

Modeling Loss-Versus-Rebalancing in Automated Market Makers via Continuous-Installment Options

Srisht Fateh Singh ✉ 

University of Toronto, Canada

Samuel Gaskin ✉ 


Independent Researcher, Toronto, Canada

Jeffrey Klinck ✉ 

University of Toronto, Canada

Zissis Poulos ✉ 

York University, Toronto, Canada

Reina Ke Xin Li ✉ 

University of Toronto, Canada

Yuntao Wu ✉ 

University of Toronto, Canada

Panagiotis Michalopoulos ✉ 

University of Toronto, Canada

Andreas Veneris ✉ 

University of Toronto, Canada

Abstract

This paper mathematically models a constant-function automated market maker (CFAMM) position as a portfolio of exotic options, known as perpetual American continuous-installment (CI) options. This model replicates an AMM position's delta at each point in time over an infinite time horizon, thus taking into account the perpetual nature and optionality to withdraw of liquidity provision. This framework yields two key theoretical results: (a) It proves that the AMM's adverse-selection cost, loss-versus-rebalancing (LVR), is analytically identical to the continuous funding fees (the time value decay or theta) earned by the at-the-money CI option embedded in the replicating portfolio. (b) A special case of this model derives an AMM liquidity position's delta profile and boundaries that suffer approximately constant LVR, up to a bounded residual error, over an arbitrarily long forward window. Finally, the paper describes how the constant volatility parameter required by the perpetual option can be calibrated from the term structure of implied volatilities and estimates the errors for both implied volatility calibration and LVR residual error. Thus, this work provides a practical framework enabling liquidity providers to choose an AMM liquidity profile and price boundaries for an arbitrarily long, forward-looking time window where they can expect an approximately constant, price-independent LVR. The results establish a rigorous option-theoretic interpretation of AMMs and their LVR, and provide actionable guidance for liquidity providers in estimating future adverse-selection costs and optimizing position parameters.

2012 ACM Subject Classification Applied computing → Economics

Keywords and phrases blockchain, decentralized finance, automated market makers, mathematical finance, perpetual options, continuous installments

Digital Object Identifier 10.4230/LIPICs.AFT.2025.6

Related Version *Full Version:* <https://arxiv.org/abs/2508.02971> [17]

Acknowledgements The authors would like to thank the reviewers for their helpful comments.

1 Introduction

The success of blockchains supporting smart contract such as Ethereum [4], Solana [18], etc., has led to the rise of *Decentralized Finance* (DeFi) which offers alternatives to traditional financial services by removing central trusted intermediaries and replacing them with public, verifiable, and immutable computer programs. One of the pivotal components of the DeFi infrastructure stack is automated market makers, or AMMs, allowing the exchange of one token for another at prices decided by an underlying algorithm. In recent years, AMMs have seen a rapid adoption reflected in financial metrics such as total value locked (above



© Srisht Fateh Singh, Reina Ke Xin Li, Samuel Gaskin, Yuntao Wu, Jeffrey Klinck, Panagiotis Michalopoulos, Zissis Poulos, and Andreas Veneris;
licensed under Creative Commons License CC-BY 4.0

7th Conference on Advances in Financial Technologies (AFT 2025).

Editors: Zeta Avarikioti and Nicolas Christin; Article No. 6; pp. 6:1–6:23



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

\$21B) and yearly transaction volumes (above \$2T), as well as by their composability [8, 13]. Today, tens of thousands of tokens are listed and hundreds of applications are built on top of them [10].

The key participants in an AMM are traders and agents known as *liquidity providers*, or LPs. Traders exchange one token for another, where the token pair generally consists of a risky asset with volatile value, and a stable asset or numéraire.¹ On the other hand, LPs serve as counterparties to traders (sellers to buyers and buyers to sellers) by depositing both tokens upfront to the exchange. As a result of each trade, LPs receive the less favourable of the two tokens. To hedge against the adverse selection faced by AMMs, LPs can continuously rebalance an off-chain replicating portfolio by accumulating the risky token as its price rises and selling it when the price falls. However, this hedge is not perfect when the AMM is not the primary venue for price discovery because the pool's quoted price tends to lag behind the prevailing price on a primary venue such as a centralized exchange.

Under such a setting, even proactive LPs who rebalance in response to price changes are exposed to a systematic cost. Since AMM quotes lag those on primary markets, arbitrageurs can act faster than LPs and restore the pool price to the external market level. This generates a small but persistent transfer of value from LPs to arbitrageurs, and when aggregated over multiple price updates, this cost becomes significant. This cost is referred to as *loss-versus-rebalancing* (LVR) [16]. The rate at which LVR accumulates depends on the steepness of the AMM curve and the volatility of the underlying token; both amplify the arbitrage gap and thus accelerate LVR.

Despite such costs, LPs are incentivized to participate through the earning of trading fees that are proportional to the value of each trade and paid by the trader. Therefore, LPs considering whether to provide liquidity on an AMM pool must calculate their expected *payoff* by estimating and comparing their position's LVR with the anticipated trading fees in some predetermined forward time window. Recent work [16] have analyzed and quantified expressions for instantaneous LVR and retrospectively tested with historic market data. However, there is limited work that provides estimation methods for future LVR.

This work estimates the LVR for an LP that decides to provide liquidity for an arbitrarily long period and can exit at any point. It does so by mathematically modeling liquidity provision on a general class of AMMs, known as constant function AMMs (CFAMMs), as selling a continuum of perpetual American put options across continuous strikes. Perpetual American options are financial derivatives that give their holder the right to buy (known as a call) or sell (known as a put) an underlying asset at an agreed-upon price (strike) with no expiration. Unlike traditional vanilla options, this work uses exotic options in its model, known as continuous installment (CI) options, in which the holder must pay a stream of constant installment rate, referred to as *funding fees*, to keep their position alive. This funding fee is analogous to the time value decay of traditional fixed-term options. The paper uses perpetual American CI options because, unlike fixed-term vanilla options, the pricing function of these options does not change over time (assuming other market parameters are constant). Moreover, despite their exotic nature, CI options have been well-studied in the past, and this work builds on results from the existing literature [5]. This approach yields two key theoretical results.

¹ Although some token pairs consist of two stable assets, this work focuses primarily on pairs of one risky and one stable asset.

Funding Fees = LVR. In the limit where the installment rate tends to infinity (analogous to extremely short-dated fixed-term options), a CI put option has the classic hockey-stick payoff function: it pays the difference between the strike and the spot price when the underlying's spot price is below the strike price, and zero otherwise. Therefore, in this limit, a continuous distribution of CI puts *exists* that delta-replicates an arbitrary LP position's payoff *at each point in time*. Moreover, the installment rate earned on this distribution of puts (which, at a given point in time, is earned by the option whose strike equals the spot price at that time) reproduces the expression for the instantaneous LVR. Therefore, theoretically, an LP can hold such an options portfolio to stay delta-neutral – the cost due to time value decay of holding this portfolio, *i.e.*, funding fees, is the LVR.

Since the above model is theoretical, as a continuum of options cannot be reproduced in practice, the paper subsequently quantifies the approximation error when a discrete portfolio of options with discrete strike prices is used to replicate the payoff of a liquidity position.

Constant Future LVR. The second result analyzes the converse scenario where a liquidity position's payoff replicates the valuation of a single perpetual American CI option. This produces a unique liquidity profile with *almost constant* instantaneous LVR rate over a forward time window. Moreover, this rate is approximately equal to the funding fee of the replicated CI option. As a result, this yields guidance to LPs, under the model assumptions, on:

- Choosing the price boundaries and shape of liquidity provision that incurs *predictable, flat, price-path independent* future LVR.
- Estimating the forward adverse-selection cost for a planned holding period.
- Understanding the relationship between the optimal holding period and the width of the liquidity position.
- Selecting an appropriate pool based on its expected future trading fee income.

