistence, however, is non-negligible, indicating that liquidity stress episodes may last long enough to generate economically meaningful pricing effects in derivative markets.

Second, the long-run mean $\theta$ is small but positive. This suggests that, on average, a mild liquidity premium is embedded in USDC prices. From a financial perspective, this can be interpreted as compensation for redemption frictions, settlement delays, and counterparty risk associated with the issuer. Although the magnitude is modest, its presence indicates that even tightly pegged stablecoins may carry a systematic liquidity component that can be priced in no-arbitrage frameworks.

Third, the diffusion volatility parameter $\sigma$ exhibits substantial posterior dispersion. This reflects the dual nature of the stablecoin market: extended periods of near-perfect stability coexist with occasional episodes of heightened uncertainty. Continuous fluctuations capture day-to-day trading imbalances and microstructure effects, while their relatively low average magnitude confirms the strong anchoring of the peg under normal conditions.

The jump intensity $\lambda_J$ is small, indicating that extreme liquidity events are rare. Nevertheless, the estimated jump-size parameter $\beta_J$ implies economically meaningful jump magnitudes when such events occur. This combination—low frequency but sizable impact—is consistent with historical de-peg episodes triggered by exogenous shocks (e.g., banking sector stress or collateral concerns). From a risk-management perspective, this highlights the importance of modeling tail risk explicitly rather than relying solely on Gaussian diffusion dynamics.

Taken together, the estimates suggest that liquidity risk in USDC is characterized by mean-reverting behavior with infrequent but material dislocations. For derivative pricing, this implies that option premia and futures spreads will be primarily driven by expectations of temporary liquidity stress and, in particular, by the market's assessment of jump risk. For regulators and market participants, the results underscore that while stablecoins are generally stable in normal conditions, their risk profile is inherently asymmetric and dominated by rare downside events.

More broadly, the presence of a statistically significant jump component supports the view that stablecoins embed liquidity risk premia that are state-dependent and nonlinear. This reinforces the relevance of affine jump–diffusion models in capturing the financial economics of digital dollar substitutes.

Figure 2 displays the observed USDC price series together with the posterior mean of the latent liquidity risk component estimated via the Particle Markov Chain Monte Carlo procedure. The filtered latent factor closely tracks deviations of USDC from its one-dollar parity, particularly during episodes characterized by sharp price dislocations and temporary de-pegging events.

Figure 2: UDSC Price Series and Filtered Liquidity Risk - CIR-Jumps model

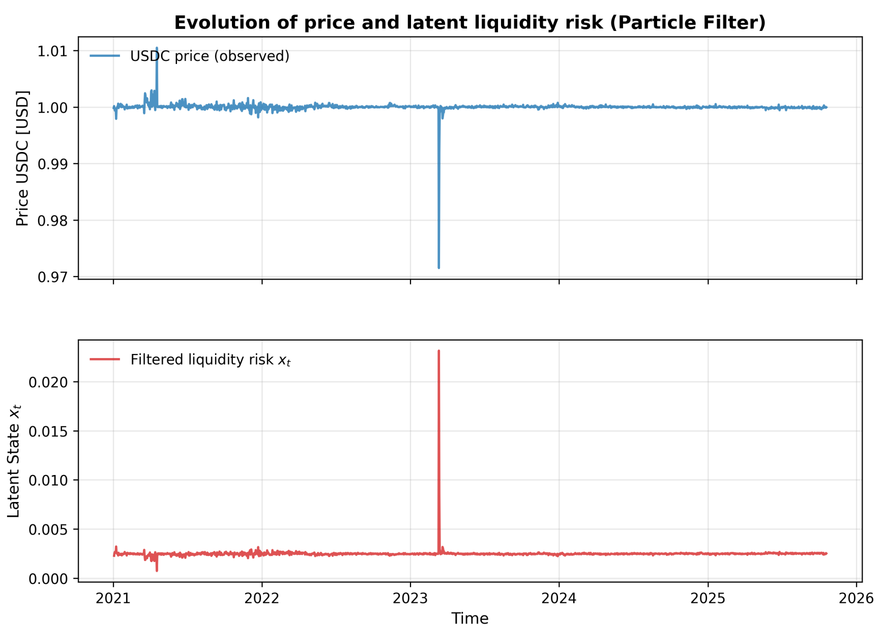


Two features emerge clearly from the figure. First, the latent factor reacts smoothly during tranquil periods, reflecting the mean-reverting structure implied by the CIR dynamics. Second, during stress episodes, the estimated process exhibits abrupt movements consistent with the presence of jump risk.

