Option Pricing on Automated Market Maker Tokens

Philip Z. Maymin*

February 2026

Abstract

We derive the stochastic price process for tokens whose sole price discovery mechanism is a constant-product automated market maker (AMM) and develop a corresponding option pricing framework. When the net flow into an AMM pool follows a diffusion process, the token price follows a constant elasticity of variance (CEV) model with exponent $\beta = w$, where w is the numeraire weight in the pool. For the standard constant-product AMM ($w = 1/2$), this yields $\beta = 1/2$, nesting the Black–Scholes model as the limiting case of infinite liquidity. We derive closed-form European option prices using the non-central chi-squared distribution, introduce liquidity-adjusted Greeks, and extend the model to incorporate token emissions. The CEV structure produces a structural leverage effect and an implied volatility skew, both controlled by pool depth, with the pricing discrepancy relative to Black–Scholes scaling as $O(k^{-1})$: for shallow pools, out-of-the-money price differences can exceed 20%, while for deep pools they fall below 1%. Applied to Bittensor's dTAO system, a cross-sectional delta-hedged backtest across 82 subnets confirms that the two models' hedging performance diverges significantly with pool depth ($p = 0.006$).

Keywords: automated market makers, option pricing, constant elasticity of variance, Black–Scholes, Bittensor, decentralized finance, liquidity

JEL Classification: G12, G13, G14

*Fairfield University Dolan School of Business. Email: pmaymin@fairfield.edu.

1 Introduction

Automated market makers (AMMs) have become the dominant mechanism for token exchange in decentralized finance (DeFi), processing hundreds of billions of dollars in cumulative trading volume since the launch of Uniswap in 2018 [Adams et al., 2021]. Unlike traditional order-book exchanges where prices emerge from the interaction of discrete buy and sell orders, AMMs determine prices algorithmically from the ratio of token reserves held in liquidity pools. The most widely adopted design, the *constant-product* AMM, maintains the invariant $x \cdot y = k$ across two token reserves x and y, with the marginal exchange rate given by the reserve ratio [Angeris et al., 2021].

A new class of AMM-native tokens has emerged in which the AMM is not merely a trading venue but the *sole price discovery mechanism*. There is no order book, no off-chain market, and no external price oracle; the bonding curve *is* the market. The most prominent example is the Bittensor network's Dynamic TAO (dTAO) system, launched in February 2025, which assigns each of its subnets an "alpha" token traded exclusively through a dedicated constant-product AMM against the network's native TAO token [Bittensor Foundation, 2025]. With over 60 active subnets, a total staked value exceeding $3 billion, and a growing ecosystem of subnet operators seeking to hedge treasury exposure, pricing derivatives on these tokens is an increasingly practical concern. Bittensor's setting also provides the cleanest possible test environment for AMM price dynamics: because there is no competing venue, the observed price process is entirely determined by the AMM mechanics, making it possible to test the theory without confounding effects from external price discovery.

The standard approach to option pricing, the Black–Scholes model [Black and Scholes, 1973], assumes the underlying asset follows geometric Brownian motion (GBM), with log-normal returns and constant volatility. This assumption is justified when the asset trades on a deep, frictionless market where price changes are driven by exogenous information arrival. For AMM-native tokens, however, the price is *endogenously determined* by the AMM's bonding curve. Every trade mechanically shifts the reserve ratio, and the resulting price dynamics inherit the nonlinear structure of the AMM itself.

The central contribution of this paper is to derive the stochastic process governing AMM token prices from first principles and to develop the resulting option pricing framework. Our main result (Theorem 1) shows that when the net staking flow into an AMM pool follows a Brownian diffusion, the token price follows a *constant elasticity of variance* (CEV) model [Cox, 1975, Cox and Ross, 1996] with exponent β equal to the weight w of the numeraire token in the pool. For the standard constant-product AMM, $\beta = 1/2$.

This result has several important implications:

1. **Black–Scholes as a limiting case.** As pool liquidity $k \to \infty$, the CEV volatility parameter vanishes and the price becomes deterministic. For large but finite k, the

process approximates GBM, and Black–Scholes applies with an effective volatility that is inversely proportional to the square root of pool depth. The result gives a precise characterization of when standard pricing models are adequate.

2. **Liquidity-dependent volatility and the leverage effect.** The instantaneous volatility of the AMM token price is $\sigma(P) = \delta P^{\beta-1}$, where δ is proportional to the flow volatility and inversely proportional to pool depth. For $\beta = 1/2$, volatility decreases with the square root of price. This produces a structural leverage effect (negative correlation between price and volatility) that is a consequence of the bonding curve mechanics, not capital structure.

3. **Novel Greeks.** Beyond the standard sensitivities, we derive a "liquidity Greek" $\Lambda = \partial V/\partial k$ measuring option value sensitivity to pool depth, and an "emission Greek" $\mathcal{E} = \partial V/\partial e$ capturing sensitivity to the token emission rate.

4. **Quantifiable pricing discrepancy.** We provide closed-form expressions for the pricing discrepancy relative to Black–Scholes as a function of pool depth and moneyness (Figure 5), enabling practitioners to assess when the standard model is inadequate and by how much.

Our work connects three strands of literature. The first is the financial theory of AMMs, including analyses of impermanent loss [Loesch et al., 2021], AMM design [Angeris et al., 2021, 2022], and the relationship between AMM liquidity provision and option payoffs [Clark, 2021, Guillaume and Schroers, 2024]. The second is the CEV option pricing literature initiated by Cox [1975], with closed-form solutions developed by Schroder [1989] and subsequent extensions [Davydov and Linetsky, 2001, Larguinho et al., 2013]. The third is the nascent literature on Bittensor's tokenomics and decentralized AI markets [Bittensor Foundation, 2025].

The remainder of the paper is organized as follows. Section 2 surveys the related literature. Section 3 describes the institutional setting, with particular attention to Bittensor's dTAO mechanism. Section 4 derives the CEV price process from AMM fundamentals. Section 5 develops the option pricing framework, including closed-form solutions, Greeks, and emission extensions. Section 6 provides numerical analysis, Monte Carlo validation, and calibration to Bittensor data. Section 7 discusses limitations and extensions. Section 8 concludes.

2 Related Literature

AMM theory. The mathematical foundations of constant-function market makers were established by Angeris et al. [2021], who characterized the set of feasible trades and the connection between AMM prices and external market prices. Angeris et al.

[2022] extended this to optimal routing across multiple pools. Park [2023] identifies conceptual flaws in constant-product pricing, including persistent arbitrage and front-running vulnerabilities. The welfare properties and fee structures of AMMs are analyzed by Roughgarden [2024], while Cartea et al. [2024] develop a continuous-time model of AMM dynamics incorporating arbitrageurs and informed traders.

AMMs and options. Clark [2021] showed that a constant-product AMM liquidity provider is effectively writing a perpetual straddle, and Loesch et al. [2021] formalized impermanent loss as a variance-dependent cost. Guillaume and Schroers [2024] developed static hedging strategies using vanilla options, and **?** used the Black–Scholes formula to estimate divergence loss magnitudes. On-chain options protocols have motivated further pricing theory: Dave [2023] and Block Scholes and Panoptic [2025] analyze perpetual options in DeFi, while Singh [2025] survey the broader DeFi options ecosystem.

CEV models. The constant elasticity of variance model was introduced by Cox [1975] and extended by Cox and Ross [1996]. Empirical estimation and testing of the CEV model against equity option data were provided by Beckers [1980] and Emanuel and MacBeth [1982]. Closed-form European option prices were derived by Schroder [1989] using the non-central chi-squared distribution. Davydov and Linetsky [2001] provided eigenfunction-based pricing, and Larguinho et al. [2013] improved the numerical stability of the CEV formula. The CEV model generates implied volatility skew controlled by the elasticity parameter, making it a parsimonious alternative to stochastic volatility models [Heston, 1993]. In the CEV literature, the elasticity parameter β is typically an empirical quantity estimated from data; our contribution is to show that for AMM tokens, β is pinned by the pool design ($\beta = w$) rather than estimated.

Concurrent work. In independent and concurrent work, Hitier [2025] models LP portfolio value in constant-product AMMs under GBM. That paper assumes the external asset price follows GBM and derives LP value dynamics, whereas we derive the *endogenous* price process of a token that trades *only* through the AMM. The approaches are complementary: Hitier's applies when the AMM is one of many trading venues; ours applies when the AMM is the sole price discovery mechanism.

Our contribution. While the existing literature has explored AMMs as option-like instruments, no prior work has derived the stochastic process governing AMM token prices from the bonding curve mechanics. Our result, that AMM token prices follow a CEV process with exponent equal to the pool weight, provides this missing link and yields a complete option pricing framework for derivatives *on* AMM tokens.

3 Institutional Background

3.1 Constant-Product Automated Market Makers

A constant-product AMM maintains two token reserves (x, y) subject to the invariant

$$x \cdot y = k, \tag{1}$$

where $k > 0$ is constant during any individual swap. A trader who deposits Δx units of token X receives Δy units of token Y, determined by

$$(x + \Delta x)(y - \Delta y) = k \quad \Longrightarrow \quad \Delta y = \frac{y \cdot \Delta x}{x + \Delta x}. \tag{2}$$

The marginal price of token Y in terms of token X is

$$P = \frac{x}{y}, \tag{3}$$

obtained by differentiating the invariant. Equation (2) implies that the effective execution price for a trade of size Δx exceeds the marginal price by a slippage term of order $\Delta x/x$, making large trades progressively more expensive. The design was introduced by Buterin [2017] and formalized by Angeris et al. [2021].

Example 1 (A simple constant-product AMM)**.** Suppose a pool holds $x = 1{,}000$ TAO and $y = 40{,}000$ ALPHA, so $k = x \cdot y = 4 \times 10^7$ and the marginal price is $P = 1{,}000/40{,}000 = 0.025$ TAO per ALPHA. A trader who stakes 100 TAO receives

$$\Delta y = \frac{40{,}000 \times 100}{1{,}000 + 100} = 3{,}636.4 \text{ ALPHA}.$$

After the trade the reserves are $x = 1{,}100$ and $y = 36{,}363.6$, the invariant is still $k = 4 \times 10^7$, and the new price is $P' = 1{,}100/36{,}363.6 \approx 0.0302$. A deposit equal to 10% of the TAO reserve moved the price by 21%. This amplification of flows into price changes is the mechanism that generates the CEV dynamics derived in Section 4. Note that a pool ten times deeper ($k = 4 \times 10^9$) would produce correspondingly smaller price impact from the same trade, foreshadowing the Black–Scholes limit of infinite pool depth.