The paper is organised as follows. Section 2 explains notations and the necessary background, Section 3 discusses prior literature and related works, Section 4 motivates the option-based interpretation, Section 5 provides the options decomposition, Section 6 proves the funding-fee–LVR identities, Section 7 measures the approximation error on delta when the continuous strip is replaced by finitely many strikes, and Section 8 presents the volatility calibration and design rules for LPs. Finally, Section 9 concludes with directions for future research.

2 Background

In this section, we provide the necessary notation, terminology and background concepts used in the remainder of the paper.

2.1 Notation

We consider two tokens: *token 0* representing a risky asset (*e.g.* BTC, ETH) and *token 1* representing a stable/safe asset (*e.g.* USDC). Let S_t denote the spot price of *token 0* in units of *token 1* at time t that follows a geometric Brownian motion (GBM) on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{Q})$ satisfying the standard assumptions for GBMs (where \mathbb{Q} is a risk-neutral probability measure), so that

$$\frac{dS_t}{S_t} = r dt + \sigma dB_t^{\mathbb{Q}}, \quad (1)$$

with constant annual risk-free rate r and volatility $\sigma > 0$. $\{B_t^{\mathbb{Q}}\}_{t \geq 0}$ is a Wiener process. As usual, we assume that AMMs constitute secondary markets and the price S_t is governed by primary markets such as centralized exchanges. In the subsequent sections, we omit the subscript t and use S and S_t interchangeably for convenience.

2.2 Constant-Function Automated Market Makers (CFAMMs)

A *constant-function automated market maker* (CFAMM) maintains token reserves (x, y) , deposited by LPs, such that each trade transforms the reserves to (x', y') and the reserves before and after the trade satisfy an invariant function $F(x, y) = F(x', y') = k$. In a *constant-product* AMM (CPAMM), such as Uniswap v2 [1], the invariant takes the form of $F(x, y) = \sqrt{xy}$. As a result, the marginal exchange price, assuming no arbitrage, takes the form: $S = -\frac{\partial y}{\partial x} = \frac{y}{x}$. This is shown in Figure 1a. Therefore, the expression for token reserves are $x = \frac{k}{\sqrt{S}}$, and $y = k\sqrt{S}$.

A CPAMM with concentrated liquidity (as in Uniswap v3 [2]) uses an invariant parameter k within a price band $[a, b]$. The token reserves of LPs in this band consist of only *token 0* when the price $S \geq b$, only *token 1* when $S \leq a$, and for prices $S \in (a, b)$, the reserves are:

$$x = k \frac{\sqrt{b} - \sqrt{S}}{\sqrt{Sb}}, \quad y = k (\sqrt{S} - \sqrt{a}). \quad (2)$$

Therefore, the reserve value (denominated in token 1) is

$$V(S) = k (\sqrt{S} - \sqrt{a}) + kS \frac{\sqrt{b} - \sqrt{S}}{\sqrt{Sb}}, \quad S \in [a, b]. \quad (3)$$

The liquidity position's sensitivity to price, known as its *delta*, is denoted by $X(S)$, and the sensitivity of delta to price, known as *gamma*, is denoted by $\Gamma(S)$. In practice, delta is a measure of exposure to small changes in the price of the risky asset, whereas gamma is a measure of exposure to large movements of the risky asset's price. These are given by the first and second derivatives of the value function $V(S)$ with respect to the underlying's price, respectively, and are expressed as follows:

$$X(S) := V'(S), \quad \Gamma(S) := V''(S) = X'(S) \leq 0 \quad (S \in (a, b)). \quad (4)$$

Figure 1b illustrates the behavior of delta and gamma across the liquidity band. As shown, the magnitudes of both delta and gamma decrease monotonically with price, as higher prices correspond to a greater allocation to the numeraire asset.

2.2.1 Loss-Versus-Rebalancing (LVR)


A continuously rebalanced, self-financing delta-hedge portfolio that holds $X(S_t)$ units of token 0 has value W_t with $dW_t = X(S_t) dS_t$. The difference

$$\text{LVR}_t := V(S_t) - W_t \quad (5)$$

quantifies the AMM's adverse-selection loss relative to a hedged trader and is referred to as loss-versus-rebalancing or LVR. The instantaneous LVR, $d\text{LVR}_t$, grows quadratically with the spot price and volatility, and linearly with the gamma of the liquidity position [16].

$$d\text{LVR}_t = \frac{1}{2} \sigma^2 S_t^2 \Gamma(S_t) dt = \frac{1}{2} \sigma^2 S_t^2 [X'(S_t)] dt. \quad (6)$$

We will later demonstrate the equivalence between the right-hand side of Eq. (6) and the funding fee of a perpetual American CI option.



(a) CPAMM Invariant with $k = 1$ and Marginal Price (Tangent at $x = 30$).

(b) $V(S)$, $100X(S)$, and $-5000\Gamma(S)$ for concentrated liquidity position in price band $[50, 200]$, and $k = 1$.

■ **Figure 1** Profiling CPAMM invariant and concentrated liquidity position, its delta and gamma.

2.3 Perpetual American Continuous-Installment Options

A *perpetual American continuous-installment put option* has no expiration date and requires the holder to pay a continuous stream of *constant funding fee* $q > rK$ per year to keep the contract alive. At any point, the holder may choose to stop paying the fee, at which point they can either exercise the option or drop the position. The option can be exercised at any time for a payoff of $\max(K - S, 0)$ [5]. In the analysis below, we assume the underlying asset (token 0) pays zero dividends.

2.3.1 Notation and Ordinary Differential Equation Formulation

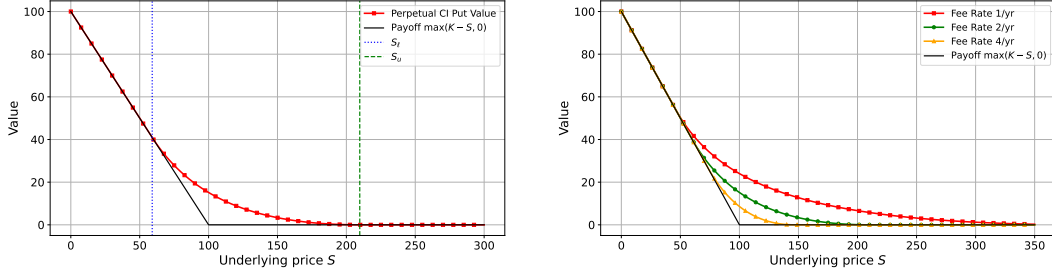
Let $P_q(S; K)$ denote the discounted put option value at spot price S , strike K , and funding fee q . Let S_ℓ denote the lower boundary, below which the option value equals its payoff, and let S_u denote the upper boundary, above which the option value is zero, as illustrated in Figure 2a. Under the risk-neutral dynamics, $P_q(S; K)$ satisfies the inhomogeneous Black-Scholes ordinary differential equation in the *continuation region* $S_\ell < S < S_u$:

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^2 P_q}{\partial S^2} + rS \frac{\partial P_q}{\partial S} - rP_q = q, \quad S \in (S_\ell, S_u). \quad (7)$$

The left and right boundaries are determined endogenously from the value-matching and delta-matching conditions

$$\begin{aligned} P_q(S_\ell; K) &= K - S_\ell, & \frac{\partial P_q}{\partial S}(S_\ell; K) &= -1, \\ P_q(S_u; K) &= 0, & \frac{\partial P_q}{\partial S}(S_u; K) &= 0. \end{aligned} \quad (8)$$

Solving (7) with the four boundary conditions in (8) yields the closed-form expressions below.