This behavior provides empirical support for the CIR jump-diffusion specification: purely diffusive dynamics would fail to capture the sudden and asymmetric deviations observed in the data, while a jump component allows the model to accommodate liquidity shocks and short-lived breakdowns in parity.

The close alignment between observed prices and the filtered latent factor also validates the use of a nonlinear state-space representation estimated via Particle MCMC. Because the model combines square-root diffusion dynamics with discontinuous jump components, the filtering problem is intrinsically nonlinear and non-Gaussian. The particle-based approach allows the latent path to adjust dynamically to extreme observations without imposing linearization or Gaussian approximations, thereby preserving the structural interpretation of liquidity risk.

Overall, the results illustrate that deviations from parity can be interpreted as manifestations of an underlying stochastic liquidity factor with both mean-reverting and jump components. The joint estimation of structural parameters and the latent path through Particle MCMC provides a coherent probabilistic framework capable of capturing both gradual liquidity adjustments and sudden de-pegging episodes.

### 5.3 Derivative Pricing Results

As an application of derivative pricing results, the Figure 3 presents the posterior predictive distribution of the 30-day USDC futures price, obtained by evaluating the semi-analytical pricing formula of the CIR jump-diffusion model across posterior draws generated by the PMMH algorithm (after burn-in). For each sampled parameter vector $(\kappa, \theta, \sigma, \lambda_J, \beta_J)$, the futures price $F(t, T)$ is computed conditional on the filtered state variable and a fixed risk-free rate. The resulting histogram therefore reflects parameter uncertainty propagated into derivative pricing, rather than mere diffusion-driven price variability.

The posterior mean provides a model-consistent point estimate of the fair futures value, while the 95% credible interval quantifies the economically plausible range of prices implied by uncertainty in mean reversion, long-run equilibrium level, volatility, and jump intensity. The dispersion and potential asymmetry of the distribution capture the contribution of liquidity shocks and jump risk to forward valuation, offering a structurally grounded measure of uncertainty surrounding short-term stablecoin futures pricing.

As a another numerical illustration, we compute the price of a 30-day European call option written on the stablecoin, using the posterior mean parameter estimates of the CIR jump-diffusion model. The parameter values are $\kappa = 0.3267$, $\theta = 0.0055$, $\sigma = 0.1647$, $\lambda_J = 0.0018$, $\beta_J = 0.1730$. The risk-free rate is set to $r = 2\%$ (annualized), the settlement lag is one day, and the time to maturity is $\Delta = 30/365$. The current latent liquidity factor is fixed at $x_t = 0.10$.

Option prices are computed using two independent numerical procedures based on the Carr–Madan Fourier inversion framework. First, we apply the Fast Fourier Transform (FFT) method to recover a vector of call prices across strikes. For a strike close to parity ($K \approx 1.0$), the FFT-based method yields $C^{\mathrm{FFT}}(K = 1.0) = 0.0006247$. Second, we compute the price for $K = 1.0$ using direct numerical integration of the damped characteristic function with adaptive refinement. The direct integration method produces $C^{\mathrm{Direct}}(K = 1.0) = 0.0006442$. The difference between the two pricing approaches is below $2 \times 10^{-5}$, confirming the numerical stability and internal consistency of the Fourier-based implementation.

Economically, the option price is small in magnitude, reflecting (i) the short maturity, (ii) the strong mean-reverting dynamics of the liquidity factor, and (iii) the relatively low posterior jump intensity. Nevertheless, the strictly positive premium captures residual uncertainty regarding short-term liquidity shocks and potential deviations from parity. This example illustrates how the CIR jump-diffusion structure translates estimated liquidity risk parameters into economically interpretable derivative prices.

Figure 3: Posterior distribution of the 30-day USDC futures price under the CIR jump-diffusion model.

