Day 1: Pricing Digital Assets

Stochastic Models for Crypto Derivatives

Prof. Jörg Osterrieder

PhD Seminar – Digital Finance Research

2026

Seminar Overview: Five Days of Digital Finance

- **D1 Crypto Derivatives** – jump-diffusions, stochastic vol, equilibrium pricing
- **D2 DeFi Mathematics** – CFMMs, impermanent loss, lending protocols
- **D3 Microstructure & MEV** – order flow, sandwich attacks, fee mechanisms
- **D4 Tokenomics & Regulation** – staking, CBDC, RWA, stablecoins
- **D5 ML for Digital Finance** – deep hedging, LOB forecasting, RL

Today's Roadmap

- Why Black-Scholes fails for BTC
- Merton jump-diffusion
- Kou double-exponential model
- Heston stochastic volatility
- Bates = Heston + jumps
- Equilibrium pricing
- Calibration & applications

Why Black-Scholes Fails for Bitcoin

Black-Scholes assumes $\ln S_T \sim \mathcal{N}\big((\mu - \frac{\sigma^2}{2})T,\ \sigma^2 T\big)$.

Empirical BTC log-returns (daily):

- Excess kurtosis $\approx$ 8–15 (vs. 0 for Gaussian)
- Negative skewness during crashes
- Volatility clustering: calm $\to$ storm $\to$ calm
- Jumps: $> 10\%$ daily moves occur $\sim$monthly

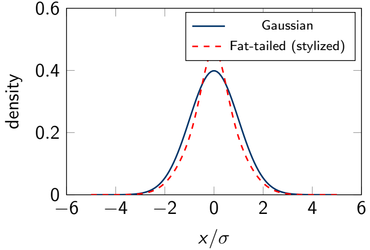

Consequence

BS misprices OTM puts by 50–200%; the implied-vol smile is steep and persistent [7].

BTC Return Distribution: Heavy Tails and Skewness

Stylized histogram of BTC daily returns:

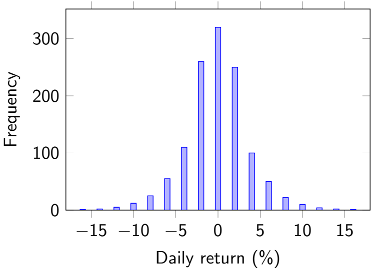

Key statistics (2017–2025):

Mean (ann.)	$\approx 60\%$
Std (ann.)	$\approx 70\%$
Skewness	-0.3 to -0.8
Kurtosis	8–15
Max daily loss	-39% (Mar 2020)
Max daily gain	$+23\%$

Takeaway: tails are *much* heavier than Gaussian. Jump models are not optional–they are necessary.

Deribit Implied Volatility Surface

Conceptual BTC IV surface:

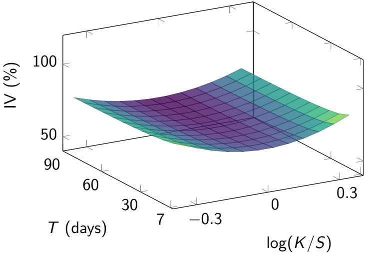

Features of the crypto smile:

- Pronounced U-shape at short tenors
- Negative skew for puts (crash risk premia)
- Positive skew for calls at long tenors ("FOMO premium")
- ATM IV $\approx$ 50–80% (vs. $\approx$ 15–20% for SPX)
- Term structure: inverted during crises

Source: Deribit exchange data; BTC options comprise $\sim$90% of crypto options volume.

Stylized Facts of Crypto Returns

Market structure:

- 24/7/365 trading – no overnight gap, no weekends
- No circuit breakers on most venues
- Fragmented across CEX/DEX (arbitrage: [6])
- High retail participation [3]

Return properties:

- Heavy tails (power-law index $\alpha \approx 3$)
- Volatility clustering (GARCH effects)
- Leverage effect (weaker than equities)
- Long memory in absolute returns

Extreme events:

- March 2020: -39% single day
- May 2021: -30% in 48h (China ban)
- Nov 2022: FTX collapse cascade
- Flash crashes on low-liquidity venues

Implication for Modeling

Need: (i) jumps for tail events, (ii) stochastic volatility for clustering, (iii) possibly correlated jumps in price and variance.

What We Need: Beyond Black-Scholes

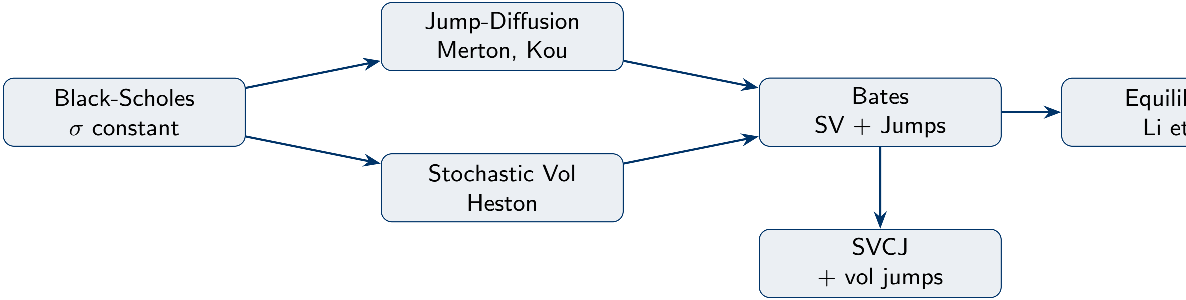

- Each arrow adds a **degree of freedom** to match an empirical feature
- **Jumps** $\rightarrow$ heavy tails, sudden moves
- **Stochastic vol** $\rightarrow$ volatility clustering, smile dynamics
- **Correlated jumps** $\rightarrow$ crash-induced vol spikes
- **Equilibrium** $\rightarrow$ endogenous risk premia, liquidity risk