3.2 Generalized Constant-Function AMMs

The constant-product design is a special case of the *constant-weighted-product* family:

$$x^w \cdot y^{1-w} = K, \tag{4}$$

where $w \in (0,1)$ is the weight of token X (the numeraire) and $K > 0$. For $w = 1/2$, this reduces to the constant-product AMM with $K = \sqrt{k}$. Platforms such as Balancer implement arbitrary weights, enabling asymmetric exposure [Martinelli and Mushegian, 2019]. The marginal price under (4) is

$$P = \frac{1-w}{w} \cdot \frac{x}{y}. \tag{5}$$

3.3 Bittensor and Dynamic TAO

Bittensor is a decentralized network for AI services organized into *subnets*, each specializing in a particular machine learning task. Since February 2025, the network employs Dynamic TAO (dTAO), under which each subnet i maintains an independent constant-product AMM with reserves (x_i, y_i), where x_i is the TAO (native currency) reserve and y_i is the subnet-specific "alpha" (α_i) reserve [Bittensor Foundation, 2025].

Users "stake" TAO into a subnet by swapping TAO for alpha through the AMM, and "unstake" by swapping alpha back for TAO. The alpha price in TAO is $P_i = x_i/y_i$. Three features distinguish this setting from standard AMMs:

1. **No external market.** Alpha tokens trade exclusively through the on-chain AMM. There is no order book, no off-chain market, and no external price oracle. The AMM is the sole price discovery mechanism.

2. **Pool liquidity injection.** Each block (approximately every 12 seconds), the protocol injects TAO into the subnet's AMM reserve. The TAO allocated to subnet i is

$$e_{\tau,i} = E_{\text{block}} \cdot \frac{\max(S_i - L, 0)}{\sum_j \max(S_j - L, 0)}, \tag{6}$$

   where E_{block} is the total TAO block emission (currently 0.5 TAO per block, following the December 2025 halving from 1 TAO per block), S_i is the exponentially weighted moving average of net TAO flows into subnet i, and L is a lower threshold. Simultaneously, alpha is injected into the pool in proportion $\Delta\alpha_i = \Delta\tau_i/P_i$, preserving the current spot price while deepening liquidity [Bittensor Foundation, 2025]. This grows the invariant $k_i = x_i \cdot y_i$ over time without changing the price, mechanically dampening price volatility.

3. **Alpha participant emissions.** Independently of the pool injection, each subnet emits alpha to *participants* at a base rate of approximately 1 alpha per block, subject to its own halving schedule (both TAO and each alpha are capped at 21 million). At the end of each tempo (360 blocks), this participant alpha is distributed: 41% to miners, 41% to validators and their stakers, and 18% to the subnet owner. This alpha

does not enter the pool. It increases circulating supply outside the pool and can exert selling pressure if recipients swap their alpha back to TAO through the AMM.

Example 2 (Emission mechanics)**.** Suppose subnet i has reserves $x_i = 1{,}000$ TAO and $y_i = 40{,}000$ ALPHA, so $P_i = 0.025$ TAO per ALPHA. In a single block, the pool injection channel adds $\Delta\tau_i = 0.01$ TAO[1] and $\Delta\alpha_i = 0.01/0.025 = 0.4$ ALPHA to the reserves. The new reserves are $x_i = 1{,}000.01$ and $y_i = 40{,}000.4$, the invariant grows from $k = 4 \times 10^7$ to $k' \approx 4.0001 \times 10^7$, and the price is unchanged at $P_i = 0.025$. Separately, the protocol emits 1 ALPHA to participants (0.41 to miners, 0.41 to validators and stakers, 0.18 to the subnet owner). This alpha does not enter the pool but increases circulating supply. If recipients sell it through the AMM, it exerts downward pressure on the alpha price. Over one tempo (360 blocks, roughly 72 minutes), the pool injection deepens reserves by 3.6 TAO and 144 ALPHA, while 360 ALPHA is distributed to participants.

4 The Model

4.1 Setup and Notation

Consider a single AMM pool with reserves $(x(t), y(t))$ satisfying the constant-weighted-product invariant (4). Let $P(t) = \frac{1-w}{w} \cdot \frac{x(t)}{y(t)}$ denote the marginal price at time t. Define the *net flow process* $F(t)$ as the cumulative net TAO staked into the pool by time t.

Definition 1 (Stochastic Flow Process)**.** The net flow process satisfies

$$\mathrm{d}F(t) = \mu_F \mathrm{d}t + \sigma_F \mathrm{d}W(t), \tag{7}$$

where $\mu_F \in \mathbb{R}$ is the drift (expected net inflow rate), $\sigma_F > 0$ is the flow volatility, and $W(t)$ is a standard Brownian motion on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$.

The assumption that flow follows a Brownian diffusion is a continuous-time approximation to discrete staking and unstaking events. When individual staking amounts are small relative to pool size (the standard "many small traders" assumption), this is justified by the functional central limit theorem. We relax this assumption in Section 7 by considering jump-diffusion flows.

4.2 Reserve Dynamics

In the absence of emissions, the reserves evolve according to the AMM mechanics:

$$\mathrm{d}x(t) = \mathrm{d}F(t), \tag{8}$$

[1]The total emission is 0.5 TAO per block (post-halving). With roughly 60 active subnets competing via (6), 0.01 TAO per block is illustrative of an average subnet.

and the constant-weighted-product constraint (4) determines $y(t)$ implicitly:

$$y(t) = \left(\frac{K}{x(t)^w}\right)^{1/(1-w)}. \tag{9}$$

Differentiating via Itô's lemma yields

$$\mathrm{d}y = -\frac{w}{1-w} \cdot \frac{y}{x}\mathrm{d}x + \frac{w(2w-1)}{2(1-w)^2} \cdot \frac{y}{x^2}(\mathrm{d}x)^2. \tag{10}$$

4.3 Derivation of the Price Process

We now derive the central result: the stochastic differential equation governing the AMM token price.

Theorem 1 (AMM Token Price Process)**.** *Under the constant-weighted-product AMM (4) with net flow process (7), the marginal token price $P(t)$ satisfies the CEV stochastic differential equation*

$$\mathrm{d}P = \mu(P)\mathrm{d}t + \delta P^w \mathrm{d}W(t), \tag{11}$$

where the CEV exponent is $\beta = w$ (the numeraire weight), the volatility parameter is

$$\delta = \frac{1}{1-w}\left(\frac{1-w}{w}\right)^{1-w} K^{-1}\sigma_F, \tag{12}$$

and the drift is

$$\mu(P) = \frac{1}{1-w}\left(\frac{1-w}{w}\right)^{1-w} K^{-1}\mu_F P^w + \frac{w}{2(1-w)^2}\left(\frac{1-w}{w}\right)^{2(1-w)} K^{-2}\sigma_F^2 P^{2w-1}. \tag{13}$$

Proof. From (4), the price $P = \frac{1-w}{w} \cdot \frac{x}{y}$ can be expressed purely as a function of x:

$$P(x) = \frac{1-w}{w} \cdot x \cdot \left(\frac{x^w}{K}\right)^{1/(1-w)} = \frac{1-w}{w} \cdot K^{-1/(1-w)} \cdot x^{1/(1-w)}. \tag{14}$$

Let $\alpha \equiv 1/(1-w)$ and $B \equiv \frac{1-w}{w} \cdot K^{-\alpha}$, so that $P = Bx^\alpha$. By Itô's lemma,

$$\begin{aligned} \mathrm{d}P &= B\alpha x^{\alpha-1}\mathrm{d}x + \tfrac{1}{2}B\alpha(\alpha-1)x^{\alpha-2}(\mathrm{d}x)^2 \\ &= B\alpha x^{\alpha-1}(\mu_F \mathrm{d}t + \sigma_F \mathrm{d}W) + \tfrac{1}{2}B\alpha(\alpha-1)x^{\alpha-2}\sigma_F^2 \mathrm{d}t. \end{aligned} \tag{15}$$

Substituting $x = (P/B)^{1/\alpha}$:

$$x^{\alpha-1} = (P/B)^{(\alpha-1)/\alpha} = (P/B)^{1-1/\alpha} = (P/B)^w = B^{-w}P^w, \tag{16}$$

$$x^{\alpha-2} = (P/B)^{(\alpha-2)/\alpha} = (P/B)^{2w-1} = B^{1-2w}P^{2w-1}. \tag{17}$$

Collecting terms, the diffusion coefficient of $\mathrm{d}P$ is

$$B\alpha \cdot B^{-w}P^{w} \cdot \sigma_F = \alpha B^{1-w}\sigma_F \cdot P^{w}.$$

Defining $\delta \equiv \alpha B^{1-w}\sigma_F$ and noting $B^{1-w} = \left(\frac{1-w}{w}\right)^{1-w} K^{-\alpha(1-w)} = \left(\frac{1-w}{w}\right)^{1-w} K^{-1}$, we obtain $\delta = \frac{1}{1-w}\left(\frac{1-w}{w}\right)^{1-w} K^{-1}\sigma_F$. The drift follows analogously. □

Corollary 2 (Constant-Product AMM)**.** *For the standard constant-product AMM ($w = 1/2$, $K = \sqrt{k}$ where $k = xy$), the price process is*

$$\mathrm{d}P = \left(\frac{2\mu_F}{\sqrt{k}}\sqrt{P} + \frac{\sigma_F^2}{k}\right)\mathrm{d}t + \frac{2\sigma_F}{\sqrt{k}}\sqrt{P}\,\mathrm{d}W(t). \tag{18}$$

The CEV exponent is $\beta = 1/2$ and the volatility parameter is $\delta = 2\sigma_F/\sqrt{k}$.

Corollary 3 (Black–Scholes Limit)**.** *As pool depth $K \to \infty$, the volatility parameter $\delta \to 0$ and the price becomes deterministic. For large but finite K, with P near P_0, the process approximates GBM:*

$$\frac{\mathrm{d}P}{P} \approx \tilde{\mu}\,\mathrm{d}t + \sigma_{\text{eff}}\,\mathrm{d}W(t), \qquad \sigma_{\text{eff}} = \delta P_0^{w-1}, \tag{19}$$

and Black–Scholes applies with volatility $\sigma = \sigma_{\text{eff}}$.