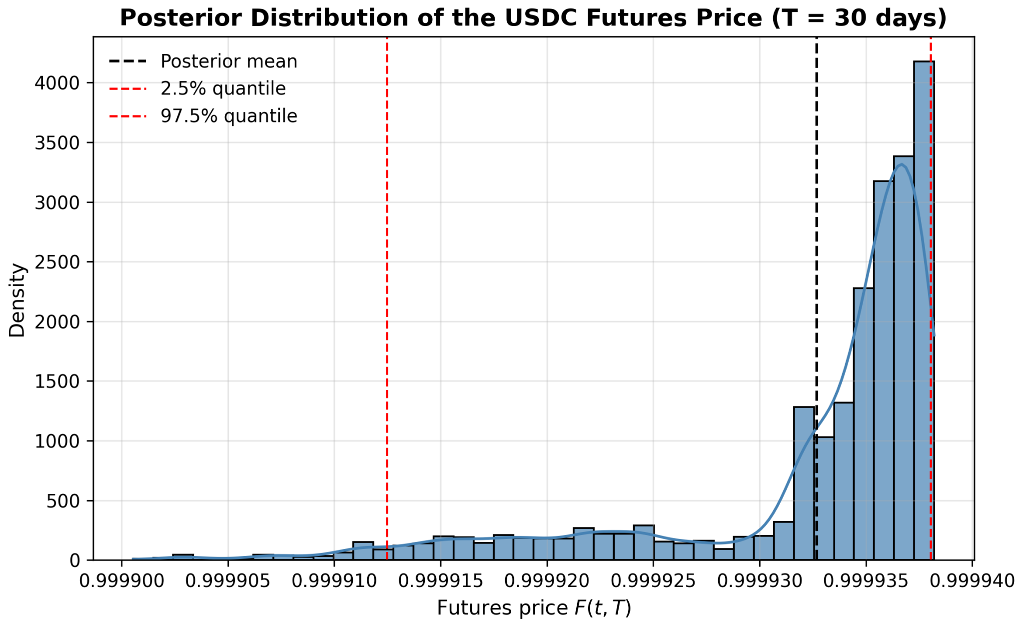


Note: The dashed black line denotes the posterior mean, while the red dashed lines indicate the 2.5% and 97.5% quantiles, forming the 95% credible interval.

## 5.4 Simulation-Based Risk Management: VaR and Expected Shortfall

Beyond derivative pricing, the estimated CIR jump-diffusion model can be used to simulate forward stablecoin price dynamics and quantify short-horizon risk measures. Using the posterior mean parameter vector obtained from the PMMH estimation procedure, we generate $N_{\text{sim}} = 10{,}000$ Monte Carlo paths over a one-day horizon ($T = 1$ day). Each simulation evolves the latent liquidity factor according to the estimated square-root jump-diffusion dynamics, and observed prices are constructed through the measurement equation including observational noise.

For each simulated path, we record the terminal price $S_T$ and compute the associated log-return

$$r_T = \log\left(\frac{S_T}{S_0}\right),$$

where $S_0$ denotes the initial simulated price. Losses are defined as

$$L_T = -r_T,$$

so that positive values correspond to adverse price movements (downward deviations from parity). On the basis of the empirical loss distribution, we introduce two widely used quantitative risk measures: Value at Risk (VaR) and Expected Shortfall (ES). In what follows, we provide formal definitions of these two measures.

**Value at Risk (VaR).** At confidence level $\alpha$, the Value at Risk is defined as the $\alpha$-quantile of the loss distribution,

$$\mathrm{VaR}_\alpha = \inf \{\ell \in \mathbb{R} : \mathbb{P}(L_T \leq \ell) \geq \alpha\}.$$

This represents the maximum expected loss over the given horizon that will not be exceeded with probability $\alpha$.

**Expected Shortfall (ES).** The Expected Shortfall (also known as Conditional VaR) is defined as

$$\mathrm{ES}_\alpha = \mathbb{E}[L_T \mid L_T > \mathrm{VaR}_\alpha],$$

and measures the average loss in the tail beyond the VaR threshold. Unlike VaR, ES is a coherent risk measure and captures tail severity.