(a) Discounted value of a perpetual CI put option with lower and upper boundaries.

(b) Discounted value of a perpetual CI put option with fee rates 1, 2, 4 per annum.

■ **Figure 2** Profiling of perpetual CI put option with $K = 100.0$, $\sigma = 0.25$, $q = 2$ (in the first figure) and $r = 0.01$ (all parameters are annualized).

2.3.2 Closed-form solution

The expression for the option price $P_q(S; K)$ takes the following closed form as derived in [5].

$$P_q(S; K) = \alpha_p S + \beta_p S^{\gamma_p} + \frac{q}{r} \quad (9)$$

The delta of the put option, $X_q(S; K) = \frac{\partial}{\partial S} P_q(S; K)$, thus takes the form:

$$X_q(S; K) = \alpha_p + \beta_p \gamma_p S^{\gamma_p - 1}. \quad (10)$$

Moreover, the upper and lower boundaries, S_u and S_ℓ respectively, have the following closed form:

$$S_\ell = \frac{q}{r + \sigma^2/2} [g - g^{1/\gamma_p}], \quad (11)$$

$$S_u = \frac{q}{r + \sigma^2/2} [g^{1-1/\gamma_p} - 1]. \quad (12)$$

Here, α_p , β_p , γ_p , and g are expressions that depend on parameters r , σ , K , and q and their expressions are provided in Appendix (A).

The Black–Scholes partial differential equation in (7) contains no derivative with respect to time. Consequently, the value of a perpetual CI option is *time-invariant*. This makes it unique from vanilla finite expiry American or European options, whose pricing has a time-varying component. The analog of expiration in CI options is the funding rate, where high funding rates make it “behave similar” to short expiration option and vice versa for small fee rates. This is depicted in Figure 2b, where increasing the fee rate reduces the price of the CI option closer to its payoff. The lower and upper boundaries, S_ℓ and S_u , characterize the holder’s optimal policy. When the spot price first falls below S_ℓ , it is optimal to *exercise* the option; when it first exceeds S_u , it is optimal to *drop* the option – *i.e.*, to exit the position with zero payoff. This is because, in both cases, the expected benefit of continued funding is outweighed by its cost. In either scenario, the holder stops paying the funding fee immediately upon exit. For this reason, S_ℓ and S_u are also called *optimal exercise* and *dropping boundaries*, respectively. Conversely, the option seller receives the continuous funding fee only while the spot price remains in the continuation region $S_\ell < S < S_u$. We will exploit this fact in the sections that follow.

3 Related Work

Early work on studying LP positions in CFAMMs focus on mitigating risks associated with impermanent loss – the loss experienced by a liquidity provider compared to simply holding the asset, so that an LP can earn trading fees without exposure to this loss. Deng et al. [9] and Fukasawa et al. [12] study static replication strategies for impermanent loss in finite-time CFAMM positions using European options and variance swaps, respectively. Lipton et al. [14] further studies model-based dynamic replication. However, it can be argued that a comparison against a buy-and-hold strategy is insufficient as it does not account for adverse selection, where informed traders extract value from passive LPs over time.

Another approach to quantifying LP loss is the concept of LVR, which captures the loss incurred by a liquidity provider compared to continuously rebalancing their portfolio at market prices, due to adverse selection. LVR was formalized by Millionis et al. [16]. Their work offers closed-form expressions for instantaneous LVR in a CFAMM and provides empirical validation using historical market data. However, they focus on instantaneous and historical LVR and do not address forward-looking or long-term LVR estimation. Meanwhile, our approach provides a framework that captures forward-looking LVR by delta-replication with options portfolios.

Maire and Wunsch [15] make a case for the hedging of the LP position value instead of the impermanent loss. They study the problem of market-neutral liquidity provision by constructing a static replication portfolio that matches the AMM position's dollar value over time, effectively achieving a constant value position for the finite lifetime of the position. The replicating portfolio's margin requirement is then itself hedged by shorting a perpetual or dated futures contract to offset changes in the margin value, enabling a market-neutral LP strategy that generates interest from LP trading fees and futures funding fees. Meanwhile, Clark [6, 7] investigates the replicating portfolio of the payoff of a constant product AMM position and shows how an LP position's terminal value can be fully statically replicated using a portfolio of European options. The approach focuses on fixed, finite time horizons and seeks to hedge only the terminal value of the liquidity position. On the other hand, our approach neither relies on dynamic hedging, nor assumes finite time horizons. Instead, we model the LP's position over an indefinite time horizon using a portfolio of perpetual American CI put options, providing a theoretical framework that statically captures the path-dependency of forward-looking LVR.

4 Work Motivation

The value profile of a concentrated liquidity position in a finite price band, as shown in Figure 1b, closely resembles that of a portfolio consisting of cash (a constant payoff) and short put options (the negative of the payoff shown in Figure 2a): flat on one wing, linear on the other, and smoothly curved in between. Figure 3a makes this visual similarity precise by comparing the *value* of a CPAMM position with $k = 1$ and band $(80, 125)$ to a portfolio comprising cash and a one-month European put, with strike at the geometric mean of the price boundaries. The put is valued using the standard Black–Scholes model [3].

Despite the superficial similarity, key differences emerge. As shown in Figure 3b, the two profiles diverge meaningfully. More critically, the European option's value is inherently time-variant: even a single day's passage erodes its time value (*theta*), while the CPAMM's value remains time-stationary. This contrast is illustrated in Figure 3c, where the solid line represents the option value as a function of the time-to-maturity. In addition, the CPAMM offers a flexible, perpetual holding period, whereas fixed-term options require periodic rolling – selling expiring contracts and buying new ones – to replicate a liquidity position.

These challenges raise a natural question: *Can one construct a static portfolio that tracks a perpetual AMM band without daily rebalancing?* The answer is affirmative, provided we replace European options with a class of *American continuous-installment options*. This is because the ongoing constant funding fee, $q \, dt$, charged by an active CI option offsets the time decay found in European options. In the limit of infinite maturity (the perpetual CI variant), the mark-to-market value of a CI put or call becomes *time-invariant*, producing a flat line as shown by the dashed line in Figure 3c.

This observation enables a decomposition of a CFAMM band into a *perpetual strip* of CI puts that both matches the pool's delta and offers stationarity. Beyond its conceptual appeal, this decomposition yields two practical insights:

1. The instantaneous *funding fee* of the active CI option equals the LP's *loss-versus-rebalancing* cost. Since a CI option is perfectly hedgeable with a Black-Scholes type rebalancing portfolio [5] and both the CPAMM and the CI option strip have the same delta, they share the same rebalancing portfolio, and hence the LVR is precisely the difference between the two instruments, *i.e.* instantaneous funding fee.
2. If the CFAMM's delta and price boundaries are calibrated to match those of a single CI put, then the LVR over a future time window becomes nearly constant and equal to the funding fee q . Notably, this construction relies on implied volatility, rather than instantaneous volatility, meaning it can be formulated using observed market data.

The following section formalizes the CI decomposition, proving the equivalence between funding fees and LVR.