Today's Roadmap

1. Jump-Diffusion Models
2. Stochastic Volatility Models
3. Equilibrium Pricing
4. Extensions
5. Applications & Industry

Geometric Brownian Motion: The Baseline

GBM under $\mathbb{Q}$ (risk-neutral measure)

$$\frac{dS_t}{S_t} = (r-q)\,\mathrm{d}t + \sigma\,\mathrm{d}W_t^{\mathbb{Q}}, \qquad S_0 > 0$$

Solution:

$$S_T = S_0 \exp\left[\left(r - q - \tfrac{\sigma^2}{2}\right)T + \sigma W_T^{\mathbb{Q}}\right]$$

Log-return: $\ln(S_T/S_0) \sim \mathcal{N}\left(\left(r-q-\tfrac{\sigma^2}{2}\right)T,\ \sigma^2 T\right)$

Properties:

- Continuous paths – no jumps possible
- Constant volatility σ – no smile
- Gaussian returns – thin tails
- One parameter (σ) for all strikes and maturities

BS European call: $C = S_0 e^{-qT}\Phi(d_1) - Ke^{-rT}\Phi(d_2), \quad d_{1,2} = \frac{\ln(S_0/K)+(r-q\pm\sigma^2/2)T}{\sigma\sqrt{T}}$

Merton (1976): Adding Jumps to GBM

Merton Jump-Diffusion SDE

$$\frac{dS_t}{S_{t-}} = (\mu - \lambda\bar{k})\,\mathrm{d}t + \sigma\,\mathrm{d}W_t + (e^J - 1)\,\mathrm{d}N_t$$

Components:

- W_t: standard Brownian motion (diffusion component)
- N_t: Poisson process with intensity λ (jump arrivals)
- $J \sim \mathcal{N}(\mu_J, \sigma_J^2)$: log-jump size (normally distributed)
- $\bar{k} = \mathbb{E}[e^J - 1] = e^{\mu_J + \sigma_J^2/2} - 1$: compensator

Parameters: $\theta_{\text{Merton}} = (\sigma, \lambda, \mu_J, \sigma_J)$ (4 params vs. 1 for BS)

Solution:

$$S_T = S_0 \exp\left[\left(\mu - \lambda\bar{k} - \tfrac{\sigma^2}{2}\right)T + \sigma W_T + \sum_{i=1}^{N_T} J_i\right]$$

Poisson Jumps: Mechanics

Poisson process $N(t)$ with intensity λ:

- $\Pr(N(t+\Delta t) - N(t) = 1) = \lambda \Delta t + o(\Delta t)$
- $\Pr(\text{no jump in } \Delta t) = 1 - \lambda \Delta t$
- Inter-arrival times: $\tau_i \sim \mathrm{Exp}(\lambda)$
- Expected jumps per year: λ

Jump size J (Merton):

$$J \sim \mathcal{N}(\mu_J, \sigma_J^2)$$

- $\mu_J < 0$: downward jumps on average
- σ_J: jump size dispersion
- Multiplicative impact: $S \to S \cdot e^J$

Sample path with jumps:

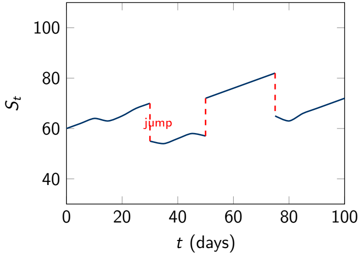

Merton Model: Characteristic Function

The **characteristic function** of $\ln S_T$ under $\mathbb{Q}$ is the key to Fourier pricing.

Merton CF

$$\phi_T^{\text{Merton}}(u) = \mathbb{E}^{\mathbb{Q}}\left[e^{iu\ln S_T}\right] = \exp\left[iu\left(\ln S_0 + (r - q - \tfrac{\sigma^2}{2} - \lambda\bar{k})T\right) - \tfrac{u^2\sigma^2 T}{2} + \lambda T\left(e^{iu\mu_J - u^2\sigma_J^2/2} - 1\right)\right]$$

Derivation sketch:

1. Conditional on $N_T = n$ jumps: $\ln S_T \sim \mathcal{N}(m_n, s_n^2)$ with

$$m_n = \ln S_0 + (r - q - \tfrac{\sigma^2}{2} - \lambda\bar{k})T + n\mu_J, \quad s_n^2 = \sigma^2 T + n\sigma_J^2$$

2. CF of a Gaussian is $\exp(ium - u^2 s^2/2)$
3. Sum over n: $\phi_T(u) = \sum_{n=0}^{\infty} \frac{(\lambda T)^n e^{-\lambda T}}{n!} \cdot e^{ium_n - u^2 s_n^2/2}$
4. Closed-form summation yields the result above

Fourier Pricing: The Carr-Madan Formula

Given the CF $\phi_T(u)$, we can price **any** European option via Fourier inversion.

Carr-Madan (1999) formula for European call

$$C(K) = \frac{e^{-\alpha \ln K}}{\pi} \int_0^{\infty} e^{-iv \ln K}\, \psi(v)\, \mathrm{d}v$$

where $\alpha > 0$ is a damping parameter and

$$\psi(v) = \frac{e^{-rT}\, \phi_T\big(v - (\alpha + 1)i\big)}{\alpha^2 + \alpha - v^2 + i(2\alpha + 1)v}$$

Implementation:

1. Choose $\alpha \in [1, 2]$ (controls integrability)
2. Discretize: $v_j = j\Delta v$, $j = 0, \dots, N - 1$
3. Apply FFT: $O(N \log N)$ for all strikes simultaneously
4. Interpolate onto observed strikes

Advantage: Works for *any* model with known CF – Merton, Kou, Heston, Bates, VG, ...