Figure 4 presents the simulated distribution of one-day USDC prices. The distribution is tightly centered around parity, reflecting strong mean reversion in the liquidity factor. However, the left tail exhibits asymmetry induced by the jump component, generating rare but economically meaningful de-pegging scenarios. These tail realizations drive the magnitude of the VaR and Expected Shortfall measures.

Figure 4: VaR and ES - Simulated Prices

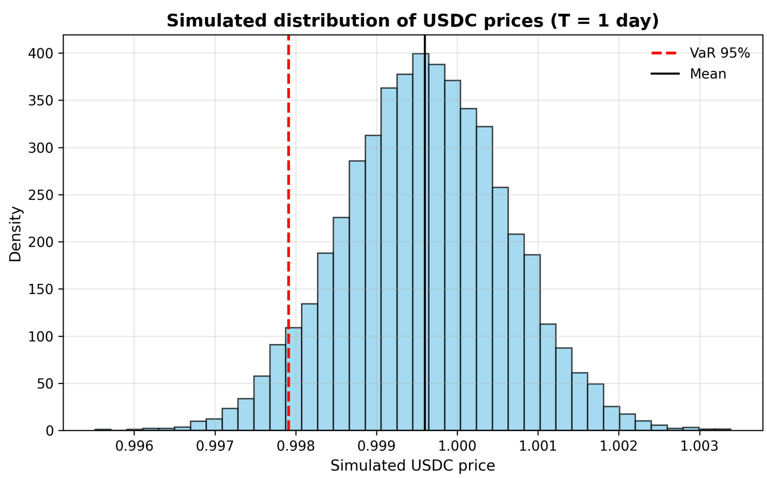


From a risk management perspective, this simulation framework provides a structurally grounded assessment of short-term stablecoin risk. Rather than relying on historical empirical quantiles alone, the Monte Carlo distribution incorporates Mean-reverting liquidity dynamics (CIR component), Discrete liquidity shocks (jump intensity $\lambda_J$ and jump size $\beta_J$) and Parameter uncertainty embedded in the estimated structural model.

Numerically, the simulation exercise yields a one-day Value at Risk at the 95% confidence level of $\mathrm{VaR}_{0.95} = 0.000273$, and a corresponding Expected Shortfall of $\mathrm{ES}_{0.95} = 0.000314$. These magnitudes indicate that, under the estimated structural dynamics, a one-day loss exceeding approximately 2.73 basis points occurs with probability 5%, while the average loss conditional on breaching this threshold is about 3.14 basis points. Although these values are small in absolute terms—consistent with the strong mean-reverting mechanism anchoring the stablecoin around parity—the gap between VaR and ES highlights the presence of non-negligible tail asymmetry. This tail thickness is primarily driven by the jump component of the liquidity factor, which generates rare but economically meaningful downward deviations from parity.

Hence, even in an environment characterized by low continuous volatility, the model-implied risk measures capture the impact of liquidity shocks on short-horizon downside risk. This reinforces the importance of incorporating jump risk into stablecoin risk management frameworks, as diffusive models alone would underestimate the severity of extreme de-pegging scenarios.

# 6 Diffusion Model with Vasicek Dynamics

While the baseline specification employs a square-root (CIR-type) process to model liquidity risk, an alternative and economically meaningful formulation consists of a Vasicek jump-diffusion process for the latent liquidity factor. In this setting, the liquidity component evolves according to an Ornstein–Uhlenbeck dynamic augmented with a jump term, allowing the state variable to take both positive and negative values.

The primary motivation for adopting a Vasicek specification lies in the economic interpretation of liquidity risk in stablecoin markets. Deviations from parity may occur in both directions: downward deviations reflect liquidity shortages or de-pegging pressure, whereas upward deviations (prices above one dollar) arise from excess demand, collateral frictions, or temporary exchange segmentation. A square-root diffusion imposes a non-negativity constraint that may be unnecessarily restrictive when the latent factor is interpreted as a net liquidity spread or imbalance measure. In contrast, the Vasicek process naturally accommodates symmetric fluctuations around a long-run equilibrium level.

Allowing negative values for the liquidity factor is particularly relevant when the model is specified in additive form within the discount rate or pricing kernel. A negative realization of the liquidity component can be interpreted as an environment of excess liquidity or positive demand pressure, compressing spreads and temporarily pushing the stablecoin price above parity. This interpretation is economically consistent with observed market episodes in which stablecoins trade at a premium due to short-term funding demand or arbitrage frictions.