5 Modeling a CFAMM Position with CI Options

5.1 Overview

In this section, we construct a portfolio whose delta (change in option price w.r.t. change in spot price) is the same as the delta of a CFAMM $X(S)$. This portfolio consists of a distribution of perpetual American CI put options in the limit $q \rightarrow \infty$ across a continuum of strike prices. As the funding rate tends to infinity, the exercise and dropping boundaries of each option collapse to the strike price, and the option's valuation converges to $\max(K - S, 0)$. Consequently, the option's delta converges to a step function. This property enables the construction of a portfolio that replicates the delta of an arbitrary (but smooth) CFAMM payoff function.


5.2 Portfolio Construction

We begin with the following lemma:

► **Lemma 1.** *In the limit $q \rightarrow \infty$, both the lower and upper boundary of a CI put converge to the strike K , and option's delta $X_q(S; K)$ becomes a step function $-\mathbf{1}_{\{S < K\}}$.*


This lemma encapsulates the relationship between the funding rate and the option's "effective" time to expiration. As the funding rate increases, the option behaves increasingly like a zero-time-to-expiry option, with its delta approaching a discontinuous step as $q \rightarrow \infty$.

Next, we specify the weight distribution of the options portfolio that delta-replicates a CFAMM value function $V(S)$.



(a) Value of a CPAMM band [80, 125] versus a short one-month Black-Scholes put plus cash.

(b) Difference between CPAMM value and short put plus cash.



(c) Time value decay of a European put versus the constant value of a perpetual CI call, whose funding fee offsets theta exactly.

■ **Figure 3** Comparison between CPAMM liquidity value and a portfolio of short put ($K=100$, $r=5\%$, $\sigma=50\%$) plus cash.

► **Theorem 2.** Let $V : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ be a twice continuously differentiable function. Assume that $V' \in L^1(\mathbb{R}_{>0})$, V' has bounded variation on $\mathbb{R}_{>0}$ and that $\lim_{S \rightarrow \infty} V'(S) = 0$. Define the weight $w(K) := V''(K)$ and, for each $q < \infty$,

$$\Pi_q(S) := \int_0^\infty w(K) P_q(S; K) dK,$$

and

$$\Pi(S) := \lim_{q \rightarrow \infty} \Pi_q(S).$$

Then the portfolio $\Pi(S) - V(S)$ is delta-neutral.

Thus, a perpetual CFAMM position that can be closed by its owner at any time can be perfectly modeled using a distribution of CI puts with very large funding rates². In practice, a continuous distribution is infeasible, and funding rates are finite. Therefore, one may replicate the delta profile using a discrete set of puts with different strikes. However, discretization and

² $\Pi_q(S_t)$ has the same payoff as described above at all times. Therefore, the holder must continuously reissue options that are exercised or dropped.

finite funding rates introduce non-negligible delta-replication error, dependent on inter-strike spacing and funding rate. We quantify this approximation error numerically for a constant product AMM position in Section 7.

An additional corollary of the above result is the relationship between the instantaneous LVR of a liquidity position and the funding fees of options with strikes around the spot price (also referred to as activated strikes) in the replicating portfolio. This relationship is analyzed in the following section.

6 Establishing LVR as Funding Fees

The portfolio Π delta-replicates a given CFAMM liquidity position. However, unlike the AMM position, the short options portfolio pays a continuous funding fee to the seller – arising from the active CI puts with strike around the spot price. We show that these funding fees, absent in the AMM, are precisely equal to the LP's LVR.

To compute this running funding fee, we first establish the following lemma:

► **Lemma 3.** *As $q \rightarrow \infty$, the product $q(S_u(q; K) - S_\ell(q; K))$ converges to a finite limit:*

$$\lim_{q \rightarrow \infty} q \cdot (S_u(q; K) - S_\ell(q; K)) = \frac{\sigma^2 K^2}{2}.$$

Next, we express Π as the limit of a discrete sum of options. In the following lemma, we construct such a discretization and show that, as $q \rightarrow \infty$, this portfolio converges pointwise to that of Π . We also compute the limiting funding fee contribution from the option whose holding region contains the spot price.

► **Lemma 4.** *Fix $q > 0$ and an interval $[a, b] \subset \mathbb{R}_{>0}$. Define a sequence of strikes $K_1 < K_2 < \dots < K_{N(q)}$ recursively by:*

$$\begin{cases} S_\ell(q; K_1) = a, \\ S_\ell(q; K_{i+1}) = S_u(q; K_i), \quad i = 1, \dots, N(q) - 1, \\ S_u(q; K_{N(q)}) = b. \end{cases}$$

For each i , define the weight

$$w_i := \int_{S_\ell(q; K_i)}^{S_u(q; K_i)} w(K) dK = X(S_u(q; K_i)) - X(S_\ell(q; K_i)),$$

so that

$$\sum_{i=1}^{N(q)} w_i = \int_a^b w(K) dK = -1.$$

Let K_j be the activated strike such that $S_\ell(q; K_j) \leq S_t \leq S_u(q; K_j)$. Then, as $q \rightarrow \infty$, the discrete portfolio

$$\tilde{\Pi}_q := \sum_{i=1}^{N(q)} w_i P_q(\cdot; K_i)$$

converges pointwise to the continuous payoff $\Pi(S)$. Moreover, the weighted funding fees at strike K_j satisfies

$$\lim_{q \rightarrow \infty} w_j q = \frac{\sigma^2 S^2}{2} X'(S).$$

Lemma 4 constructs a portfolio of CI puts with discrete strikes and finite funding rates, with weights chosen such that, in the limit $q \rightarrow \infty$, the portfolio converges to the CFAMM-replicating portfolio defined in Theorem (2). The index j denotes the unique option that remains active (*i.e.*, in the holding region), while all other options are either exercised or dropped. Therefore, the funding fee received by the portfolio owner is $w_j q$.

As q becomes large, Lemmas 3 and 4 together imply that this funding fee converges to the instantaneous LVR of the corresponding CFAMM position. This leads to the following result:

► **Theorem 5** (Funding fee = LVR). *Let $d\text{LVR}_t$ denote the instantaneous change in the LVR of a CFAMM position, and let $d\text{Fee}_t$ denote the instantaneous funding income of its delta-replicating CI option portfolio Π . Then,*

$$d\text{Fee}_t = d\text{LVR}_t \quad \text{and} \quad \text{Fee}|_0^T = \text{LVR}|_0^T \quad (\forall T > 0).$$

Therefore, the LVR of a liquidity position is precisely the CI funding premium of its delta-replicating options portfolio. Another way to look at this is that a CI option can be perfectly delta-hedged using a continuously rebalanced Black–Scholes-type portfolio of risky and stable assets [5]. As the funding rate q increases, the delta of this rebalancing portfolio converges to that of a CFAMM position. However, unlike the CI option, a CFAMM position does not compensate the liquidity provider via a funding stream. The discrepancy between the CFAMM position and its hedge is therefore exactly the foregone CI funding. This equality relies solely on closed-form expressions and the principle of self-financing, without invoking any additional modeling assumptions.

One advantage of this option-theoretic interpretation is that CI fee rates, as implied by the options market, provide both real-time and forward-looking estimates of LVR, enabling informed range management. Moreover, this framework permits the construction of liquidity profiles with nearly price-path-independent LVR. These features are demonstrated in the sections that follow. A further implication is that a market for CI options allows for static-weight delta-hedging of CFAMM positions – under the assumption of constant implied volatility – eliminating the need for frequent rebalancing, unlike with conventional American or European options.