Carr-Madan: The Full Pricing Integral

Starting point: The modified call price $c(k) = e^{\alpha k} C(K)$ with $k = \ln K$:

$$c(k) = \int_{-\infty}^{\infty} e^{ivk} \, \psi(v) \, \mathrm{d}v$$

So the call price is:

$$\boxed{C(K) = \frac{e^{-\alpha \ln K}}{\pi} \int_0^{\infty} e^{-iv \ln K} \cdot \frac{e^{-rT} \, \phi_T\big(v - (\alpha+1)i\big)}{\alpha^2 + \alpha - v^2 + i(2\alpha+1)v} \, \mathrm{d}v}$$

Key observations:

- The denominator is $\alpha(\alpha+1) - (v^2 - i(2\alpha+1)v) = -\big((iv+\alpha)(iv+\alpha+1)\big)$
- $\alpha > 0$ ensures $e^{\alpha k} C(K)$ is L^1-integrable
- The integral is a single FFT – evaluating N strikes costs $O(N \log N)$
- Typical: $N = 2^{12} = 4096$ grid points suffice

For puts: use put-call parity or damp with $\alpha < 0$.

Kou (2002): Double-Exponential Jump-Diffusion

Motivation: Merton's symmetric (Gaussian) jumps miss the asymmetric nature of crypto crashes vs. rallies [4].

Kou SDE

$$\frac{dS_t}{S_{t-}} = (\mu - \lambda \bar{k})\,\mathrm{d}t + \sigma\,\mathrm{d}W_t + (e^J - 1)\,\mathrm{d}N_t$$

Same structure as Merton, but J has a **double-exponential** distribution.

Jump size density:

$$f_J(y) = p\,\eta_1\,e^{-\eta_1 y}\,\mathbf{1}_{\{y \geq 0\}} + q\,\eta_2\,e^{\eta_2 y}\,\mathbf{1}_{\{y < 0\}}, \quad p + q = 1$$

Parameters: $\theta_{\text{Kou}} = (\sigma, \lambda, p, \eta_1, \eta_2)$ (5 params)

- p: probability of an *upward* jump
- $\eta_1 > 1$: rate of exponential decay for upward jumps (ensures $\mathbb{E}[e^J] < \infty$)
- $\eta_2 > 0$: rate for downward jumps
- Asymmetry: $\eta_1 \neq \eta_2$ captures different tail behavior

Kou Jump Size: Double-Exponential Density

$$f_J(y) = p\,\eta_1\,e^{-\eta_1 y}\,\mathbf{1}_{\{y\geq 0\}} + q\,\eta_2\,e^{\eta_2 y}\,\mathbf{1}_{\{y<0\}}$$

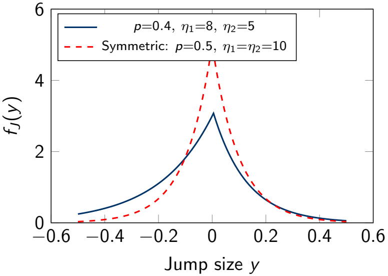

Moments:

$$\mathbb{E}[J] = \frac{p}{\eta_1} - \frac{q}{\eta_2}$$

$$\mathrm{Var}[J] = \frac{2p}{\eta_1^2} + \frac{2q}{\eta_2^2} - \big(\mathbb{E}[J]\big)^2$$

$$\bar{k} = \frac{p\eta_1}{\eta_1 - 1} + \frac{q\eta_2}{\eta_2 + 1} - 1$$

Key property: memoryless tails $\rightarrow$ leptokurtic returns with finite moments of all orders.

Kou: Characteristic Function and Pricing

Kou CF under $\mathbb{Q}$

$$\phi_T^{\text{Kou}}(u) = \exp\Big[iu\big(\ln S_0 + (r - q - \tfrac{\sigma^2}{2} - \lambda\bar{k})T\big) - \frac{u^2\sigma^2 T}{2} + \lambda T\Big(\frac{p\eta_1}{\eta_1 - iu} + \frac{q\eta_2}{\eta_2 + iu} - 1\Big)\Big]$$

Analytical tractability:

- CF has simple rational form in the jump component
- **European calls:** Carr-Madan FFT or direct Laplace inversion
- **Barrier/lookback options:** Kou derived closed-form prices using the "Laplace transform of overshoot" technique – a major advantage over Merton
- **First-passage times:** double-exponential has memoryless property $\rightarrow$ explicit formulas for first-passage distributions

Practical note: Kou requires $\eta_1 > 1$ to ensure $\mathbb{E}^{\mathbb{Q}}[e^J] < \infty$ (finite forward price).

Why Kou Beats Merton for Crypto

Empirical advantages:

1. **Asymmetric tails:** crypto crashes are sharper than rallies
   $\rightarrow \eta_2 < \eta_1$ captures heavier left tail
2. **Leptokurtic returns:** double-exponential generates higher kurtosis per unit of jump intensity
3. **Analytical barriers:** Kou gives closed-form barrier prices; Merton does not
4. **Better calibration:** fewer parameters achieve tighter smile fits

Result from [1]:

- Calibrated BS, Merton, Kou, Heston, Bates to BTC Deribit options
- **Kou best-fits BTC options** across strikes and tenors
- For ETH: Bates wins (higher vol-of-vol effect)

Key Insight

For BTC, jumps matter more than stochastic volatility. For ETH, both matter.

Empirical: Kou Best-Fits BTC Options

Calibration results from arXiv 2506.14614 [1]:

Model	**Params**	**RMSE (IV)**	**AIC**	**OOS Error**
Black-Scholes	1	12.3%	–	15.1%
Merton	4	4.8%	-210	6.2%
Kou	**5**	**2.1%**	$\mathbf{-312}$	**3.0%**
Heston	5	3.5%	-265	4.4%
Bates	8	1.9%	-298	3.3%

Key observations:

- Kou achieves near-Bates accuracy with **3 fewer parameters**
- Bates overfits at short tenors; Kou is more robust out-of-sample
- Heston alone cannot capture the steep short-tenor smile
- BS RMSE is $6\times$ larger – confirms the inadequacy of Gaussian assumptions

Note: values are illustrative of the relative ordering reported; consult the paper for exact figures.