The inclusion of jumps further enhances the realism of the specification. Stablecoin markets are susceptible to abrupt liquidity shocks triggered by regulatory announcements, exchange disruptions, collateral concerns, or large redemption flows. A pure Gaussian Ornstein–Uhlenbeck process would fail to reproduce the observed heavy tails and sudden price dislocations. The jump component allows the model to capture discrete liquidity events while maintaining tractable affine structure for pricing and estimation.

From a technical standpoint, the Vasicek jump-diffusion remains an affine process, preserving analytical tractability for bond pricing, characteristic function derivation, and Fourier-based option valuation. Moreover, in a state-space estimation framework, the Gaussian diffusion component simplifies filtering dynamics relative to the square-root case, potentially improving numerical stability in particle-based inference.

In summary, the Vasicek jump-diffusion specification offers three key advantages: (i) economic flexibility through the allowance of negative liquidity states, (ii) empirical realism via jump-driven tail behavior, and (iii) analytical tractability within the affine modeling framework. This makes it a compelling alternative when the liquidity factor is interpreted as a net imbalance measure rather than a strictly positive intensity or spread. In this setup we model the liquidity risk premium $x_t$ under the risk-neutral measure $\mathbb{Q}$ as a Vasicek (Ornstein–Uhlenbeck) diffusion augmented with compound Poisson jumps:

$$dx_t = \kappa(\theta - x_t)\,dt + \sigma\,dW_t^{\mathbb{Q}} + dJ_t, \tag{50}$$

where $W^{\mathbb{Q}}$ is a standard Brownian motion and

$$J_t = \sum_{i=1}^{N_t} Y_i, \qquad N_t \sim \text{Poisson}(\lambda_J t), \qquad Y_i \sim \text{Exp}(\beta_J).$$

All parameters satisfy $\kappa, \theta, \sigma, \lambda_J, \beta_J > 0$. Unlike the CIR specification, the Vasicek process allows $x_t \in \mathbb{R}$, thus accommodating negative liquidity premia. The arbitrage-free price of a zero-coupon claim paying one unit of USD at maturity $T$ is

$$P(t,T) = \mathbb{E}_t^{\mathbb{Q}}\left[\exp\left(-\int_t^T (r + x_s)\, ds\right)\right], \tag{51}$$

where $r$ is the constant risk-free rate and $\tau = T - t$. By the affine structure of the model, the bond price admits the representation

$$P(t,T) = \exp\big[-r\tau + A(\tau) - B(\tau)x_t\big], \tag{52}$$

where $A(\tau)$ and $B(\tau)$ solve a system of Riccati ODEs. Substituting the affine representation into the pricing PDE yields

$$\frac{dB}{d\tau}(\tau) = 1 - \kappa B(\tau), \qquad B(0) = 0, \tag{53}$$

$$\frac{dA}{d\tau}(\tau) = -\kappa\theta B(\tau) + \frac{1}{2}\sigma^2 B(\tau)^2 + \lambda_J\left(\frac{\beta_J}{\beta_J + B(\tau)} - 1\right), \qquad A(0) = 0. \tag{54}$$

The quadratic term in $B(\tau)$ arises from the diffusion component, while the last term reflects the Laplace transform of the exponential jump size. The linear ODE admits the closed-form solution

$$B(\tau) = \frac{1 - e^{-\kappa\tau}}{\kappa}. \tag{55}$$

Integrating the Riccati equation yields

$$A(\tau) = A_{\text{diff}}(\tau) + A_{\text{jump}}(\tau), \tag{56}$$

where the diffusion component is

$$A_{\text{diff}}(\tau) = \left(\theta - \frac{\sigma^2}{2\kappa^2}\right)\left(\tau - \frac{1 - e^{-\kappa\tau}}{\kappa}\right) - \frac{\sigma^2}{4\kappa}\left(\frac{1 - e^{-\kappa\tau}}{\kappa}\right)^2, \tag{57}$$

and the jump contribution is

$$A_{\text{jump}}(\tau) = \lambda_J \int_0^\tau \left(\frac{\beta_J}{\beta_J + B(s)} - 1\right) ds. \tag{58}$$