Remark 1 (Elasticity spectrum)**.** The CEV exponent $\beta = w$ reveals a fundamental connection between AMM design and price dynamics:

- $w = 1/2$ (constant-product): $\beta = 1/2$, variance decreases with price. This is the standard Uniswap/Bittensor case.
- $w \to 1$ (pool dominated by numeraire): $\beta \to 1$, approaching GBM and Black–Scholes.
- $w \to 0$ (pool dominated by alpha): $\beta \to 0$, approaching the Bachelier (normal) model.

AMM designers thus implicitly select a volatility structure through their choice of pool weights.

4.4 Properties of the CEV Price Process

Proposition 4 (Volatility structure)**.** *Under (11), the instantaneous return volatility is*

$$\sigma_{\text{ret}}(P) = \delta P^{w-1} = \delta P^{\beta-1}. \tag{20}$$

For the constant-product AMM ($\beta = 1/2$), $\sigma_{\text{ret}}(P) = \delta/\sqrt{P}$: volatility is inversely proportional to the square root of price.

This property has a natural economic interpretation. When the alpha token price is high, the TAO reserve x is large (since $P \propto x^2/k$), meaning the pool is deep in TAO terms. A given staking flow $\mathrm{d}F$ then produces a smaller proportional change in x, hence a smaller proportional price impact. Conversely, when the price is low, the TAO reserve is shallow, and the same flow produces larger price swings.

Proposition 5 (Implied volatility skew)**.** *The CEV model with $\beta < 1$ generates a negative implied volatility skew: out-of-the-money puts have higher Black–Scholes implied volatility than out-of-the-money calls. For AMM tokens, this skew steepens as pool depth decreases.*

Proof. This follows from the well-known result that the CEV model with $\beta < 1$ produces a downward-sloping implied volatility curve as a function of strike [Cox and Ross, 1996, Davydov and Linetsky, 2001]. The skew magnitude is controlled by $1 - \beta$ and δ; both the departure from GBM (via $\beta = w < 1$) and the volatility level (via $\delta \propto K^{-1}$) increase as pool depth decreases. □

Proposition 6 (Boundary behavior at zero)**.** *For $\beta = 1/2$, the CEV process (11) has $P = 0$ as an* absorbing *boundary. The boundary classification follows from the Feller test: the scale function $s(P) \propto P^{-2r/\delta^2}$ is integrable near zero, while the speed measure diverges, so zero is an entrance boundary under the risk-neutral measure but is attainable in finite time from the physical measure when δ is large relative to r. Economically, $P = 0$ corresponds to a fully drained TAO reserve ($x = 0$), at which point the AMM cannot quote a price. Once the TAO reserve is exhausted, no further trades are possible without external injection (e.g., emissions).*

Remark 2 (Leverage effect)**.** The CEV structure with $\beta < 1$ generates a negative correlation between the price level and return volatility: as P falls, $\sigma_{\text{ret}}(P) = \delta P^{\beta-1}$ rises. In equity markets, this "leverage effect" is typically attributed to increased financial leverage as firm value declines [Black and Scholes, 1973]. For AMM tokens, the mechanism is purely structural. When the alpha price falls, the TAO reserve $x = \sqrt{kP}$ decreases, making the pool shallower in TAO terms. The same dollar-equivalent staking flow then moves the price proportionally more. The AMM bonding curve thus provides a first-principles derivation of the leverage effect, grounded in market microstructure rather than capital structure.

5 Option Pricing

5.1 Risk-Neutral Dynamics

Under the risk-neutral measure $\mathbb{Q}$, the CEV price process becomes

$$\mathrm{d}P = rP\mathrm{d}t + \delta P^{\beta}\mathrm{d}W^{\mathbb{Q}}(t), \tag{21}$$

where r is the risk-free rate and $W^{\mathbb{Q}}$ is a $\mathbb{Q}$-Brownian motion. The change of measure from $\mathbb{P}$ to $\mathbb{Q}$ is effected via the Girsanov kernel $\theta(P) = [\mu(P) - rP]/(\delta P^{\beta})$. For $\beta \in (0,1)$ and P bounded away from zero, θ is bounded and the Novikov condition $\mathbb{E}[\exp(\frac{1}{2}\int_0^T \theta^2 \mathrm{d}t)] < \infty$ holds. Near $P = 0$, θ can diverge, but this region has zero probability under the risk-neutral measure when the boundary is absorbing (Proposition 6). The existence of the equivalent martingale measure for the CEV process with $\beta \in (0,1)$ is established by Davydov and Linetsky [2001].

The risk-neutral pricing argument requires approximate replicability of contingent claims by dynamic trading in the underlying. For AMM tokens, this is imperfect due to slippage: a trade of size Δx incurs a price impact of order $\Delta x/x$. When individual hedge trades are small relative to the TAO reserve (i.e., $|\Delta x| \ll x$), the slippage cost is second-order and the replication error is bounded. We quantify this error in Section 5.5 and show it scales as k^{-2}, becoming negligible for deep pools.

5.2 European Option Pricing Formula

The European call price under the CEV model was derived by Cox [1975] and refined by Schroder [1989]. For $\beta < 1$ (which includes the constant-product case $\beta = 1/2$), the price of a European call with strike K_{str} and maturity T is:

Theorem 7 (CEV Call Price)**.**

$$C(P, K_{\text{str}}, T) = P\left[1 - \chi^2(a;\, b+2,\, c)\right] - K_{\text{str}}\mathrm{e}^{-rT}\chi^2(c;\, b,\, a)\,, \tag{22}$$

where $\chi^2(x;\, n,\, \lambda)$ denotes the cumulative distribution function of the non-central chi-squared distribution with n degrees of freedom and non-centrality parameter λ, and

$$\kappa = \frac{2r}{\delta^2(1-\beta)(e^{2r(1-\beta)T} - 1)}, \tag{23}$$

$$c = \kappa P^{2(1-\beta)} e^{2r(1-\beta)T}, \tag{24}$$

$$a = \kappa K_{\text{str}}^{2(1-\beta)}, \tag{25}$$

$$b = \frac{1}{1-\beta} - 1. \tag{26}$$

For the constant-product AMM ($\beta = 1/2$), these simplify to $b = 1$ and

$$\kappa = \frac{2r}{\delta^2(e^{rT} - 1)/2} = \frac{4r}{\delta^2(e^{rT} - 1)}, \quad c = \kappa P e^{rT}, \quad a = \kappa K_{\text{str}}. \tag{27}$$

The European put price follows from put-call parity:

$$\Pi(P, K_{\text{str}}, T) = C(P, K_{\text{str}}, T) - P + K_{\text{str}}\mathrm{e}^{-rT}. \tag{28}$$

Remark 3 (Convergence to Black–Scholes)**.** As $\beta \to 1$ with $\delta P_0^{\beta-1} = \sigma$ held fixed, the CEV call price (22) converges to the Black–Scholes formula

$$C_{\mathrm{BS}} = P\,\Phi(d_1) - K_{\mathrm{str}}\mathrm{e}^{-rT}\Phi(d_2), \tag{29}$$

where $d_{1,2} = \frac{\ln(P/K_{\mathrm{str}})+(r\pm\sigma^2/2)T}{\sigma\sqrt{T}}$. This provides a formal justification for using Black–Scholes when AMM liquidity is large.

5.3 The Liquidity-Adjusted Black–Scholes Formula

To make the connection to Black–Scholes explicit, we decompose the CEV call price as

$$C_{\mathrm{CEV}} = C_{\mathrm{BS}}(\sigma_{\mathrm{eff}}) + \Lambda_C, \tag{30}$$

where $\sigma_{\mathrm{eff}} = \delta P^{w-1}$ is the effective volatility at the current price, and Λ_C is the *liquidity correction*, i.e., the residual difference due to the price-dependent volatility structure.

Proposition 8 (Liquidity correction)**.** *To first order in* $(1-\beta)$*, the liquidity correction for an at-the-money call is*

$$\Lambda_C \approx -\frac{1}{2}(1-\beta)\,\sigma_{\mathrm{eff}}\,P\,\sqrt{T}\,\phi(0)\left[\sigma_{\mathrm{eff}}\sqrt{T} + \frac{1}{\sigma_{\mathrm{eff}}\sqrt{T}}\ln\!\left(\frac{P}{K_{\mathrm{str}}}\right)\right], \tag{31}$$

where ϕ *is the standard normal density. For the constant-product AMM* ($\beta = 1/2$)*, this is* $\Lambda_C \approx -\frac{1}{4}\sigma_{\mathrm{eff}}^2 PT\phi(0)$.

The liquidity correction is negative for calls (and positive for puts with low strikes), consistent with the negative skew generated by $\beta < 1$. Its magnitude scales with $\delta^2 \propto K^{-2}$, confirming that it vanishes rapidly as pool depth increases.

5.4 Liquidity-Adjusted Greeks

The standard option Greeks are modified by the CEV structure. We also introduce two new sensitivities specific to AMM tokens.

Definition 2 (AMM Greeks)**.** The *CEV delta* and *CEV gamma* of a European call are

$$\Delta_{\mathrm{CEV}} = \frac{\partial C}{\partial P} = 1 - \chi^2(a;\, b+2,\, c) + P\frac{\partial}{\partial P}\left[1 - \chi^2(a;\, b+2,\, c)\right] - K_{\mathrm{str}}e^{-rT}\frac{\partial}{\partial P}\chi^2(c;\, b,\, a), \tag{32}$$

$$\Gamma_{\mathrm{CEV}} = \frac{\partial^2 C}{\partial P^2}. \tag{33}$$

The *liquidity Greek* is

$$\Lambda = \frac{\partial C}{\partial k} = \frac{\partial C}{\partial \delta}\cdot\frac{\partial \delta}{\partial k}, \tag{34}$$

measuring the sensitivity of the option price to changes in pool depth. For the constant-product AMM, $\delta = 2\sigma_F/\sqrt{k}$, so $\partial\delta/\partial k = -\sigma_F/k^{3/2}$, and $\Lambda < 0$: deeper pools reduce option value by compressing volatility.