Model Comparison: BS vs. Merton vs. Kou

Feature	Black-Scholes	Merton	Kou
Parameters	1 (σ)	4 ($\sigma, \lambda, \mu_J, \sigma_J$)	5 ($\sigma, \lambda, p, \eta_1, \eta_2$)
Jump component	None	Log-normal	Double-exponential
Asymmetric jumps	No	Partial (μ_J)	Yes ($\eta_1 \neq \eta_2$)
Heavy tails	No	Yes	Yes (heavier)
Closed-form Euro	Yes (BS)	Yes (series)	Yes (CF/FFT)
Closed-form barrier	Yes	No	Yes
Smile generated	Flat	Moderate	Pronounced
BTC fit quality	Poor	Good	**Best**

Mathematical reason for Kou's advantage:

- The memoryless property of exponential distributions enables explicit first-passage time distributions
- $\Pr(J > x \mid J > 0) = e^{-\eta_1 x}$ – Markovian structure simplifies path-dependent pricing
- Gaussian jumps lack this property $\rightarrow$ barrier pricing requires simulation

Motivation: Volatility is Not Constant

GARCH effects in crypto:

- Autocorrelation of $|r_t|$ and r_t^2 is strong (half-life $\approx$ 20–40 days)
- Volatility "regimes": BTC realized vol ranges from 30% to 150% annualized
- Implied vol surface changes shape over time – inconsistent with constant σ

Conceptual vol time series:

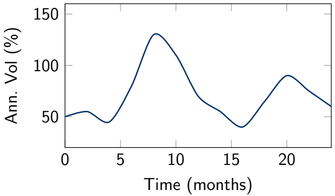

What we need:

- Variance v_t follows its own SDE
- Mean-reverting: vol returns to long-run level
- Vol-of-vol: v_t itself is random
- Correlation ρ between price and vol shocks

$\rightarrow$ **Heston (1993)** [2]

Heston (1993): Stochastic Volatility Model

Heston SDE system under $\mathbb{Q}$

$$\frac{dS_t}{S_t} = (r - q)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^S$$

$$dv_t = \kappa(\theta - v_t)\,\mathrm{d}t + \xi\sqrt{v_t}\,\mathrm{d}W_t^v$$

$$\mathrm{corr}(dW_t^S, dW_t^v) = \rho$$

Parameters: $\Theta_{\text{Heston}} = (v_0, \kappa, \theta, \xi, \rho)$ (5 params)

Param	Meaning	Typical BTC
v_0	Current variance	$0.50^2 = 0.25$
κ	Mean-reversion speed	2–5
θ	Long-run variance	$0.60^2 = 0.36$
ξ	Vol-of-vol	0.5–2.0
ρ	Price-vol correlation	-0.3 to -0.7

Feller condition: $2\kappa\theta > \xi^2$ ensures $v_t > 0$ a.s. (often violated in crypto calibrations).

Heston: Characteristic Function

Heston CF (Gatheral formulation)

$$\phi_T^{\text{Heston}}(u) = \exp\Big[iu(\ln S_0 + (r-q)T) + A(u,T) + B(u,T)\,v_0\Big]$$

where:

$$A(u,T) = \frac{\kappa\theta}{\xi^2}\left[(\kappa - \rho\xi\, iu - d)\,T - 2\ln\!\left(\frac{1-g\,e^{-dT}}{1-g}\right)\right]$$

$$B(u,T) = \frac{\kappa - \rho\xi\, iu - d}{\xi^2}\cdot\frac{1-e^{-dT}}{1-g\,e^{-dT}}$$

$$d = \sqrt{(\rho\xi\, iu - \kappa)^2 + \xi^2(iu + u^2)}$$

$$g = \frac{\kappa - \rho\xi\, iu - d}{\kappa - \rho\xi\, iu + d}$$

Pricing: plug ϕ_T^{Heston} into Carr-Madan $\rightarrow$ European prices via FFT.

Note: use the "rotation trick" ($d \rightarrow -d$ when $\Re(g) > 1$) for numerical stability.

Calibration: Fitting Models to Market Data

Objective: find Θ^* minimizing distance between model and market IVs.

Calibration as optimization

$$\Theta^* = \arg\min_{\Theta} \sum_{i=1}^{M} w_i \left(\sigma_i^{\text{model}}(\Theta) - \sigma_i^{\text{mkt}}\right)^2 + \lambda_{\text{reg}} \|\Theta - \Theta_0\|^2$$

Components:

- M observed options (strike, maturity, market IV)
- Weights w_i: by vega, volume, or 1/bid-ask spread
- Regularization $\lambda_{\text{reg}} \|\Theta - \Theta_0\|^2$: Tikhonov penalty for stability
- Inner loop: for each Θ, compute model prices via FFT, invert to model IV

Optimizers for crypto:

- Differential evolution (global, gradient-free) – robust for multimodal landscapes
- Levenberg-Marquardt (local, fast) – for refinement after global search
- Typical runtime: $\sim$2s for 100 options with FFT + DE on GPU

Bates (1996): Heston + Merton Jumps

Idea: Combine stochastic volatility (Heston) with price jumps (Merton).