Since $B(s)$ is explicit, the integral may be computed analytically or via numerical quadrature. Therefore, the price is

$$P(t,T) = \exp\Big[-r\tau + A(\tau) - B(\tau)x_t\Big]. \tag{59}$$

Given an observed price $P^{\text{obs}}(t,T)$, the instantaneous liquidity premium can be inverted as

$$x_t = \frac{A(\tau) - r\tau - \ln P^{\text{obs}}(t,T)}{B(\tau)}. \tag{60}$$

## 6.1 Discrete-Time Transition

The exact discretization of the OU component over interval $\Delta t$ is

$$x_{t+\Delta t} = x_t e^{-\kappa \Delta t} + \theta(1 - e^{-\kappa \Delta t}) + \sigma \sqrt{\frac{1 - e^{-2\kappa \Delta t}}{2\kappa}}\, \varepsilon_t + J_{t+\Delta t}, \tag{61}$$

with $\varepsilon_t \sim \mathcal{N}(0,1)$ and $J_{t+\Delta t} = \sum_{i=1}^{N_{\Delta t}} Y_i$.

## 6.2 Futures Pricing

We present the Futures price using the Vasicek-Jumps model. The option price can be obtained using analogous methods. Let $F(t,T)$ denote the futures price for settlement at time $T$:

$$F(t,T) = \mathbb{E}_t^{\mathbb{Q}}\big[P(T, T+\tau_s)\big]. \tag{62}$$

Using the affine representation,

$$F(t,T) = \exp(-r\tau_s + A(\tau_s))\ \mathbb{E}_t^{\mathbb{Q}}\left[e^{-B(\tau_s)x_T}\right]. \tag{63}$$

## 6.3 Conditional Laplace Transform

Define

$$M(u, \Delta) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-u x_{t+\Delta}} \mid x_t\right].$$

The affine property implies

$$M(u,\Delta) = \exp\big(A(\Delta; u) - B(\Delta; u) x_t\big),$$

where

$$\frac{dB}{d\tau} = -\kappa B(\tau), \qquad B(0) = u, \tag{64}$$

$$\frac{dA}{d\tau} = -\kappa\theta B(\tau) + \frac{1}{2}\sigma^2 B(\tau)^2 + \lambda_J \left(\frac{\beta_J}{\beta_J + B(\tau)} - 1\right), \qquad A(0) = 0. \tag{65}$$

The solution for $B(\tau; u)$ is

$$B(\tau; u) = u e^{-\kappa\tau}.$$

The futures price becomes

$$F(t,T) = \exp\Big(-r\tau_s + A(\tau_s) + A(\Delta; B(\tau_s)) - B(\Delta; B(\tau_s))\, x_t\Big). \tag{66}$$

Table 3: Posterior summary – Vasicek-Jumps model

| Parameter | Mean | Median | Std. Dev. | 2.5% | 97.5% |
|---|---|---|---|---|---|
| $\kappa$ | 0.8062 | 0.7879 | 0.1939 | 0.4796 | 1.2812 |
| $\theta$ | 0.0243 | 0.0184 | 0.0158 | 0.0032 | 0.0608 |
| $\sigma$ | 0.0252 | 0.0202 | 0.0193 | 0.0026 | 0.0695 |
| $\lambda_J$ | 0.0076 | 0.0034 | 0.0094 | 0.0007 | 0.0350 |
| $\beta_J$ | 1.0151 | 0.5391 | 1.2256 | 0.0508 | 4.9620 |

## 6.4 Results for the Vasicek-Jumps

The posterior estimates of the Vasicek–jumps model for the same dataset, reported in Table 3 indicate a strongly mean-reverting liquidity risk premium, with $\kappa$ implying rapid shock dissipation and a small positive long-run level $\theta$ close to parity conditions. The diffusion volatility $\sigma$ is moderate, suggesting that continuous fluctuations alone do not generate large de-pegging episodes. Instead, tail risk is primarily driven by the jump component: while the estimated jump intensity $\lambda_J$ is low, the uncertainty surrounding the jump-size parameter $\beta_J$ allows for potentially sizable liquidity shocks.