The *emission Greek* is

$$\mathcal{E} = \frac{\partial C}{\partial e}, \tag{35}$$

measuring sensitivity to the emission rate e that governs the growth of k over time (see Section 5.6). Using the integrated variance (42), $\mathcal{E}$ can be computed via the chain rule: $\mathcal{E} = (\partial C/\partial \bar{v}^2)(\partial \bar{v}^2/\partial \dot{k})(\partial \dot{k}/\partial e)$. Since $\partial \bar{v}^2/\partial \dot{k} < 0$ and $\partial C/\partial \bar{v}^2 > 0$ (option prices increase with variance), the emission Greek is negative: higher emissions reduce option value by deepening the pool over the option's lifetime.

5.5 Hedging Error from AMM Friction

Delta-hedging an option on an AMM token requires trading through the AMM, incurring slippage. For a hedge trade of ΔP units of alpha, the slippage cost in a constant-product AMM is approximately

$$S(\Delta P) \approx \frac{(\Delta P)^2}{2k/P^2} = \frac{P^2(\Delta P)^2}{2k}. \tag{36}$$

Over a hedging interval Δt with rebalancing, the cumulative expected hedging cost is

$$\mathbb{E}\left[\int_0^T S(\mathrm{d}\Delta)\right] \approx \frac{P^2\Gamma^2\sigma_{\text{ret}}^2}{2k} \cdot T, \tag{37}$$

which scales as k^{-2}, an additional friction cost beyond the standard model. This cost should be added to the option price as a *replication premium*.

Proposition 9 (Replication premium)**.** *The replication premium for a European call on a constant-product AMM token is bounded by*

$$R \leq \frac{\delta^2 P^{2\beta}}{2k} \cdot \mathbb{E}^{\mathbb{Q}}\left[\int_0^T \Gamma_{\text{CEV}}^2(t) P(t)^2 \mathrm{d}t\right]. \tag{38}$$

As $k \to \infty$, $R \to 0$ and exact replication is recovered.

Remark 4 (Replication–relevance tension)**.** A conceptual tension arises: the CEV correction to Black–Scholes is largest for shallow pools (small k), but the replication premium is also largest for small k since it scales as k^{-2}. In the shallow-pool regime where the model matters most, the replication argument is weakest. For such pools, the CEV prices in Theorem 7 should be interpreted as *fair value benchmarks* under the idealized continuous-hedging framework, rather than strict arbitrage-free prices. In practice, a market maker would add the replication premium (38) to account for slippage, widening

the bid-ask spread. For deep pools, both the CEV correction and the replication premium are negligible, and Black–Scholes applies directly.

5.6 Extension: Token Emissions

When emissions inject TAO and alpha into the pool at rates e_τ and e_α per unit time, the pool invariant grows deterministically:

$$\frac{\mathrm{d}k}{\mathrm{d}t} = y(t)e_\tau + x(t)e_\alpha. \tag{39}$$

At equilibrium, with P near P_0 and reserves approximately constant, this simplifies to

$$k(t) \approx k_0 + (y_0 e_\tau + x_0 e_\alpha)t \equiv k_0 + \dot{k} \cdot t. \tag{40}$$

The CEV volatility parameter becomes time-dependent:

$$\delta(t) = \frac{2\sigma_F}{\sqrt{k(t)}} = \frac{2\sigma_F}{\sqrt{k_0 + \dot{k}t}}. \tag{41}$$

Proposition 10 (Option pricing with emissions)**.** *Under deterministically time-varying $\delta(t)$ and the assumption that $k(t)$ evolves slowly relative to the option horizon (so that the time-change from calendar time to "variance time" is approximately deterministic), the CEV call price formula (22) remains valid with $\delta^2 T$ replaced by the integrated variance:*

$$\bar{v}^2 = \int_0^T \delta(t)^2 \mathrm{d}t = 4\sigma_F^2 \int_0^T \frac{\mathrm{d}t}{k_0 + \dot{k}t} = \frac{4\sigma_F^2}{\dot{k}} \ln\left(1 + \frac{\dot{k}T}{k_0}\right). \tag{42}$$

As $\dot{k} \to 0$ (no emissions), $\bar{v}^2 \to 4\sigma_F^2 T/k_0 = \delta_0^2 T$, recovering the constant case.

Emissions have a dampening effect: higher emission rates increase $\dot{k}$, reducing $\bar{v}^2$ and hence option prices. This captures the intuition that growing liquidity compresses the range of possible price outcomes.

Remark 5 (Emissions as a dividend yield)**.** The emission effect can be interpreted as an effective continuous dividend yield. Expanding $\bar{v}^2$ for small $\dot{k}T/k_0$:

$$\bar{v}^2 \approx \delta_0^2 T \left(1 - \frac{\dot{k}T}{2k_0} + \cdots\right),$$

so the integrated variance is reduced by a factor that is first-order in $\dot{k}/k_0$. Defining an effective dividend yield $q_{\text{eff}} = \dot{k}/(2k_0)$, the emission-adjusted CEV price approximately equals the zero-emission price on an underlying with continuous yield q_{eff}. In practice, one can estimate $\dot{k}$ from the emission schedule and discount the option value accordingly.

Note that Bittensor's emission allocation (6) introduces a feedback loop not captured by the deterministic $\dot{k}$ assumption: high staking flows raise S_i, which increases the subnet's emission share $e_{\tau,i}$, deepening the pool, compressing volatility, and potentially attracting further staking. Modeling this endogeneity would require a stochastic $\dot{k}$ coupled to the flow process, an extension we leave for future work.

6 Numerical Analysis

6.1 Parameter Calibration

We calibrate the model to Bittensor subnet data as of February 2026, retrieved from the Taostats API. We select three representative subnets spanning the range of pool depths observed across the network:

Table 1: Representative Bittensor subnet parameters (median values over the sample period August 2025 to February 2026). Reserves are in native token units; $\hat{\sigma}_F$ is the annualized standard deviation of daily TAO reserve changes.

Subnet	x_0 (TAO)	y_0 (α)	k ($\times 10^9$)	P_0 (TAO/α)	$\hat{\sigma}_F$
Shallow (SN58)	4,019	1,835,800	7.4	0.0022	2,293
Medium (SN1)	22,568	2,340,831	52.8	0.0096	3,571
Deep (SN3)	54,445	2,151,385	117.1	0.0253	8,250

The flow volatility $\hat{\sigma}_F$ is estimated from the standard deviation of daily net TAO reserve changes, annualized ($\hat{\sigma}_F = \hat{s}_{\Delta F} \cdot \sqrt{365}$, where $\hat{s}_{\Delta F}$ is the sample standard deviation of daily flow changes). This estimator is consistent under the diffusion assumption and can be computed directly from on-chain reserve data. The risk-free rate is set to $r = 5\%$ (approximate stablecoin lending rate in DeFi).

Remark 6 (Distributional properties of staking flows)**.** Shapiro–Wilk tests reject the normality of daily TAO reserve changes at the 5% level for all 98 subnets in our sample. The median excess kurtosis is 10.7 (range $[0.6, 189.6]$) and the median skewness is -0.62, indicating heavy-tailed, left-skewed flow distributions. The Brownian diffusion assumption (Definition 1) is thus an approximation, as is standard for continuous-time models applied to discrete financial data. The heavy tails are consistent with occasional large staking events ("whale" trades) and support the jump-diffusion extension discussed in Section 7. That the cross-sectional backtest (Section 6.6) still confirms the predicted relationship between pool depth and hedging performance suggests the CEV framework is robust to moderate departures from normality.

Remark 7 (Testable predictions)**.** Even in the absence of traded options, the CEV model generates testable predictions about the physical price process. The return variance

should be proportional to $P^{2(\beta-1)} = P^{-1}$ for the constant-product case, a relationship that can be estimated from realized variance regressions on price levels using on-chain data.

6.2 Monte Carlo Validation

To validate the closed-form CEV pricing formula, we simulate the reserve dynamics directly. For each scenario:

1. Generate $N = 100{,}000$ paths of the flow process $\mathrm{d}F = \mu_F \mathrm{d}t + \sigma_F \mathrm{d}W$ using Euler–Maruyama discretization with time step $\Delta t = 1/(365 \times 24)$ (hourly).
2. At each step, update TAO reserve: $x_{t+\Delta t} = x_t + \Delta F_t$.
3. Compute alpha reserve from the invariant: $y_{t+\Delta t} = k/x_{t+\Delta t}$.
4. Compute terminal price: $P_T = x_T/y_T = x_T^2/k$.
5. Price the call as $C_{\mathrm{MC}} = e^{-rT}\mathbb{E}[\max(P_T - K_{\mathrm{str}}, 0)]$.

Figure 1 illustrates the qualitative difference between CEV and GBM sample paths, driven by the same Brownian increments. The CEV paths exhibit heteroskedasticity: volatility compresses as price rises and expands as price falls, a direct consequence of the $P^{1/2}$ volatility structure.

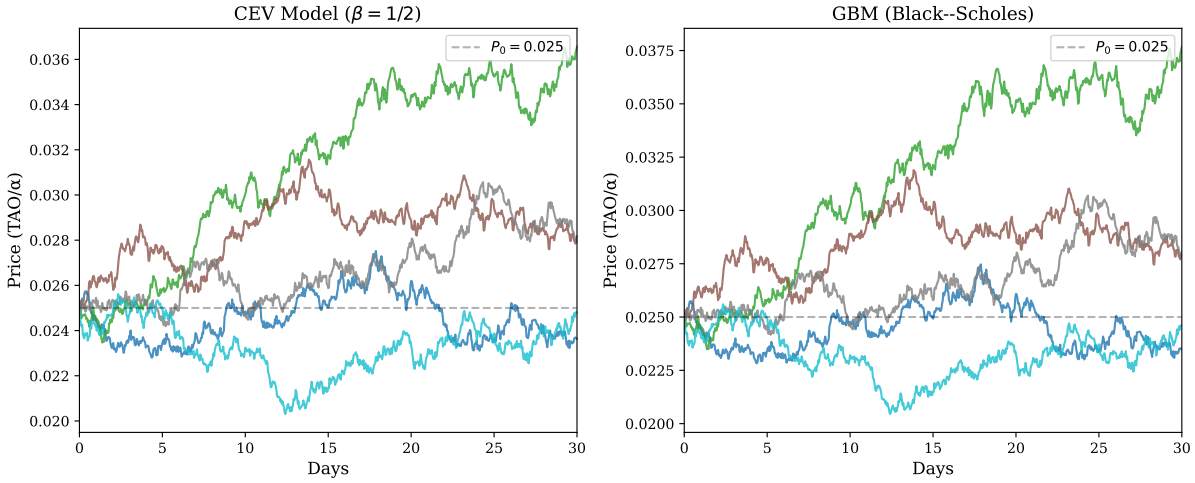

Figure 1: Simulated price paths under the CEV model ($\beta = 1/2$, left) and GBM (right), using identical Brownian increments. Illustrative parameters: $P_0 = 0.025$, $k = 10^6$, $\sigma_F = 48.7$, $T = 30$ days. The CEV paths exhibit asymmetric volatility: larger swings at low prices, dampened moves at high prices.