6.1 CFAMM Position Replicating a Unit CI Option

Consider a concentrated-liquidity CFAMM band whose delta, at every price level, matches the delta of a *single* perpetual American CI *put* option. Specifically, choose the liquidity bounds $a < b$ such that

$$X(S) = V'(S) \equiv X_q(S; K_*) \quad (S \in [a, b]), \quad (13)$$

$$a = S_l(q, K_*), \quad (14)$$

$$b = S_u(q, K_*). \quad (15)$$

for some strike K_* and finite fee rate q .

► **Theorem 6.** *The instantaneous rate of change of the LVR of the above CFAMM position is approximately equal to the funding fees of the unit put option, up to a bounded approximation error. That is, there exists a residual function $\epsilon(t)$ with $|\epsilon(t)| \leq rK_*$, such that*

$$d\text{LVR}_t = q \, dt + \epsilon(t)dt,$$

Thus, the AMM liquidity position with the above delta profile suffers an *almost flat, price-path-independent, volatility-independent* LVR. Note that because $X_q(S; K_*)$ depends on the volatility parameter σ , as shown in Eq (10), the calibrated boundaries a, b – and thus the entire delta curve $X(S)$ – remain implicitly volatility-dependent. The residual error term, $\epsilon(\cdot)$, is bounded in magnitude by a constant and its relative magnitude is reported and discussed in Section 8. The above liquidity profile is useful for LPs who want to estimate forward LVR and compare it with the *expected future trading fees*. Lastly, constructing such an AMM profile requires estimates of future volatility. This can be approximated using a term structure of implied volatility gathered from the fixed-term options market. Section 8 discusses this in detail and estimates the approximation error arising from the calibration between fixed-term and perpetual options' volatilities.

7 Error Analysis of Discrete CI-Option Replication

In this section we quantify the approximation error that arises when the continuous-strike decomposition of a concentrated CFAMM is replaced by a *discrete* strip of perpetual American CI put options with *finite* q .

7.1 Sources of error

We isolate two drivers of error: (i) The installment rate q (large but finite), and (ii) the inter-strike spacing ΔK of the discrete strip. For a given pair $(q, \Delta K)$, we construct a strip of discrete options and measure the absolute difference between the target delta (of the CFAMM) and the strip's delta. This is done over the active price band $S \in [a, b]$ and its maximum and root-mean-square values are plotted.

7.2 Experimental methodology

We chose a concentrated liquidity AMM as our target CFAMM. Thus, the analytical delta $X(S) = L(1/\sqrt{S} - 1/\sqrt{b})$ for uniform liquidity on a price band $[a, b]$. We choose $a = 80$, and $b = 125$, and L chosen so that $X(a) = 1$, $X(b) = 0$. We discretize the replication weight such that on each strike interval $[K_i, K_{i+1}]$ we set

$$w_i = \int_{K_i}^{K_{i+1}} X'(K) dK = X(K_{i+1}) - X(K_i), \quad (16)$$

ensuring that the discrete weights sum to the continuous integral. For each strike, we compute the short CI put delta, clipped to $\{-1, 0\}$ outside its continuation band. Let $X_{\text{strip}}(S; q, \Delta K) = \sum_i w_i X_q(S; K_i)$ be the delta of the strip at S . Define the error

$$\varepsilon(S_j; q, \Delta K) = |X(S_j) - X_{\text{strip}}(S_j; q, \Delta K)|$$

evaluated on grid $\{S_j\}_{j=1}^N$ where $N = 2000$. We consider two error metrics:

- *Maximum absolute error*: $\max_j \varepsilon(S_j; q, \Delta K)$.
- *Root-mean-square error (RMSE)*: $\sqrt{N^{-1} \sum_j \varepsilon(S_j; q, \Delta K)^2}$.

We evaluate a parameter sweep $(q, \Delta K) \in \{8, 16, 32, 64, 125, 250, 500, 1000, 2000, 4000\} \times \{0.25, 0.5, 1.0, 2.0, 4.0\}$.

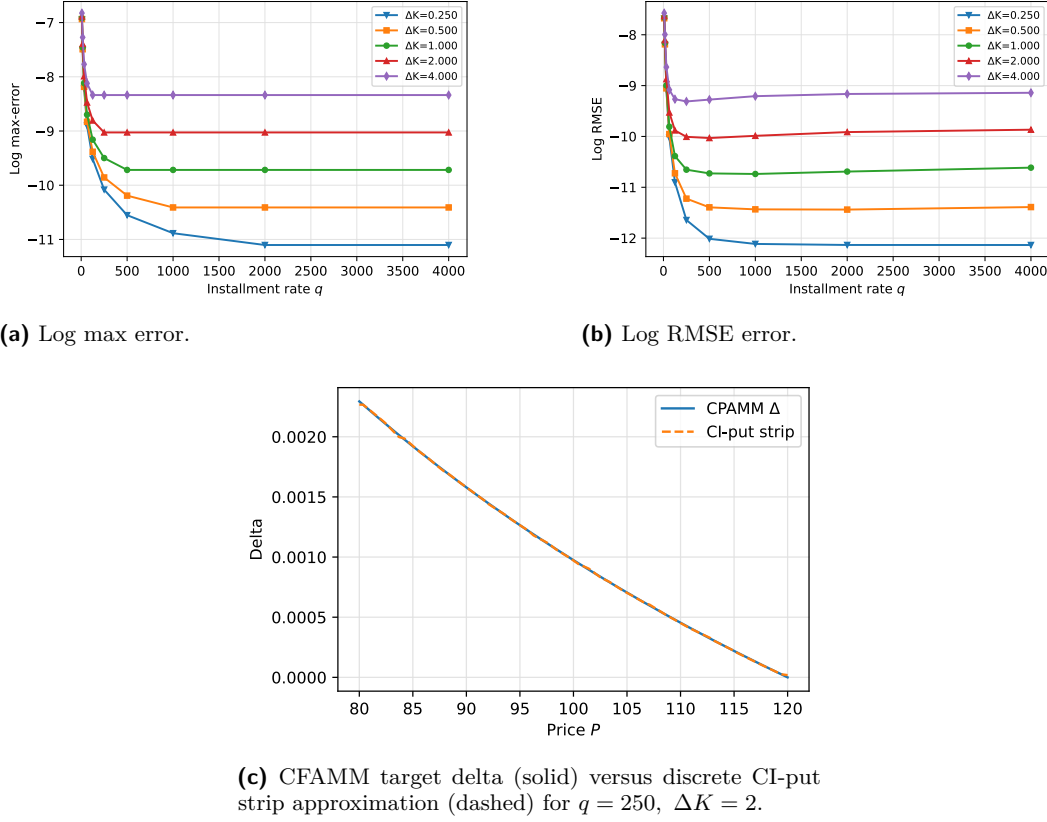


Figure 4 Log-Max absolute and RMSE delta-replication error versus installment rate q under different parameter settings.

7.3 Results


Figure 4a shows the logarithm of maximum absolute error versus q for five strike spacings. Similarly, Figure 4b displays the RMSE versus q for the same strike spacings. Both errors increase strictly with strike spacing for a given q across all installment rates. On the other hand, for a given strike spacing, both errors generally decrease with q with some exceptions. For large strike spacing, $\Delta K = \{1, 2, 4\}$, the RMSE error increases with q for large values, $q \geq 128$. Lastly, Figure 4c plots a representative delta curve ($q=250$, $\Delta K=2$) against the CPAMM target delta.

7.4 Discussion

Both errors are below 10^{-3} and shrink as expected: a larger q steepens each put's delta, while a finer ΔK better resolves the continuous weight. The trade-off between the capital cost of a large q and the operational cost of a finer strike mesh can be balanced according to the LP's precision requirements.