Bates SDE system under $\mathbb{Q}$

$$\frac{dS_t}{S_{t^-}} = (r - q - \lambda\bar{k})\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^S + (e^J - 1)\,\mathrm{d}N_t$$

$$dv_t = \kappa(\theta - v_t)\,\mathrm{d}t + \xi\sqrt{v_t}\,\mathrm{d}W_t^v$$

$$\mathrm{corr}(dW^S, dW^v) = \rho, \qquad J \sim \mathcal{N}(\mu_J, \sigma_J^2), \qquad N_t \sim \mathrm{Poisson}(\lambda)$$

Parameters: $\Theta_{\text{Bates}} = (v_0, \kappa, \theta, \xi, \rho, \lambda, \mu_J, \sigma_J)$ (8 params)

Characteristic function:

$$\phi_T^{\text{Bates}}(u) = \phi_T^{\text{Heston}}(u) \cdot \exp\Big[\lambda T\left(e^{iu\mu_J - u^2\sigma_J^2/2} - 1\right) - iu\lambda\bar{k}\,T\Big]$$

Modular: Heston CF $\times$ Merton jump CF – elegant multiplicative structure.

Bates: Full System and Dynamics

Complete Bates dynamics (explicit):

$$S_t = S_0 \exp\left[\int_0^t \left(r - q - \lambda\bar{k} - \tfrac{v_s}{2}\right) \mathrm{d}s + \int_0^t \sqrt{v_s}\, \mathrm{d}W_s^S + \sum_{i=1}^{N_t} J_i\right]$$

$$v_t = v_0 + \kappa\theta\, t - \kappa \int_0^t v_s \, \mathrm{d}s + \xi \int_0^t \sqrt{v_s}\, \mathrm{d}W_s^v$$

Three sources of randomness:

1. W_t^S: diffusive price shocks (continuous)
2. W_t^v: variance shocks (correlated with price via ρ)
3. $N_t, \{J_i\}$: jump arrivals and sizes (independent of W^S, W^v)

Monte Carlo simulation (Euler scheme for v_t):

$$v_{t+\Delta t} = v_t + \kappa(\theta - v_t)\Delta t + \xi\sqrt{v_t^+}\,\sqrt{\Delta t}\, Z^v, \qquad v_t^+ = \max(v_t, 0)$$

Full truncation or QE scheme (Andersen 2008) for variance positivity.

Key Result: Bates Best-Fits ETH Options

From [1]:

Model	**BTC RMSE**	**ETH RMSE**
Black-Scholes	12.3%	14.1%
Merton	4.8%	5.5%
Kou	**2.1%**	3.2%
Heston	3.5%	3.0%
Bates	1.9%	**1.6%**

Why Bates wins for ETH:

- ETH has higher vol-of-vol (ξ) due to DeFi-related shocks
- Stochastic vol captures time-varying smile shape
- Jumps alone (Kou) insufficient for ETH smile dynamics

Economic interpretation:

- BTC $\approx$ "digital gold": jump-dominated
- ETH $\approx$ "DeFi platform token": regime-dependent vol from TVL, gas, DeFi events
- Different assets $\rightarrow$ different optimal models

Practical Rule

Calibrate *multiple* models. Select via AIC/BIC and out-of-sample performance.

SVCJ: Stochastic Volatility with Correlated Jumps

Extension of Bates: add **simultaneous jumps in variance** correlated with price jumps.

SVCJ SDE system

$$\frac{dS_t}{S_{t-}} = (r - q - \lambda\bar{k})\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^S + (e^{J^S} - 1)\,\mathrm{d}N_t$$

$$dv_t = \kappa(\theta - v_t)\,\mathrm{d}t + \xi\sqrt{v_t}\,\mathrm{d}W_t^v + J^v\,\mathrm{d}N_t$$

Joint jump distribution:

$$J^v \sim \mathrm{Exp}(\mu_v), \quad \mu_v > 0 \quad \text{(variance always jumps up)}$$
$$J^S \mid J^v \sim \mathcal{N}(\mu_J + \rho_J J^v,\ \sigma_J^2) \quad \text{(conditional on vol jump)}$$

Key feature: ρ_J links price and vol jumps. When $\rho_J < 0$: a large *negative* price jump coincides with a large *positive* vol jump – exactly what happens in crypto crashes.

Parameters: $\Theta_{\text{SVCJ}} = (v_0, \kappa, \theta, \xi, \rho, \lambda, \mu_J, \sigma_J, \mu_v, \rho_J)$ (10 params)

Model Hierarchy: Nesting Structure

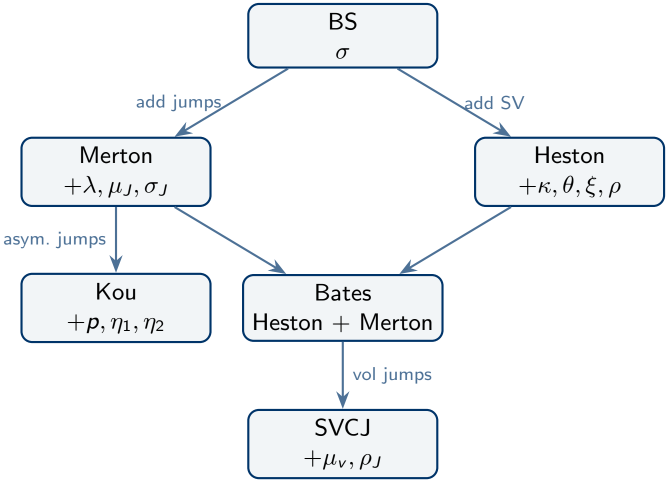

Nesting: BS $\subset$ Merton $\subset$ Bates $\subset$ SVCJ. Each level adds parameters and realism.
Trade-off: more parameters $\rightarrow$ better fit $\rightarrow$ harder calibration $\rightarrow$ potential overfitting.