Figure 2 compares the Monte Carlo prices with the CEV closed-form formula across strikes for two pool depths ($T = 30$ days). For the deep pool ($k = 10^9$), the maximum

absolute deviation is less than 0.5% of the spot price. For the shallow pool ($k = 10^6$), the MC prices systematically exceed the CEV formula by 1–3% of spot. This bias arises because the Euler–Maruyama discretization of the reserve process is less accurate when the flow volatility σ_F is large relative to the reserve level $x_0 = \sqrt{kP_0}$: for the shallow pool, a single hourly flow shock can represent several percent of the reserve, violating the small-increment assumption underlying the diffusion limit. The bias is positive (MC above CEV) because the reflecting boundary at $x > 0$ truncates large negative flow shocks, creating a slight upward skew in terminal prices.

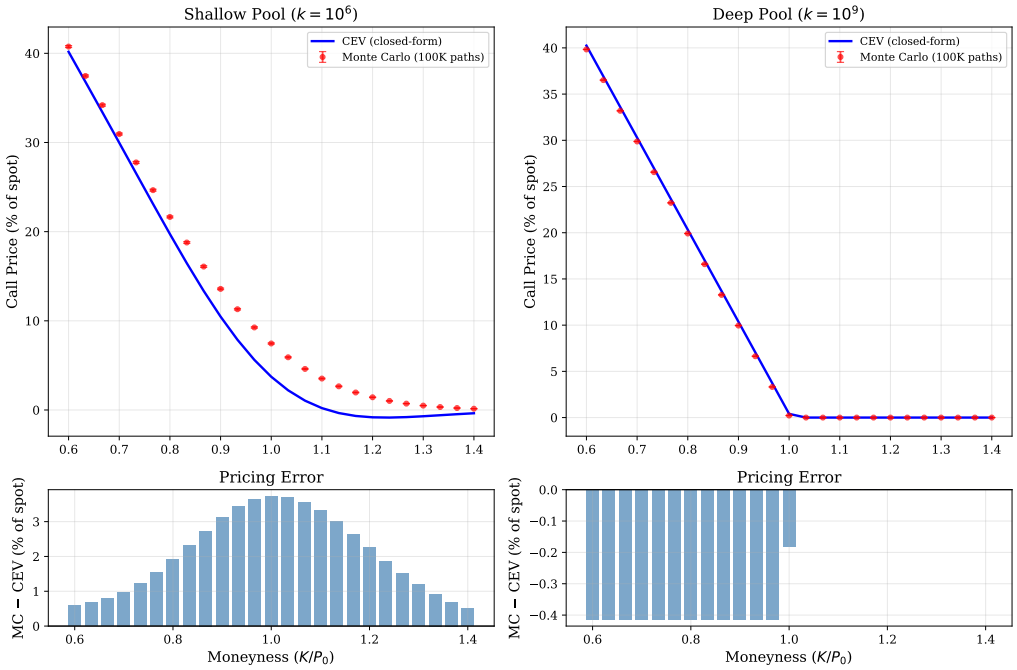

Figure 2: Monte Carlo validation of the CEV pricing formula for a shallow pool ($k = 10^6$, left) and a deep pool ($k = 10^9$, right). Top: closed-form CEV call prices (line) vs. Monte Carlo estimates with 95% confidence intervals (points). Bottom: pricing error (MC minus CEV) as a percentage of spot. The shallow pool shows a systematic positive bias of 1–3% of spot, reflecting the Euler–Maruyama discretization error that is amplified when flow volatility is large relative to the reserve. The deep pool shows near-perfect agreement (errors $< 0.5\%$). Illustrative parameters: $P_0 = 0.025$, $\sigma_F = 48.7$, $T = 30$ days, 100,000 paths.

6.3 Comparative Statics

Figure 3 illustrates how the call option price varies with pool depth k for an at-the-money option with $T = 30$ days. As k increases, the CEV price converges to Black–Scholes from below (for calls), consistent with the negative liquidity correction (Proposition 8). The convergence is approximately $O(k^{-1})$.

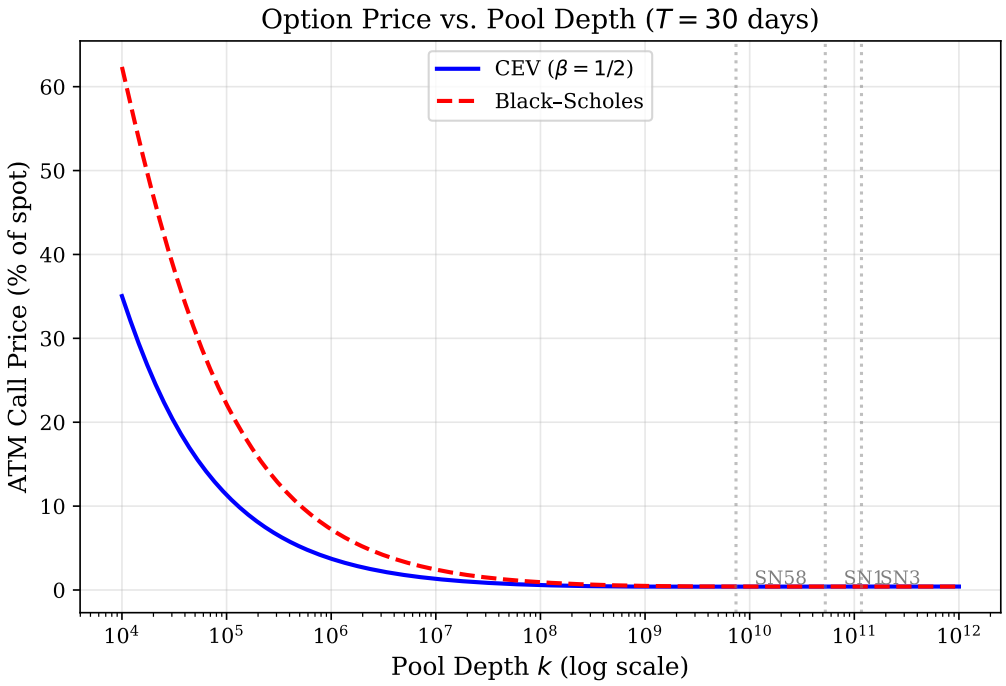

Figure 3: At-the-money call option price (as % of spot) vs. pool depth k. The CEV price (solid blue) converges to Black–Scholes (dashed red) as pool depth increases. Vertical dotted lines mark the three calibrated subnets from Table 1; at these depths, both models produce similar prices. Illustrative parameters: $P_0 = K = 0.025$, $\sigma_F = 48.7$, $T = 30$ days, $r = 5\%$.

Figure 4 plots the Black–Scholes implied volatility as a function of moneyness for different pool depths. To clearly illustrate the skew structure, we use illustrative pool depths ($k = 5 \times 10^5$ to 5×10^7) that span the range where the CEV correction is most visible; the calibrated subnets in Table 1 lie in the deeper portion of this range. Shallow pools generate a pronounced negative skew, which flattens as k increases. This is a structural prediction of the model that could be tested against observed option prices if and when options markets for AMM tokens develop.

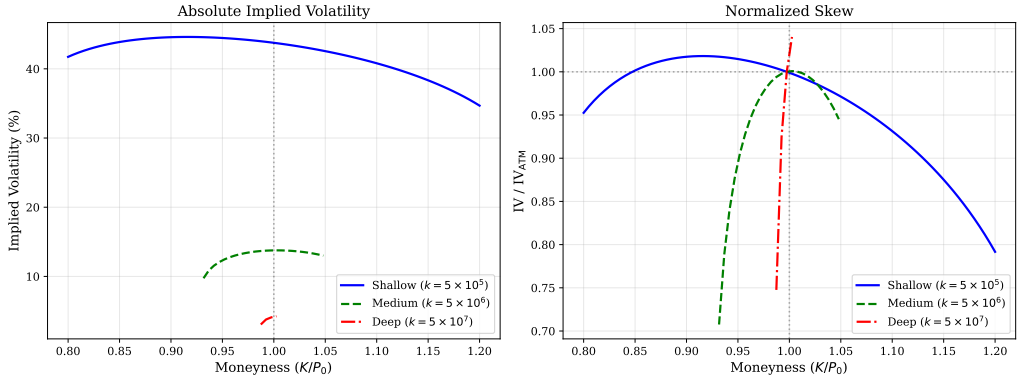

Figure 4: Left: absolute Black–Scholes implied volatility extracted from CEV prices. Right: implied volatility normalized by the at-the-money level, isolating the skew shape. Shallower pools produce both higher absolute volatility and steeper negative skew. Illustrative parameters: $P_0 = 0.025$, $\sigma_F = 48.7$, $T = 90$ days.

6.4 Pricing Discrepancy Relative to Black–Scholes

Figure 5 plots the CEV implied volatility smile for 90-day European options at three pool depths, with Black–Scholes (flat implied vol) shown for comparison. Two competing effects shape the smile:

1. **Leverage effect.** Because $\sigma(P) = \delta P^{\beta-1}$ increases as P falls, the CEV model assigns higher implied volatility to strikes below spot. This produces a left skew: Black–Scholes underprices out-of-the-money puts and overprices out-of-the-money calls.

2. **Absorbing barrier.** The CEV process with $\beta = 1/2$ reaches $P = 0$ in finite time (Proposition 6), truncating the left tail. For deep out-of-the-money puts, this barrier effect dominates, *reducing* implied volatility below the ATM level.