8 Volatility Calibration for Perpetual CI Put Options

Perpetual CI *put* pricing requires a *constant* volatility parameter σ . Market quotes instead supply a term structure $\hat{\sigma}(\tau)$ of implied volatility with time to expiration τ . In the following, we derive the effective implied volatility $\sigma_{\text{eff}}(q)$ for a perpetual American CI option with funding fee q using the market-implied term structure of at-the-money options.



■ **Figure 5** Distribution of first-exit time τ for $r = 2\%$, $\sigma = 67\%$, $K = 100$, $q = 5$.

8.1 Effective time horizon

For a CI put with rate q , the continuation band is $(S_\ell(q), S_u(q))$. It stays alive as long as the underlying spot price S_t remains within the continuation band. Otherwise, when $S_t = S_\ell$, the option is dropped, or when $S_t = S_u$, it is exercised by the holder.

Define the *first-exit time*

$$\tau(q) := \inf\{t > 0 : S_t \notin (S_\ell(q), S_u(q))\},$$

i.e., the random horizon at which the CI position terminates. In probabilistic terms, $E[\tau(q)] = \bar{\tau}(q)$ is the *mean first-exit time* of a GBM between two absorbing boundaries.

► **Theorem 7.** *The closed-form solution of $\bar{\tau}(q)$ is*

$$\bar{\tau}(q) = \begin{cases} \frac{1}{\sigma^2} \ln\left(\frac{S_0}{S_\ell(q)}\right) \ln\left(\frac{S_u(q)}{S_0}\right) & \text{if } a = 0 \\ \frac{1}{a} \left[\ln\left(\frac{S_\ell(q)}{S_0}\right) + \ln\left(\frac{S_\ell(q)}{S_u(q)}\right) \frac{S_0^\kappa - S_\ell(q)^\kappa}{S_\ell(q)^\kappa - S_u(q)^\kappa} \right] & \text{if } a \neq 0 \end{cases} \quad (17)$$

where $a = r - \frac{\sigma^2}{2}$ and $\kappa = -\frac{2a}{\sigma^2}$.

Figure 5 plots the distribution of the first-exit time, τ for a CI put with $r = 2\%$, $\sigma = 67\%$, $S_0 = K = 100$, and $q = 5$. Here, $\bar{\tau} = 1.6$ months, $\sqrt{\text{Var}(\tau)} = 0.11$, and $\mathbb{E}[|\tau - \bar{\tau}|] = 0.08$.

8.2 Practical estimation from ATM IVs

Given At-The-Money (ATM) implied volatilities $\hat{\sigma}(T)$ for fixed terms $T_1 < \dots < T_n$ and a desired effective time horizon $\tau \in [T_i, T_{i+1})$, the squared constant volatility implied by the perpetual contract can be approximated by linearly interpolating the total variances ($T\hat{\sigma}^2(T)$) derived by the market implied volatilities:

$$\sigma_{\text{eff}}^2 \tau \approx \tilde{\sigma}_{\text{eff}}^2 \tau = \hat{\sigma}^2(T_i)T_i + \frac{\hat{\sigma}^2(T_{i+1})T_{i+1} - \hat{\sigma}^2(T_i)T_i}{T_{i+1} - T_i}(\tau - T_i) \equiv w(\tau) \quad (18)$$

The effective squared volatility can be estimated ex ante for a desired first-exit time distribution:

$$\mathbb{E}[\tilde{\sigma}_{\text{eff}}^2] = \mathbb{E}\left[\frac{w(\tau)}{\tau}\right] \approx \frac{w(\bar{\tau})}{\bar{\tau}} \quad (19)$$

When $\bar{\tau}$ is a function of σ_{eff}^2 , the problem becomes a fixed-point equation. Specifically, Eq. (19) must be solved for σ_{eff}^2 as a function of itself: $\sigma_{\text{eff}}^2 = \frac{w(\bar{\tau}(\sigma_{\text{eff}}^2))}{\bar{\tau}(\sigma_{\text{eff}}^2)}$.

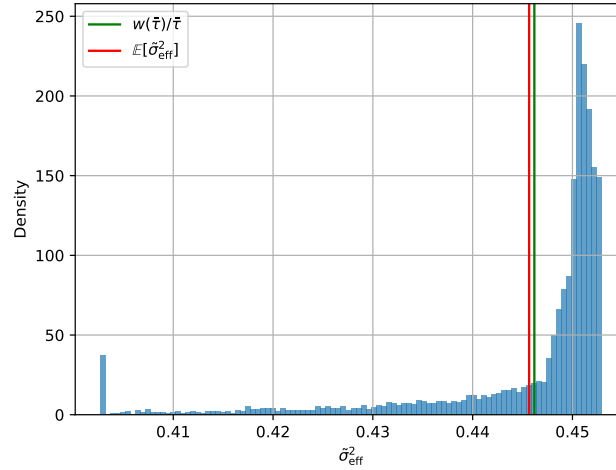
► **Theorem 8.** *The estimate $\tilde{\sigma}_{\text{eff}}^2 \approx \frac{w(\bar{\tau})}{\bar{\tau}}$ yields root mean squared error and mean absolute deviation*

$$RMSE \leq M \sqrt{\text{Var}(\tau)} \quad (20)$$

$$MAD \leq M \mathbb{E}[|\tau - \bar{\tau}|] \quad (21)$$

where $M = \max_i \sup_{\tau \in [T_i, T_{i+1})} \left| \frac{d}{d\tau} \left(\frac{w(\tau)}{\tau} \right) \right|$.

Hence, when the total variance derived from market-implied volatilities (which are approximately linear in log-Moneyness $\log(K/S)$) has a small slope, the approximation error is small. Figure 6 plots the distribution of $\tilde{\sigma}_{\text{eff}}^2$, its mean, and the approximation $\frac{w(\bar{\tau})}{\bar{\tau}}$ (for the same CI put as in Figure 5) using ETH ATM 7-day, 30-day, 90-day, and 180-day IVs. Figure 7 plots the $RMSE$ and MAD (as a percentage of $\frac{w(\bar{\tau})}{\bar{\tau}}$) for the approximation $\tilde{\sigma}_{\text{eff}}^2 \approx \frac{w(\bar{\tau})}{\bar{\tau}}$ for the period of Jan 2024-Feb 2024.




■ **Figure 6** Distribution of σ^2 for $r = 2\%$, $\sigma = 61\%$, $K = 100$, $q = 5$ using ETH ATM IVs.


8.3 Interpretation for Liquidity Providers

The volatility calibration framework enables liquidity providers (LPs) to estimate future loss-versus-rebalancing (LVR) of a concentrated AMM position using observable option market information. By associating the funding fee q of a perpetual CI put with its expected lifetime $\bar{\tau}(q)$, and mapping this to market-implied volatilities, LPs can extract an estimate for the effective squared volatility $\tilde{\sigma}_{\text{eff}}^2$ that governs the dynamics of the position and its underlying asset.

As the funding rate q increases, the continuation band $[S_\ell(q), S_u(q)]$ narrows, leading to shorter expected lifetimes $\tau(q)$ for the CI put. Conversely, smaller q implies wider bands and longer-lived options. Since market-implied volatilities are typically flatter at long durations, we can observe that:



■ **Figure 7** MAD and RMSE for $r = 5\%$, $K = 100$, $q = 40$, where σ_{eff}^2 is derived by fixed point methods. IV data is from ETH ATM IVs for Jan-Feb 2024.