Quantitatively, the Vasicek specification differs from the CIR–jumps model in three main respects. First, it permits negative liquidity premia, producing symmetric dynamics around the long-run mean, whereas the CIR model enforces non-negativity. Second, the Vasicek diffusion variance is constant, while the CIR variance is state-dependent and increases during stress episodes. Third, in the Vasicek–jumps model tail risk is mainly attributable to discrete jumps, whereas in the CIR–jumps model extreme outcomes arise from both jumps and endogenous volatility amplification.

Overall, the evidence suggests that liquidity risk in stablecoin markets is predominantly mean-reverting and low under normal conditions, with extreme events arising from infrequent but impactful liquidity disruptions. The Vasicek–jumps framework isolates this tail behavior within the jump component while maintaining analytical tractability.

## 6.5 Filtered Liquidity Risk under the Vasicek–Jumps Specification

The filtered path of the latent liquidity factor under the Vasicek–jumps model , showed in Figure 5 captures both positive and negative liquidity premia over time. This flexibility is particularly important in stablecoin markets, where deviations from parity occur in both directions. Negative realizations of the liquidity component may be interpreted as periods of excess demand or favorable funding conditions, while positive values reflect liquidity shortages and de-pegging pressure. By allowing the state variable to evolve on the entire real line, the Vasicek specification accommodates symmetric deviations around the long-run equilibrium without imposing artificial non-negativity constraints.

Empirically, the filtered series tracks well the observed fluctuations in stablecoin prices, identifying episodes of liquidity stress as well as periods of normalization. The jump component isolates abrupt dislocations, whereas the mean-reverting diffusion captures gradual adjustments back to equilibrium. This decomposition provides an economically interpretable representation of liquidity dynamics.

From a modeling perspective, the Vasicek–jumps framework offers a parsimonious yet flexible structure. With a limited number of parameters, it captures mean reversion, continuous volatility, and discrete liquidity shocks within an affine setup that preserves analytical tractability. Moreover, embedding the liquidity factor in a no-arbitrage pricing framework ensures internal consistency between the stochastic dynamics of the premium and the observed term structure of stablecoin prices.

As such, the Vasicek–jumps model provides an alternative no-arbitrage approach to modeling stablecoin prices, balancing empirical realism with structural discipline. It captures both symmetric liquidity conditions and tail risk events while remaining sufficiently tractable for estimation, simulation,

Figure 5: UDSC Price Series and Filtered Liquidity Risk - Vasicek-Jumps model

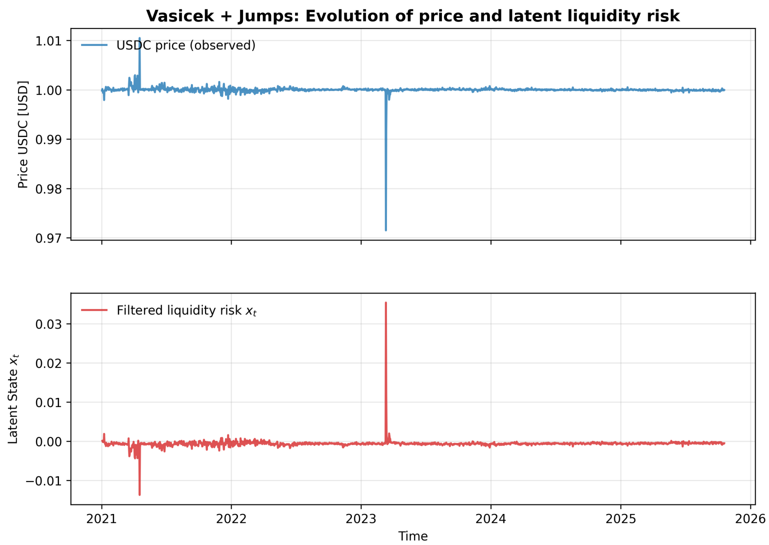


and risk management applications.