For shallow pools ($k = 10^5$), the leverage effect dominates across most of the moneyness range, producing a pronounced left skew with implied volatility up to 10% above ATM for moderately out-of-the-money puts. For medium pools ($k = 5 \times 10^5$), the barrier effect emerges for deep out-of-the-money puts ($K/P_0 < 0.83$), creating a smile that peaks below ATM and declines in both directions. For deep pools ($k = 10^6$), the smile collapses toward the Black–Scholes flat line, consistent with the $k \to \infty$ limiting behavior.

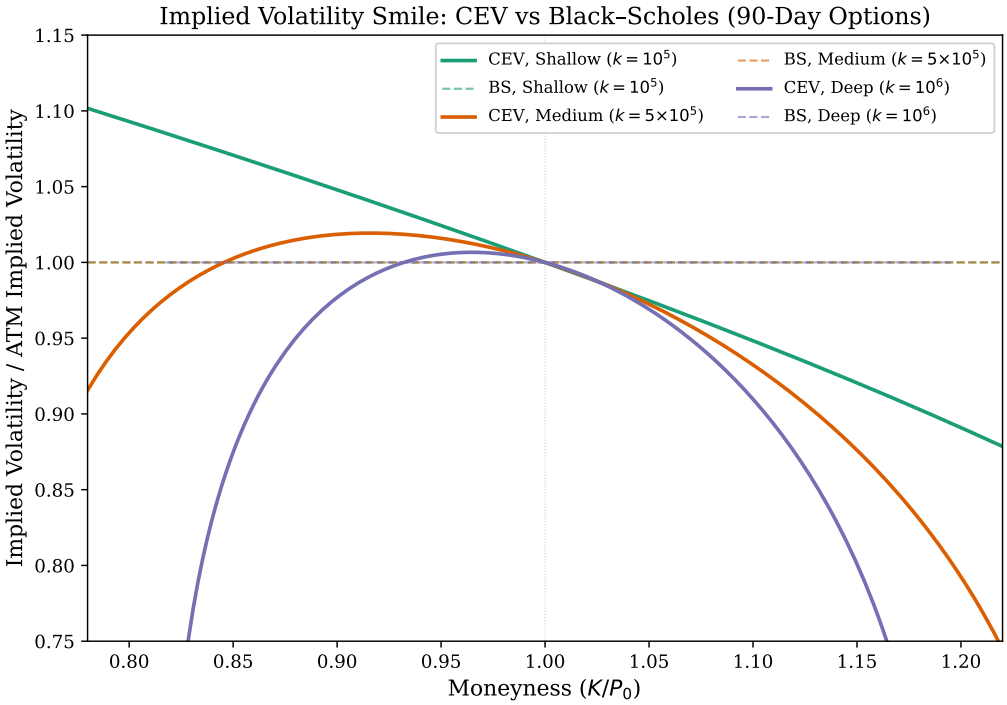

Figure 5: Implied volatility smile under the CEV model (solid curves) vs. Black–Scholes (dashed, flat at 1.0), normalized by ATM implied vol, for 90-day European options at three pool depths. The shallow pool ($k = 10^5$) exhibits a pronounced left skew driven by the leverage effect. The medium pool ($k = 5 \times 10^5$) shows a smile: the leverage effect elevates near-ATM puts, but the absorbing barrier at $P = 0$ depresses deep OTM puts. The deep pool ($k = 10^6$) collapses toward the BS flat line. Illustrative parameters: $P_0 = 0.025$, $\sigma_F = 48.7$.

Figure 6 compares the CEV and Black–Scholes delta and gamma. For an ATM call with strike $K = 0.025$, the CEV delta is steeper for low prices (reflecting higher volatility) and flatter for high prices. Both gammas peak below the strike (a consequence of the high volatility regime), but the CEV gamma is sharper and more concentrated, peaking further below the strike than BS.

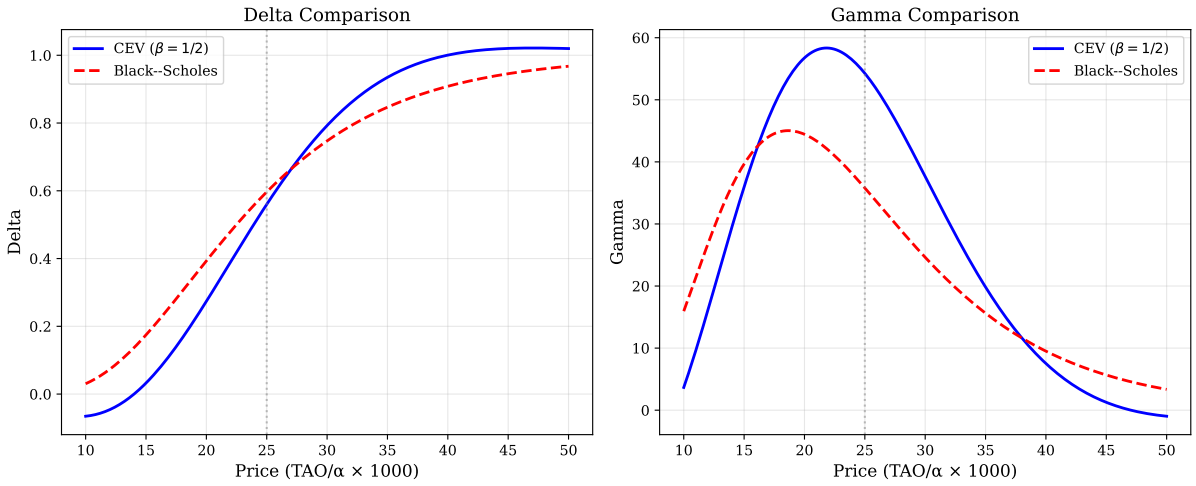

Figure 6: Comparison of CEV ($\beta = 1/2$) and Black–Scholes Greeks for an ATM European call ($K = 0.025$) on a shallow pool. Left: delta. Right: gamma. Both gammas peak below the strike, but the CEV gamma is sharper and peaks further below, reflecting its concentration in the high-volatility (low-price) region. Illustrative parameters: $k = 5 \times 10^5$, $\sigma_F = 48.7$, $T = 90$ days.

6.5 Effect of Emissions

Figure 7 shows how token emissions affect option prices. Using the integrated variance formula (42), we plot call prices as a function of T for different emission rates. Higher emissions compress option prices, particularly at longer maturities where the cumulative liquidity deepening is greatest.

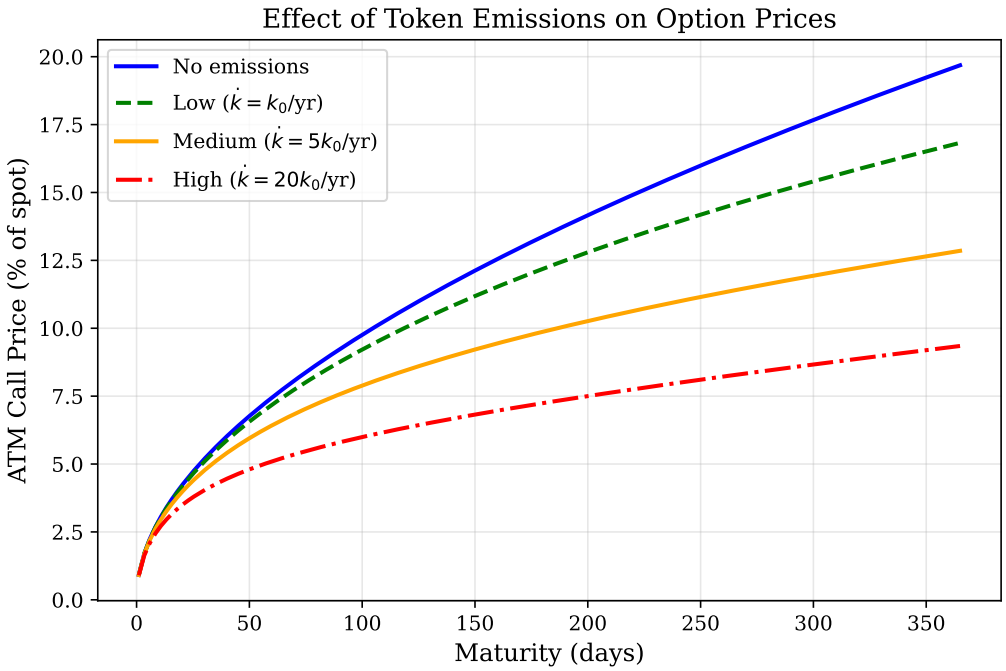

Figure 7: ATM call option price (as % of spot) vs. maturity under different emission rates, expressed as multiples of the initial pool invariant per year. Higher emissions deepen the pool over time, compressing volatility and reducing option prices at longer maturities. The effect is most pronounced for shallow pools with high emission-to-liquidity ratios. Illustrative parameters: $P_0 = K = 0.025$, $k_0 = 5 \times 10^5$, $\sigma_F = 48.7$.

Finally, Figure 8 plots the liquidity Greek $\Lambda = \partial C/\partial k$ as a function of pool depth. The Greek is negative throughout (deeper pools reduce option value) and decays rapidly, indicating that liquidity sensitivity is primarily a concern for shallow pools.