■ **Figure 8** Residual upper bound rK as a percentage of q for various time horizons and σ_{eff} exposures.

- **Short durations** (high q) correspond to low maturity implied volatilities, where the IV curve is more curved and error-prone. More fine-grained market data is required here for better estimates.
- **Long durations** (low q) correspond to long-dated IVs, where the volatility curve is typically flatter. The estimate $\tilde{\sigma}_{\text{eff}}^2 = \frac{w(\tau)}{\tau}$ is then less sensitive to the exact value of τ , so M is small and the approximation of the mean is more robust to variation in τ . Long durations also tend to have a tighter concentration of realised $\tilde{\sigma}_{\text{eff}}^2$ around its mean for

■ **Table 1** Funding fee q , resulting CFAMM band (S_l, S_u) , and residual bound rK , for $r = 0.05$, for desired $\bar{\tau}$ under varying σ_{eff} exposures.

$\bar{\tau}$	σ_{eff}	% of K			Width	% of q rK
		q (token ₁ /yr)	$S_l(q)$	$S_u(q)$		
1 d	60%	284%	97%	103%	6%	2%
1 d	80%	380%	96%	104%	8%	1%
1 d	100%	475%	95%	105%	10%	1%
1 wk	60%	106%	92%	109%	17%	5%
1 wk	80%	142%	90%	112%	22%	4%
1 wk	100%	178%	87%	115%	28%	3%
2 wk	60%	74%	89%	113%	24%	7%
2 wk	80%	99%	86%	118%	32%	5%
2 wk	100%	125%	83%	123%	40%	4%
1 mo	60%	49%	85%	120%	35%	10%
1 mo	80%	66%	80%	128%	47%	8%
1 mo	100%	84%	76%	136%	60%	6%
2 mo	60%	34%	79%	130%	51%	15%
2 mo	80%	46%	74%	142%	69%	11%
2 mo	100%	58%	68%	157%	88%	9%

the same reason, meaning the mean is a good ex ante estimate.

If a liquidity band is chosen to replicate the delta of a single perpetual CI put – using the squared volatility estimate – then the LP incurs a predictable LVR almost equal to the funding rate q . This transforms an otherwise stochastic adverse-selection cost into a predictable fixed cost per unit time, simplifying the LP’s decision-making. For instance, given a desired expected time-horizon, an LP can estimate the effective term volatility, which in turn informs the liquidity band selection. Table 1 shows the different band widths corresponding to expected time horizons and effective term volatilities and Figure 8 shows the LVR – q residual term upper bound. Conversely, given a liquidity band, an LP can estimate the position’s expected effective time horizon using numerical fixed point methods. Combined with the results of Section 6, the effective term volatility estimation provides a practical framework for LPs to choose bands that realize a desired holding period $\tau(q)$ and predictable LVR, or to estimate expected forward-looking LVR for a desired liquidity band.

9 Conclusion

This work introduces a novel decomposition of a concentrated CFAMM position into a continuum of perpetual American continuous-installment (CI) put options, offering the first closed-form equivalence between the funding mechanics of perpetual American CI options and the loss-versus-rebalancing cost faced by concentrated AMM liquidity providers, providing as option-theoretic interpretation of LVR.

By exploiting the limiting behavior of CI option valuations as the installment rate grows large, the paper constructs delta-replicating portfolios that match the AMM exposure exactly. The analysis shows that the funding income from this replicating portfolio, absent in the AMM, is analytically equal to the LVR cost borne by the LP in the AMM. This time-invariant correspondence permits a forward-looking estimation of LVR and provides actionable design rules for selecting position width and shape.

Beyond theoretical insight, the framework yields practical tools for LPs. The discrete error analysis confirms that a small collection of finite- q CI puts suffices to replicate AMM delta within tight error bounds, making implementation of the replicating portfolio feasible for LPs wishing to immunize against LVR. Crucially, the analysis also shows that if a liquidity band is chosen to replicate the delta of a single perpetual CI put, the LP incurs a predictable LVR equal approximately to the funding rate q . This converts a stochastic adverse-selection cost into a predictable fixed cost per unit time, simplifying LP decision-making when it comes to position shape and width: using market-implied volatility curves, LPs can calibrate the position shape, width, and implied volatility to a desired expected holding period and LVR, or conversely, estimate a position's effective time horizon and LVR given a liquidity band.

This framework opens several avenues for further research. While the present analysis focuses only on liquidity bands centered around the current price (ATM), LPs may, in practice, deploy liquidity asymmetrically about the spot price to expose their positions to greater or lesser volatility. A resulting mismatch of spot and the replicated strike requires extended analysis to capture the entire volatility surface, as opposed to only considering the ATM volatility curve. Alternatively, if on-chain markets for CI options were developed, they could serve as direct hedging instruments and sources of IV data. Furthermore, the Black-Scholes model assumes the underlying asset experiences a constant volatility, which is not supported by market data. In reality, volatilities may be time-dependent or even stochastic. Models like the Heston model or Hull-White model, which extend upon Black-Scholes, may be applied here for analysis under dynamic volatility surfaces and to assess the model's sensitivity to deviations from constant volatility assumptions. Finally, the model assumes continuous-time trading and perfect liquidity. Future work will relax these assumptions to quantify the impact of transaction costs, slippage, and gas fees on the CI funding fee-LVR equivalence.

References

- 1 Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap V2 Core. <https://uniswap.org/whitepaper-v2.pdf>, 2020. Accessed: 2025-05-25.
- 2 Hayden Adams, Noah Zinsmeister, Moody Salem, River Keefer, and Dan Robinson. Uniswap v3 Core Whitepaper. <https://uniswap.org/whitepaper-v3.pdf>, 2021. Accessed: 2025-05-25.
- 3 Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973.
- 4 Vitalik Buterin. Ethereum white paper. *GitHub repository*, 1(22-23):5–7, 2013.
- 5 Pierangelo Ciurlia and Camilla Caperdoni. A note on the pricing of perpetual continuous-installment options. *Mathematical Methods in Economics and Finance*, 4(1):11–26, 2009.
- 6 Joseph Clark. The replicating portfolio of a constant product market. *Available at SSRN*, 2020. doi:10.2139/ssrn.3550601.
- 7 Joseph Clark. The replicating portfolio of a constant product market with bounded liquidity. *Available at SSRN*, 2021. doi:10.2139/ssrn.3898384.
- 8 DeFiLlama. DEX protocols. <https://defillama.com/protocols/dexs>. Accessed: 2025-05-26.
- 9 Jun Deng, Hua Zong, and Yun Wang. Static replication of impermanent loss for concentrated liquidity provision in decentralised markets. *Operations Research Letters*, 51(3):206–211, 2023. doi:10.1016/j.orl.2023.03.002.
- 10 Dune Analytics. Total number of liquidity pools. <https://dune.com/queries/3243692/5427065>. Accessed: 2025-05-25.
- 11 Gerald B Folland. *Real analysis: modern techniques and their applications*. John Wiley & Sons, 1999.