Parameter Interpretation: What Each Captures

Parameter	Captures	Observable Consequence
σ (or $\sqrt{v_0}$)	ATM vol level	BS-equivalent flat vol
λ	Jump frequency	Tail probability, kurtosis
μ_J / p, η_1, η_2	Jump size dist.	Skew of smile, asymmetric tails
κ	Vol mean-reversion	Term structure slope of IV
θ	Long-run vol	Far-maturity IV level
ξ	Vol-of-vol	Curvature (convexity) of smile
ρ	Price-vol correlation	Skew of smile (leverage effect)
μ_v	Vol jump size	Short-term IV spike magnitude
ρ_J	Jump correlation	Crash-induced vol clustering

Identification strategy:

- Fix v_0 from ATM IV $\rightarrow$ reduces by 1 DOF
- κ, θ from term structure; ξ, ρ from smile shape
- λ, μ_J from short-tenor OTM puts (jump risk)
- Sequential calibration: first Heston, then add jumps

From Reduced-Form to Equilibrium

So far: reduced-form models – specify dynamics exogenously under $\mathbb{Q}$, calibrate to options.

Limitations:

- Risk premia are *assumed*, not derived
- No connection between physical and risk-neutral measures
- Why does the jump risk premium have a particular sign/magnitude?
- Liquidity risk is absent

Equilibrium approach:

- Model the *economy*: preferences, endowments, market structure
- **Derive** the pricing kernel (SDF) from first principles
- Risk premia *emerge* from the interaction of agents and market frictions
- Connects: physical dynamics $\xrightarrow{\text{SDF}}$ risk-neutral dynamics $\xrightarrow{\text{pricing}}$ option prices

Key reference: [5] – equilibrium model for BTC options with SV, correlated jumps, and liquidity risk.

Li et al. (JFM 2026): Equilibrium BTC Option Pricing

Framework [5]: representative agent, CRRA utility, BTC as risky asset with:

- Stochastic volatility (Heston-type)
- Correlated jumps in price and volatility (SVCJ-type)
- **Liquidity risk**: stochastic bid-ask spread, illiquidity factor

Key equations (physical measure $\mathbb{P}$):

$$
\begin{aligned}
\frac{dS_t}{S_{t^-}} &= \mu_S(v_t, \ell_t)\,\mathrm{d}t + \sqrt{v_t}\,\mathrm{d}W_t^S + J^S\,\mathrm{d}N_t \\
dv_t &= \kappa^{\mathbb{P}}(\theta^{\mathbb{P}} - v_t)\,\mathrm{d}t + \xi\sqrt{v_t}\,\mathrm{d}W_t^v + J^v\,\mathrm{d}N_t \\
d\ell_t &= \kappa_\ell(\bar{\ell} - \ell_t)\,\mathrm{d}t + \sigma_\ell\sqrt{\ell_t}\,\mathrm{d}W_t^\ell
\end{aligned}
$$

where ℓ_t is the **liquidity process** (e.g., bid-ask spread).

The drift $\mu_S(v_t, \ell_t)$ is **endogenous**: it includes risk premia for volatility, jump, and liquidity risks derived from the equilibrium.

Equilibrium: Deriving Risk Premia

Pricing kernel (stochastic discount factor):

$$\frac{dM_t}{M_{t-}} = -r_f\,\mathrm{d}t - \lambda_S\sqrt{v_t}\,\mathrm{d}W_t^S - \lambda_v\xi\sqrt{v_t}\,\mathrm{d}W_t^v - \lambda_\ell\sigma_\ell\sqrt{\ell_t}\,\mathrm{d}W_t^\ell + (\zeta-1)\,\mathrm{d}\tilde{N}_t$$

Market prices of risk:

- λ_S: diffusive equity risk premium per unit vol
- λ_v: variance risk premium (typically $\lambda_v > 0$ – investors pay for vol protection)
- λ_ℓ: **liquidity risk premium –** novel in crypto context
- ζ: jump risk premium ($\zeta < 1$ means negative jump risk premium)

Measure change $\mathbb{P} \to \mathbb{Q}$:

$$\kappa^{\mathbb{Q}} = \kappa^{\mathbb{P}} + \lambda_v\xi, \qquad \theta^{\mathbb{Q}} = \frac{\kappa^{\mathbb{P}}\theta^{\mathbb{P}}}{\kappa^{\mathbb{Q}}}$$

$$\lambda^{\mathbb{Q}} = \zeta \cdot \lambda^{\mathbb{P}}$$

Under $\mathbb{Q}$, option pricing proceeds via CF + FFT as before, but parameters are now **economically grounded**.

Liquidity Risk Premium in Crypto Options

Novel contribution of [5]: liquidity risk is priced in BTC options.

Mechanism:

1. BTC markets exhibit time-varying liquidity (bid-ask spreads widen during stress)
2. Illiquidity correlates with volatility spikes ("liquidity spirals")
3. Risk-averse investors demand compensation for bearing liquidity risk
4. This generates an **additional risk premium** absent in standard models

Impact on option prices:

- OTM puts are *more* expensive than predicted by SVCJ alone
- The liquidity premium accounts for ∼5–15% of the OTM put premium
- Term structure: liquidity premium decays with maturity (mean-reversion of ℓ_t)

Empirical finding

Models ignoring liquidity risk systematically **underprice** OTM BTC puts, especially during stress periods (FTX collapse, Luna crash).

Implications: Why Market-Makers Care About Equilibrium

For dealers / market-makers:

- Equilibrium models explain *why* OTM puts carry high premia
- Better delta-hedging: vol and jump risk premia affect hedge ratios
- Liquidity-adjusted Greeks: $\Delta^{\text{liq}} \neq \Delta^{\text{SVCJ}}$

For risk managers:

- Stress testing: scenario-consistent risk premia
- VaR under physical measure uses equilibrium-implied $\mu_S(v_t, \ell_t)$
- Tail risk: equilibrium predicts co-movement of vol, jumps, liquidity

For researchers:

- Connects option pricing to macro/DeFi fundamentals
- Testable predictions: how TVL shocks affect option smile
- Can incorporate blockchain-specific frictions (MEV, gas costs)

Open Question

Can we endogenize ℓ_t from DEX microstructure (Uniswap depth, gas prices)?

Extension: Regime-Switching Models (Self-Study)

Idea: crypto markets alternate between bull/bear/consolidation regimes.