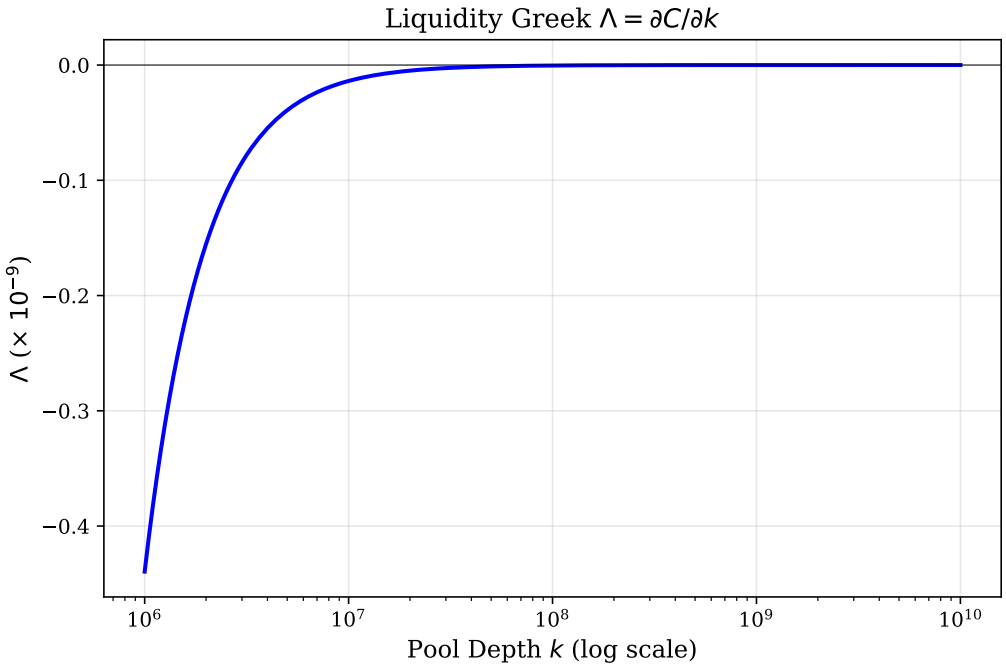

Figure 8: The liquidity Greek $\Lambda = \partial C/\partial k$ for an ATM call as a function of pool depth. Negative values indicate that increasing pool depth reduces option value. The sensitivity is concentrated in shallow pools and becomes negligible for $k > 10^9$. Illustrative parameters: $P_0 = K = 0.025$, $\sigma_F = 48.7$, $T = 30$ days.

6.6 Empirical Backtest

We conduct a cross-sectional backtest across all active Bittensor subnets to test whether the theoretical divergence between CEV and Black–Scholes pricing varies systematically with pool depth. Using daily on-chain data from 98 subnets retrieved via the Taostats API (August 2025 to February 2026), we execute the following procedure for each subnet and each rolling start date t:

1. Estimate $\hat{\sigma}_F$ from the trailing 14-day standard deviation of daily TAO reserve changes, annualized.

2. Compute the pool invariant $k_t = x_t \cdot y_t$ from observed reserves.

3. Sell a 14-day ATM European call ($K = P_t$) at the model price under both the CEV model ($\beta = 1/2$, $\delta_t = 2\hat{\sigma}_F/\sqrt{k_t}$) and Black–Scholes with matched ATM volatility ($\sigma_{\text{eff}} = \delta_t P_t^{-1/2}$).

4. Delta-hedge daily for 14 days using each model's delta, updating k and recomputing deltas from observed reserves at each rebalance.

5. At expiry, compute the hedged P&L: premium collected plus cumulative hedge gains minus the realized payoff $\max(P_{t+14} - K, 0)$.

We aggregate each subnet's trades into a single mean absolute hedging error (MAE, as % of spot) for each model. Because the 14-day option windows overlap, per-subnet MAE estimates exhibit serial correlation; we address this by using the cross-sectional regression (one observation per subnet), which is free of this overlap bias. After filtering 16 subnets with degenerate price paths (MAE $> 50\%$, typically from near-zero reserves or extreme price dislocations), 82 subnets remain.

Figure 9 presents the cross-sectional results. The left panel plots each subnet's CEV hedging error against its BS hedging error, with color indicating pool depth $\log_{10}(k)$. Most points cluster near the 45-degree line. The right panel quantifies the relationship between relative hedging performance and pool depth. An OLS regression of the CEV/BS error ratio on $\log_{10}(k)$ yields

$$\frac{\text{MAE}_{\text{CEV}}}{\text{MAE}_{\text{BS}}} = 1.83 - 0.080 \cdot \log_{10}(k), \quad R^2 = 0.091, \quad p = 0.006.$$

The negative slope is statistically significant: the two models diverge more as pool depth decreases, consistent with the theory that the CEV correction grows with $1/k$. Within the observed range ($\log_{10}(k) \approx 7$ to 12), the fitted ratio ranges from roughly 1.27 for the shallowest pools to 0.87 for the deepest. The regression crosses 1.0 at $\log_{10}(k) \approx 10.4$ ($k \approx 2.4 \times 10^{10}$), indicating that CEV hedging outperforms on average only for pools above this depth. For deep pools, CEV deltas achieve a modest hedging advantage, as expected from the more accurate local volatility specification. For shallow pools, the ratio exceeds 1 on average, indicating that CEV hedging is *worse* despite the larger theoretical correction. This likely reflects the heavier-tailed flows documented in Remark 6: when the diffusion assumption is most violated (small k, large σ_F/x), the CEV delta's additional dependence on k amplifies estimation error rather than improving the hedge. The low R^2 is expected, as many factors beyond pool depth (emission rates, staking patterns, subnet-specific events) affect hedging performance.

Of the 82 subnets, 38 (46%) show lower hedging error under CEV than BS. Among subnets with $\log_{10}(k) < 9.5$ (the shallower half), 19 of 37 (51%) favor CEV; among those with $\log_{10}(k) \geq 9.5$, only 19 of 45 (42%) do. The near-even split for shallow pools, combined with the above-1 mean ratio, suggests that CEV wins and losses are roughly balanced in count but the losses are larger in magnitude. Overall, both models produce modest hedging errors (median MAE $\approx 3\%$ of spot).

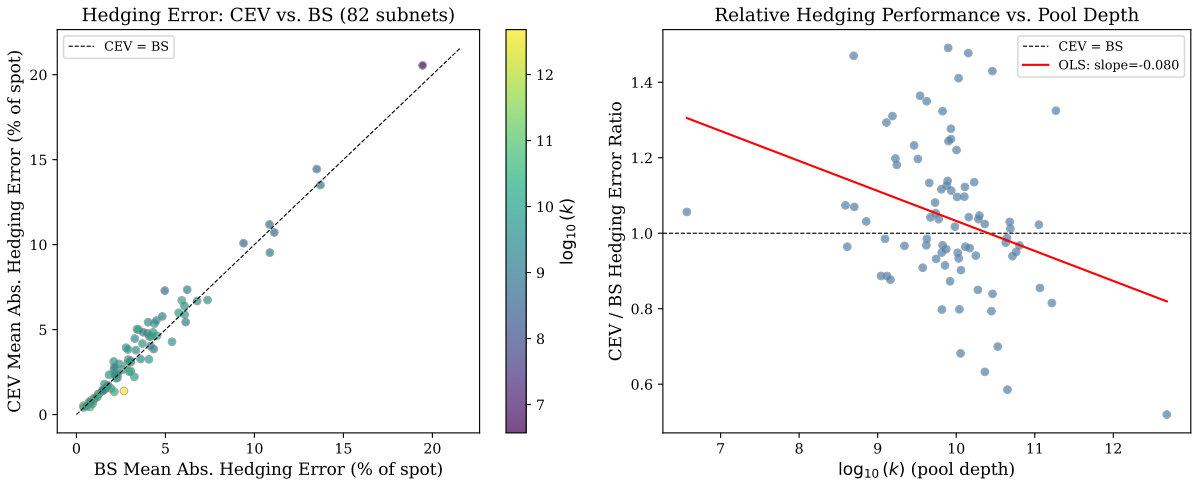

Figure 9: Cross-sectional delta-hedged backtest of 14-day ATM calls across 82 Bittensor subnets (August 2025 to February 2026). Left: each subnet's mean absolute hedging error under CEV (y-axis) vs. BS (x-axis), colored by pool depth $\log_{10}(k)$. Right: the CEV/BS error ratio declines with pool depth (OLS slope $= -0.080$, $p = 0.006$), confirming that the two models diverge more for shallower pools. Data source: Taostats API, daily pool snapshots.

7 Discussion

7.1 Limitations

Several limitations merit discussion. First, most AMMs charge a swap fee (e.g., 0.3% on Uniswap). Fees modify the effective invariant: a trade of Δx yields $\Delta y = y(1-\phi)\Delta x/(x+(1-\phi)\Delta x)$, where ϕ is the fee rate. This introduces a bid-ask spread but does not alter the qualitative CEV structure; the main effect is to reduce the effective σ_F by a factor of $(1-\phi)$ in the diffusion limit. Bittensor's dTAO pools currently charge no explicit swap fee, making our zero-fee model directly applicable.

Second, the diffusion assumption for staking flows is an approximation. In practice, large staking events ("whale" transactions) can produce jump-like price movements. A Merton-type jump-diffusion extension [Merton, 1976] would be straightforward: augment the flow process with a compound Poisson component, leading to a jump-diffusion CEV model. The option pricing formula would then involve a weighted sum of CEV prices across possible jump scenarios.

Third, the risk-neutral pricing argument requires the ability to delta-hedge, which is imperfect due to AMM slippage. Our replication premium (Proposition 9) provides a bound on this friction, but a more rigorous treatment would employ the utility-based framework of Davis et al. [1993] for markets with transaction costs.

Fourth, the model treats flow volatility σ_F as constant. In practice, staking activity exhibits time-of-day effects, momentum, and regime changes. A stochastic volatility

extension that layers Heston-type dynamics [Heston, 1993] onto the flow process would capture these features at the cost of analytical tractability.

Fifth, the model treats staking flows as exogenous, but in practice they may respond to the TAO/USD exchange rate. Since TAO trades on centralized exchanges, a decline in the TAO/USD rate may trigger unstaking (alpha-to-TAO swaps) as participants rebalance their USD-denominated portfolios. This feedback loop is absent from our model. For derivatives denominated in TAO (i.e., the numeraire of the AMM), the framework remains valid as stated; the CEV dynamics describe the alpha/TAO price conditional on a given flow process. For USD-denominated derivatives, one would need to jointly model the TAO/USD price and the flow process, likely introducing stochastic correlation.

Sixth, AMM token prices are vulnerable to manipulation near option expiry. Because the AMM is the sole price discovery mechanism, a well-capitalized agent can move the settlement price by staking or unstaking a large amount just before expiry. The cost of moving the price by a fraction ϵ from the current level is approximately $\epsilon \cdot x$ (the TAO reserve), which for shallow pools may be small relative to the option payoff gained. If the counterparty's ability to respond is constrained (e.g., by lock-up periods, gas costs, or information delays), this creates an asymmetric manipulation opportunity. Any practical implementation of options on AMM tokens must therefore incorporate safeguards such as time-weighted average prices (TWAPs) for settlement, minimum settlement windows, or oracle-based price feeds that aggregate over multiple blocks.