# 7 Conclusion

This paper develops a no-arbitrage framework for modeling stablecoin prices by explicitly introducing a latent liquidity risk premium evolving under affine jump–diffusion dynamics. We consider two alternative specifications for the liquidity component, a CIR–jumps process and a Vasicek–jumps process, each offering distinct economic and statistical advantages. The CIR specification ensures non-negativity and embeds state-dependent volatility, capturing stress amplification during severe liquidity episodes. The Vasicek specification, in contrast, permits negative liquidity premia, allowing symmetric deviations around parity and providing a parsimonious representation of net liquidity imbalances.

A central contribution of the paper is the joint estimation of the latent liquidity factor and structural parameters using a particle Markov chain Monte Carlo (PMCMC) algorithm. The state-space representation accommodates nonlinearity and jump components, enabling accurate filtering of the unobserved liquidity process while preserving the no-arbitrage structure of the model. The particle filtering step recovers the time-varying liquidity premium, whereas the MCMC layer provides coherent posterior inference for structural parameters. This approach allows the model to capture both gradual liquidity adjustments and abrupt de-pegging events observed in stablecoin markets.

The affine structure of both specifications ensures analytical tractability for pricing applications. Zero-coupon claims, futures, and other derivatives can be valued in closed or semi-closed form through Riccati equations and conditional Laplace transforms. This tractability makes the framework suitable not only for empirical analysis but also for practical risk management. In particular, the model can be used to simulate forward price distributions and compute risk measures such as VaR and ES, incorporating both diffusive fluctuations and jump-induced tail risk.

Overall, the results indicate that stablecoin prices are largely driven by mean-reverting liquidity conditions punctuated by infrequent but economically significant shocks. By integrating affine term structure methods, jump dynamics, and particle-based Bayesian estimation, this paper provides a coherent structural approach to modeling stablecoin pricing dynamics. The proposed framework bridges the gap between reduced-form empirical models and arbitrage-free asset pricing theory, offering a flexible and internally consistent tool for valuation, inference, and risk assessment in digital asset markets.

## References


Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(3):269–342.

Björk, T. (2019). *Arbitrage Theory in Continuous Time*. Oxford University Press, Oxford, 3 edition.

Bluhm, M., Vasiljević, A. C., Derivaux, S., and Jessen, S. T. H. (2024). Real-time risk metrics for programmatic stablecoin crypto asset-liability management (calm). *arXiv preprint arXiv:2401.13399*.

Campbell R. Harvey, C. R., Ramachandran, A., and Santoro, editors (2021). *DeFi and the Future of Finance*. Wiley, New York.

Carr, P. and Madan, D. B. (1999). Option valuation using the fast fourier transform. *Journal of Computational Finance*, 2(4):61–73.

Cont, R. and Tankov, P. (2003). *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, Boca Raton.

Cox, J. C., Ingersoll, J. E., and Ross, S. A. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.

d'Avernas, A., Maurin, V., and Vandeweyer, Q. (2022). Can stablecoins be stable? *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2022-131).

Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential monte carlo sampling methods for bayesian filtering. *Statistics and Computing*, 10(3):197–208.

Duffie, D., Filipović, D., and Schachermayer, W. (2003). Affine processes and applications in finance. *The Annals of Applied Probability*, 13(3):984–1053.

Duffie, D., Pan, J., and Singleton, K. J. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.

Eichengreen, B., Nguyen, M. T., and Viswanath-Natraj, G. (2025). Stablecoin devaluation risk. *The European Journal of Finance*, 31(11):1469–1496.

Filipović, D. (2009). *Term-Structure Models: A Graduate Course*. Springer Finance. Springer, Berlin, Heidelberg.

Gorton, G. B. and Zhang, J. Y. (2023). Taming wildcat stablecoins. *U. Chi. L. Rev.*, 90:909.

Karatzas, I. and Shreve, S. E. (1991). *Brownian Motion and Stochastic Calculus*. Springer, New York, 2nd edition.

Musiela, M. and Rutkowski, M. (2005). *Martingale Methods in Financial Modelling*. Springer, Berlin, 2 edition.

Pitt, M. K. and Shephard, N. (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599.

Schär, F. (2021). Decentralized finance: On blockchain-and smart contract-based financial markets. *FRB of St. Louis Review*.

Wen, H. and Lau, R. (2025). A risk mitigation model of monetary ecosystem with stablecoins. *arXiv preprint arXiv:2510.10469*.