Hidden Markov Model (HMM) for regime detection

Latent state $Z_t \in \{1, 2, \ldots, K\}$ with transition matrix $\boldsymbol{P}$:

$$P_{ij} = \Pr(Z_{t+1} = j \mid Z_t = i)$$

Conditional dynamics:

$$r_t \mid Z_t = k \sim \mathcal{N}(\mu_k, \sigma_k^2) \quad \text{or more generally a jump-diffusion in regime } k$$

Key features:

- Each regime has its own $(\mu_k, \sigma_k, \lambda_k)$ – separate drift, vol, jump intensity
- Regime transitions: estimated via Baum-Welch (EM) algorithm
- **Crypto application:** 2-state model captures "risk-on" ($\sigma \approx 40\%$, $\lambda \approx 1$) vs. "risk-off" ($\sigma \approx 100\%$, $\lambda \approx 10$) regimes
- Pricing: weighted average of within-regime option prices (via CF for each regime)

Self-study: Hamilton (1989) for econometrics; [7] for crypto option hedging with regimes.

Extension: Variance Gamma and Lévy Processes

Beyond compound Poisson: infinite-activity Lévy processes have infinitely many small jumps per unit time – no diffusion component needed.

Variance Gamma (VG) process

$$X_t^{\mathsf{VG}} = \theta\, G_t + \sigma\, W_{G_t}$$

where $G_t \sim \mathrm{Gamma}(t/\nu, \nu)$ is a random time change ("business time").

VG characteristic function:

$$\phi_T^{\mathsf{VG}}(u) = \left(1 - iu\theta\nu + \tfrac{1}{2}u^2\sigma^2\nu\right)^{-T/\nu}$$

Parameters: (σ, θ, ν) – volatility, drift, variance rate of the Gamma clock.

Advantages: pure-jump process with finite variation; generates heavier tails than Merton without a diffusion component. **CGMY** generalizes further.

For crypto: VG fits well at short tenors where jump clustering dominates diffusion.

Extension: Bayesian MCMC for Regime Detection

Problem: MLE for HMM-jump-diffusion is multimodal; frequentist confidence intervals unreliable.
Bayesian approach:

- Prior: $\pi(\Theta)$ on all parameters including regime transition probabilities
- Likelihood: $L(\{r_t\} \mid \Theta, \{Z_t\})$ – Gaussian mixture conditional on regimes
- Posterior: $\pi(\Theta, \{Z_t\} \mid \{r_t\}) \propto L \cdot \pi(\Theta)$
- Sampling: Gibbs sampler alternating:
  1. Draw $\{Z_t\} \mid \Theta, \{r_t\}$ (forward-filtering backward-sampling)
  2. Draw $\Theta \mid \{Z_t\}, \{r_t\}$ (conjugate updates where possible)

Advantages for crypto:

- Full posterior uncertainty on regime probabilities
- Credible intervals for option prices (not just point estimates)
- Robust to short samples (important for altcoins with limited history)

Self-study: Eraker, Johannes, Polson (JF 2003) for MCMC estimation of SV-jump models.

Applications & Industry

From theory to practice: Deribit, DVOL, calibration

Deribit: The Crypto Options Market

Market dominance:

- $\sim$90% of all crypto options volume
- BTC + ETH options (European-style, cash-settled)
- Daily + weekly + monthly + quarterly expiries
- Inverse contracts: denominated *in* BTC/ETH

Microstructure:

- Central limit order book (CLOB)
- Maker/taker fees: 0.03%/0.03%
- Block trades for >25 BTC notional
- Mark price = model IV (not last trade)

Typical liquidity:

- ATM: tight spreads ($\sim$0.5% IV)
- OTM puts ($\delta < 0.1$): wide spreads ($\sim$5% IV)
- Short tenors: most liquid
- Open interest: $15–30B notional (2024–2025)

Data access:

- Free API: real-time quotes, trades
- Historical: Tardis, Kaiko, Genesis Vol
- Challenge: 24/7 $\rightarrow$ massive data volumes

Bitcoin DVOL: The "VIX of Crypto"

Definition: Deribit Volatility Index (DVOL)

$$\text{DVOL} = 100 \times \sqrt{\frac{2}{T} \sum_i \frac{\Delta K_i}{K_i^2} e^{rT} \, O(K_i) \cdot \frac{1}{T}}$$

Methodology (VIX-like):

- 30-day constant maturity
- Weighted strip of OTM calls and puts
- Interpolated between two nearest expiries
- Model-free: no parametric assumption

Typical DVOL levels:

Regime	DVOL
Calm bull market	40–55
Normal	55–75
Elevated	75–100
Crisis	100–150+
VIX (SPX) normal	12–20

BTC DVOL $\approx$ 3–5$\times$ VIX on average.

Key difference from VIX: DVOL includes 24/7 variance; no overnight/weekend gaps to exclude.

Empirical Vol Surfaces: Crypto Smile Characteristics

Key empirical features [7]:

Short tenors (1–7 days):

- Steep smile – pronounced U-shape
- Negative put skew (crash fear)
- Sometimes positive call skew ("FOMO")
- IV range: 40%–200%

Long tenors (30–90 days):

- Flatter smile
- Skew less pronounced
- Term structure: usually upward-sloping
- Inverts during stress (backwardation)

Compared to equity (SPX) smile:

- BTC smile is **more symmetric** (SPX has strong downside skew only)
- ATM IV is 3–5$\times$ higher
- Smile **flattens less** with maturity
- No clear "volatility skew" – more of a "volatility smile" (U-shape)

Implication: models need both left and right tail richness – Kou's asymmetric double-exponential is well-suited.