7.2 Extensions

Several extensions are promising. *Concentrated liquidity.* Uniswap V3's concentrated liquidity positions restrict the constant-product invariant to a bounded price range $[P_a, P_b]$. Within this range, the local reserve dynamics are equivalent to a constant-product AMM with effective invariant k_{eff}, so our CEV result applies locally. The continuous-time AMM framework of Cartea et al. [2024], which models informed and uninformed flow in Uniswap V3 pools, provides a natural starting point for this extension: their local dynamics within each tick range exhibit the constant-product structure to which our CEV characterization applies. At the range boundaries, the position becomes single-sided, creating a discontinuity in the price process that could be modeled as an absorbing or reflecting barrier. *Cross-subnet options.* Correlations between staking flows into different Bittensor subnets could be modeled by correlating the Brownian motions, enabling the pricing of basket options or spread options on multiple alpha tokens. *American and perpetual options.* American options can be priced via the free-boundary formulation of the CEV model [Detemple and Tian, 2002]. Perpetual options, recently proposed as a DeFi primitive [Dave, 2023], fit naturally into the CEV framework through the stationary solution of the pricing PDE.

7.3 Practical Implications

For practitioners, the key takeaway is that Black–Scholes *underprices* downside protection on shallow AMM pools. The leverage effect (**??**) elevates implied volatility for out-of-the-money puts: as the token price falls toward the strike, the AMM's bonding curve amplifies volatility, making further declines more likely than the lognormal model predicts. A market maker using Black–Scholes to sell puts on a low-liquidity alpha token would systematically charge too little, exposing themselves to losses in downward moves. Conversely, Black–Scholes overprices out-of-the-money calls, since the CEV model's declining volatility on the upside suppresses the right tail. The CEV model, or at minimum the liquidity correction (Proposition 8), should be used whenever the pool depth is small relative to the option notional.

8 Conclusion

We have shown that the price of a token traded on a constant-weighted-product automated market maker follows a constant elasticity of variance process, with the CEV exponent equal to the numeraire weight. This result is derived from first principles: the AMM's bonding curve mechanics and the stochastic nature of staking flows together determine the price dynamics, with no exogenous distributional assumptions required. The Black–Scholes model emerges as the limiting case of infinite pool depth, providing a precise characterization of when standard pricing tools are adequate and when they are not.

The framework yields closed-form European option prices via the non-central chi-squared distribution, novel liquidity and emission Greeks, and a quantitative decomposition of the pricing discrepancy relative to Black–Scholes as a function of pool depth and moneyness. The CEV structure also provides a first-principles derivation of the leverage effect for AMM tokens: the negative correlation between price and volatility arises directly from the bonding curve, making it a structural prediction rather than an empirical regularity.

Applied to 82 Bittensor subnets, a cross-sectional delta-hedged backtest confirms that the divergence between CEV and Black–Scholes hedging performance varies significantly with pool depth ($p = 0.006$). For deep pools, where the diffusion approximation holds well, CEV deltas achieve a modest hedging advantage. For shallow pools, both models hedge poorly due to heavy-tailed flows that violate the diffusion assumption, and CEV's additional parameter sensitivity amplifies estimation error. The low cross-sectional R^2 underscores that pool depth is one of several factors driving hedging performance.

The emergence of AMM-native tokens, not only in Bittensor but across a growing number of decentralized protocols, creates a genuine need for derivative pricing tools

tailored to these instruments. As options markets develop on AMM tokens (whether through on-chain protocols or centralized venues), the CEV framework developed here provides a foundation for pricing, hedging, and risk management that respects the structural constraints of the underlying market mechanism.

References

Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 core. *Uniswap Labs Technical Report*, 2021.

Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra. An analysis of Uniswap markets. *Cryptoeconomic Systems*, 1(1), 2021.

Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. Optimal routing for constant function market makers. *ACM DeFi*, 2022.

Stan Beckers. The constant elasticity of variance model and its implications for option pricing. *Journal of Finance*, 35(3):661–673, 1980.

Bittensor Foundation. Dynamic TAO whitepaper. https://bittensor.com/dtao-whitepaper, 2025. Accessed February 2026.

Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.

Block Scholes and Panoptic. Perpetual options — a research report. Block Scholes Research, 2025. Published August 2025.

Vitalik Buterin. On path independence. Blog post, 2017. https://vitalik.ca/general/2017/06/22/marketmakers.html.

Álvaro Cartea, Fayçal Drissi, and Marcello Monga. Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. *SIAM Journal on Financial Mathematics*, 15(3):931–961, 2024.

Joseph Clark. Replicating market makers. *arXiv preprint arXiv:2103.14769*, 2021.

John C. Cox. Notes on option pricing I: Constant elasticity of variance diffusions. *Working Paper, Stanford University*, 1975. Reprinted in Journal of Portfolio Management, 1996.

John C. Cox and Stephen A. Ross. The constant elasticity of variance option pricing model. *Journal of Portfolio Management*, 22:15–17, 1996. Special Issue.

Sachin Dave. Perpetual options in decentralized finance. *Panoptic Research Report*, 2023.

Mark H. A. Davis, Vassilios G. Panas, and Thaleia Zariphopoulou. European option pricing with transaction costs. *SIAM Journal on Control and Optimization*, 31(2): 470–493, 1993.

Dmitry Davydov and Vadim Linetsky. Pricing options on scalar diffusions: An eigenfunction expansion approach. *Operations Research*, 49(2):185–197, 2001.

Jérôme Detemple and Weidong Tian. The valuation of American options for a class of diffusion processes. *Management Science*, 48(7):917–937, 2002.

David C. Emanuel and James D. MacBeth. Further results on the constant elasticity of variance call option pricing model. *Journal of Financial and Quantitative Analysis*, 17 (4):533–554, 1982.

Florence Guillaume and Dennis Schroers. A unified approach for hedging impermanent loss of liquidity provision. *arXiv preprint arXiv:2407.05146*, 2024.

Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, 1993.

Sébastien Hitier. The dynamics of constant product market makers: A geometric Brownian motion approach. *SSRN Working Paper 5404433*, 2025.

Manuela Larguinho, José Carlos Dias, and Carlos A. Braumann. A note on the computation of the CEV option pricing formula. *Quantitative Finance*, 13(6):877–886, 2013.

Stefan Loesch, Nate Hindman, Mark B. Richardson, and Nicholas Welber. Impermanent loss in Uniswap v3. *arXiv preprint arXiv:2111.09192*, 2021.

Fernando Martinelli and Nikolai Mushegian. A non-custodial portfolio manager, liquidity provider, and price sensor. *Balancer Labs Technical Report*, 2019.

Robert C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1–2):125–144, 1976.

Andreas Park. The conceptual flaws of decentralized automated market making. *Management Science*, 69(11):6731–6751, 2023.

Tim Roughgarden. Transaction fee mechanism design in a post-MEV world. *ACM SIGecom Exchanges*, 21(1):2–18, 2024.

Mark Schroder. Computing the constant elasticity of variance option pricing formula. *Journal of Finance*, 44(1):211–219, 1989.

Fateh Singh. Option contracts in the DeFi ecosystem: Opportunities, solutions, and technical challenges. *International Journal of Network Management*, 35(2):e70005, 2025.

A Proofs

A.1 Proof of Proposition 8

Write the CEV call price (22) as $C(\beta)$ and note that $C(1) = C_{\mathrm{BS}}(\sigma_{\mathrm{eff}})$ when $\sigma_{\mathrm{eff}} = \delta P^{\beta-1}$ is held fixed. The liquidity correction is $\Lambda_C = C(\beta) - C(1) \approx C'(1)(1-\beta)$, so we need $\partial C/\partial\beta$ evaluated at $\beta = 1$.

The non-central chi-squared parameters a, b, c in (23)–(26) depend on β through the exponents $2(1-\beta)$ and $1/(1-\beta)$. As $\beta \to 1$, $b = 1/(1-\beta) - 1 \to \infty$, and the non-central chi-squared distribution converges to a normal distribution. Writing $\chi^2(x; n, \lambda)$ in terms of the normal CDF plus $O(1/\sqrt{n})$ corrections [Larguinho et al., 2013] and differentiating with respect to β at $\beta = 1$ gives

$$C'(1) = -\sigma_{\mathrm{eff}} P\sqrt{T}\,\phi(d_2)\left[\frac{\sigma_{\mathrm{eff}}\sqrt{T}}{2} + \frac{d_1}{2}\right],$$

where d_1, d_2 are the standard Black–Scholes quantities. At the money ($P = K_{\mathrm{str}}$), $d_1 = \sigma_{\mathrm{eff}}\sqrt{T}/2$ and $d_2 = -\sigma_{\mathrm{eff}}\sqrt{T}/2$, so $\phi(d_2) = \phi(0)\cdot\exp(-\sigma_{\mathrm{eff}}^2 T/8) \approx \phi(0)$ for small $\sigma_{\mathrm{eff}}^2 T$, yielding $\Lambda_C \approx -\frac{1}{2}(1-\beta)\sigma_{\mathrm{eff}}^2 PT\phi(0)$, which for $\beta = 1/2$ gives $\Lambda_C \approx -\frac{1}{4}\sigma_{\mathrm{eff}}^2 PT\phi(0)$.

A.2 Proof of Proposition 9

The hedging error over $[t, t+\Delta t]$ from trading $\Delta_{\mathrm{CEV}}\cdot\Delta P$ units through the AMM with price impact (36) is

$$\epsilon_t = S\big(\Gamma_{\mathrm{CEV}}(\Delta P)^2/(2P)\big) \approx \frac{P^2\Gamma_{\mathrm{CEV}}^2}{2k}\cdot\delta^2 P^{2\beta}\Delta t,$$

using $(\Delta P)^2 \approx \delta^2 P^{2\beta}\Delta t$. Integrating and taking expectations under $\mathbb{Q}$ gives (38).

A.3 Proof of Proposition 10

For time-dependent $\delta(t)$, the CEV transition density depends on the total integrated variance $\bar{v}^2 = \int_0^T \delta(t)^2 \mathrm{d}t$ when the time-change technique is applied. Substituting $\delta(t) = 2\sigma_F/\sqrt{k_0 + \dot{k}t}$ and evaluating the integral yields (42).