Calibration Walkthrough: Kou to BTC Options

Step-by-step procedure (Python notebook demo):

1. **Data:** pull Deribit BTC options snapshot via API
   $\to$ filter: $\delta \in [0.05, 0.95]$, volume > 0, tenors 1d–90d
2. **Initial guess:** $\sigma_0 = \text{ATM IV}/100$, $\lambda_0 = 5$, $p_0 = 0.4$, $\eta_{1,0} = 10$, $\eta_{2,0} = 5$
3. **Inner loop:** for given Θ, compute model prices via Kou CF + FFT
   $\to$ invert to model IVs via bisection on BS formula
4. **Outer loop:** minimize $\sum_i w_i(\sigma_i^{\text{model}} - \sigma_i^{\text{mkt}})^2$
   $\to$ differential evolution (scipy), then Nelder-Mead refinement
5. **Validate:** out-of-sample on held-out strikes/tenors

Constraints: $\sigma > 0$, $\lambda > 0$, $p \in (0, 1)$, $\eta_1 > 1$, $\eta_2 > 0$

Runtime: $\sim$3s on laptop (100 options, FFT with $N = 4096$).

Calibration Results: Model Comparison on IV Fits

BTC options, Deribit snapshot (typical results):

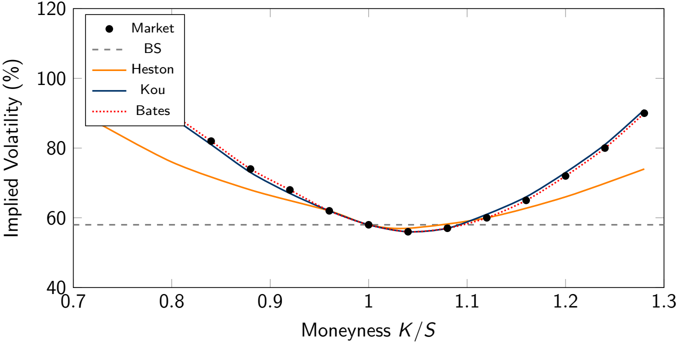

Kou and Bates both track the market smile closely; BS fails entirely at wings.

Model Selection: AIC, BIC, and Out-of-Sample

In-sample fit is necessary but not sufficient. Use information criteria and OOS testing.

Information criteria

$$\text{AIC} = -2\ln\hat{L} + 2k$$
$$\text{BIC} = -2\ln\hat{L} + k\ln n$$

where k = number of parameters, n = number of options, $\hat{L}$ = maximized likelihood.

Out-of-sample protocol:

1. Calibrate on day t options
2. Price day $t+1$ options using day-t parameters
3. Compute RMSE, MAE on day $t+1$
4. Repeat rolling over 250 trading days (= 250 calendar days for crypto)

Typical ranking for BTC:

$$\underbrace{\text{Kou}}_{\text{best OOS}} \approx \text{Bates} > \text{Heston} > \text{Merton} \gg \text{BS}$$

Practical Challenges in Crypto Option Pricing

Illiquid strikes:

- Deep OTM options: wide bid-ask, few trades
- Stale quotes (>minutes old) are common
- Filtering: discard options with $\delta < 0.02$ or spread $> 20\%$ IV
- Interpolation: SVI or SABR to fill gaps

24/7 markets:

- No close price – use 00:00 UTC snapshot?
- Volatility is intraday-heterogeneous
- Expiry: 08:00 UTC (Deribit) – not EOD
- Time-to-expiry: calendar days, not trading days

Infrastructure:

- Inverse contracts: $P_{\text{call}} = \max(S_T^{-1} - K^{-1}, 0)$
- Settlement: in crypto, not fiat
- Margin: denominated in BTC/ETH
- Interest rate r: use USDC lending rate? Treasury rate?

Model risk:

- Calibration instability near roll dates
- Parameter jumps across days
- Kou vs. Bates: different models for different tokens
- Regulatory uncertainty affects tail risk

Summary and Required Reading

Today we covered:

1. **Why BS fails:** heavy tails, skewness, volatility clustering in crypto
2. **Jump-diffusion:** Merton (symmetric), Kou (asymmetric) – CF + Carr-Madan FFT
3. **Stochastic volatility:** Heston (mean-reverting vol), Bates (SV + jumps)
4. **SVCJ:** correlated jumps in price and vol
5. **Equilibrium:** Li et al. – endogenous risk premia + liquidity risk
6. **Applications:** Deribit, DVOL, calibration, model selection

Required reading:

- [4] – Kou jump-diffusion (core paper)
- [2] – Heston stochastic volatility (essential)
- [1] – calibration to BTC/ETH options
- [5] – equilibrium pricing with liquidity risk

Next session (Day 2): DeFi – CFMMs, impermanent loss, lending protocol mathematics.

References I

[1] Saeid Babaei et al. *Pricing Options on Cryptocurrency Futures*. 2025. arXiv: `2506.14614 [q-fin.PR]`.

[2] Steven L. Heston. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options". In: *Review of Financial Studies* 6.2 (1993), pp. 327–343.

[3] Shimon Kogan, Igor Makarov, Marina Niessner, and Antoinette Schoar. "Are Cryptos Different? Evidence from Retail Trading". In: *Journal of Financial Economics* 159 (2024), p. 103893.

[4] Steven G. Kou. "A Jump-Diffusion Model for Option Pricing". In: *Management Science* 48.8 (2002), pp. 1086–1101.

[5] Zhe Li, Wei Zhang, Yi Zhang, and Haoran Yi. "Equilibrium Pricing of Bitcoin Options with Stochastic Volatility, Jumps, and Liquidity Risk". In: *Journal of Futures Markets* (2026).

[6] Igor Makarov and Antoinette Schoar. "Trading and Arbitrage in Cryptocurrency Markets". In: *Journal of Financial Economics* 135.2 (2020), pp. 293–319.

[7] Jovanka Lili Matić, Natalie Packham, and Lutz Schloegl. "Hedging Cryptocurrency Options". In: *Review of Derivatives Research* 26 (2023), pp. 91–133.