- 12 Masaaki Fukasawa, Basile Maire, and Marcus Wunsch. Weighted variance swaps hedge against impermanent loss. *Quantitative Finance*, 23:901–911, 2023. doi:10.1080/14697688.2023.2202708.
- 13 hagaetc. DEX metrics – Dune analytics dashboard. <https://dune.com/hagaetc/dex-metrics>. Accessed: 2025-05-26.
- 14 Alex Lipton, Vladimir Lucic, and Artur Sepp. Unified approach for hedging impermanent loss of liquidity provision. *CoRR*, 2024. arXiv:2407.05146.
- 15 Basile Maire and Marcus Wunsch. Market neutral liquidity provision. *Ledger*, 9:73–88, November 2024. doi:10.5195/ledger.2024.389.
- 16 Jason Milionis, Ciamac C Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. *CoRR*, 2024. arXiv:2208.06046.
- 17 Srisht Fateh Singh, Reina Ke Xin Li, Samuel Gaskin, Yuntao Wu, Jeffrey Klinck, Panagiotis Michalopoulos, Zissis Poulos, and Andreas Veneris. Modeling loss-versus-rebalancing in automated market makers via continuous-installment options, 2025. arXiv:2508.02971.
- 18 Anatoly Yakovenko. Solana: A new architecture for a high performance blockchain v0. 8.13, 2018.

A

 Closed-Form of Perpetual American CI Put Option

Below, we provide closed-form expressions for the perpetual American CI put option. Define

$$\gamma_p := -\frac{2r}{\sigma^2}, \quad (22)$$

$$g := 1 + \frac{rK}{q}. \quad (23)$$

Then, the option price is

$$P_q(S; K) = \alpha_p S + \beta_p S^{\gamma_p} + \frac{q}{r}$$

where the constants α_p and β_p are given by

$$\alpha_p = (g^{1-1/\gamma_p} - 1)^{-1}, \quad (24)$$

$$\beta_p = -\frac{1}{\gamma_p} \left(\frac{q}{r + \sigma^2/2} \right)^{1-\gamma_p} \alpha_p^{\gamma_p}. \quad (25)$$

Let $X_q(S; K) = \frac{\partial}{\partial S} P_q(S; K)$ be the delta of put value, its expression is given by

$$X_q(S; K) = \alpha_p + \beta_p \gamma_p S^{\gamma_p-1}.$$

Lastly, the lower and upper boundaries, S_ℓ and S_u respectively, have the following closed-form:

$$S_\ell = \frac{q}{r + \sigma^2/2} [g - g^{1/\gamma_p}],$$

$$S_u = \frac{q}{r + \sigma^2/2} [g^{1-1/\gamma_p} - 1].$$

B

 Proofs of Main Theorems

Proof of Lemma 1. Boundary collapse. Let $S_\ell(q)$ and $S_u(q)$ be the exercise and abandonment boundaries in (11) and (12). We will prove that

$$\lim_{q \rightarrow \infty} S_\ell(q) = \lim_{q \rightarrow \infty} S_u(q) = K.$$

Define $\varepsilon := rK/q$, so $g = 1 + \varepsilon$ and $\varepsilon \downarrow 0$ as $q \uparrow \infty$. Consider the second-order expansion $(1 + \varepsilon)^a = 1 + a\varepsilon + \frac{1}{2}a(a-1)\varepsilon^2 + \mathcal{O}(\varepsilon^3)$. Applying it to the two exponents in (11) and (12) yields

$$\begin{aligned} S_\ell(q) &= \frac{q\varepsilon}{r + \sigma^2/2} \left(1 - \frac{1}{\gamma_p}\right) + \mathcal{O}(q^{-1}), \\ S_u(q) &= \frac{q\varepsilon}{r + \sigma^2/2} \left(1 - \frac{1}{\gamma_p}\right) + \mathcal{O}(q^{-1}). \end{aligned}$$

Because $q\varepsilon = rK$ and $1 - \frac{1}{\gamma_p} = 1 + \sigma^2/(2r)$, both leading terms equal K . Therefore,

$$\begin{aligned} \lim_{q \rightarrow \infty} S_\ell(q) &= \lim_{q \rightarrow \infty} K + \mathcal{O}(q^{-1}) \\ &= K, \\ \lim_{q \rightarrow \infty} S_u(q) &= \lim_{q \rightarrow \infty} K + \mathcal{O}(q^{-1}) \\ &= K. \end{aligned}$$

Step-delta limit. The smooth-fit conditions (continuous first order derivative) on the boundaries of the holding region of the option valuation curve from Section 2.3.1 give $X_q(S_\ell; K) = -1$ and $X_q(S_u; K) = 0$. Because X_q is monotone increasing in S between the two boundaries and the interval $S_u(q) - S_\ell(q)$ collapses, we have the *pointwise* limit

$$X_\infty(S; K) := \lim_{q \rightarrow \infty} X_q(S; K) = -\mathbf{1}_{\{S < K\}}. \quad (26)$$

Thus, as the funding fees of a continuous-installment put becomes large, it transforms into a unit-step-delta contract. \blacktriangleleft

Proof of Theorem 2. For each K , the map $S \mapsto P_q(S; K)$ is continuously differentiable. Moreover, for fixed q , there exists a constant $c_q > 0$ such that $|\partial_S P_q(S; K)| \leq c_q(1 + K)^{-2}|w(K)|$. Hence, the integrand is point-wise dominated by an L^1 -function of K . Leibniz's rule [11] yields

$$\partial_S \Pi_q(S) = \int_0^\infty w(K) X_q(S; K) dK,$$

where $X_q := \partial_S P_q$ is the CI-put delta. For every K , we have

$$\begin{aligned} \lim_{q \rightarrow \infty} X_q(S; K) &= X_\infty(S; K) \\ &= -\mathbf{1}_{\{S < K\}} \end{aligned}$$

from Lemma 1. Let $\{q_n\}$ be a sequence such that $q_n \rightarrow \infty$. $|X_{q_n}(S; K)| \leq 1$ for all n , S and K , and $X_{q_n}(S; K) \rightarrow X_\infty(S; K)$ pointwise in K , so dominated convergence theorem [11] applies:

$$\begin{aligned} \partial_S \Pi(S) &= \lim_{n \rightarrow \infty} \int_0^\infty w(K) X_{q_n}(S; K) dK \\ &= \int_0^\infty w(K) X_\infty(S; K) dK \\ &= - \int_S^\infty w(K) dK \\ &= V'(S) - V'(\infty) \\ &= V'(S). \end{aligned}$$

Therefore, $\Pi(S) - V(S)$ is delta-neutral. \blacktriangleleft

Proof of Lemma 3. Using closed forms for a CI put boundaries from Appendix A, $g := 1 + \varepsilon$, $\varepsilon = rK/q$ and $\gamma_p = -2r/\sigma^2$:

$$S_\ell(q) = \frac{q}{r + \sigma^2/2} [g - g^{1/\gamma_p}],$$

$$S_u(q) = \frac{q}{r + \sigma^2/2} [g^{1-1/\gamma_p} - 1].$$

Setting $\alpha := 1/\gamma_p = -\sigma^2/(2r)$ and expanding to second order:

$$g^\alpha = 1 + \alpha\varepsilon + \frac{1}{2}\alpha(\alpha - 1)\varepsilon^2 + \mathcal{O}(\varepsilon^3),$$

$$g^{1-\alpha} = 1 + (1 - \alpha)\varepsilon - \frac{1}{2}(1 - \alpha)\alpha\varepsilon^2 + \mathcal{O}(\varepsilon^3).$$

Substituting the above into the closed forms of S_ℓ and S_u cancels the first-order terms and the second-order coefficient becomes $\alpha(\alpha - 1)$. Therefore, one obtains

$$S_u(q) - S_\ell(q) = \frac{\alpha(\alpha - 1)r^2K^2}{(r + \sigma^2/2)q} + \mathcal{O}(q^{-2}). \quad (27)$$

Hence,

$$\begin{aligned} \lim_{q \rightarrow \infty} q (S_u(q) - S_\ell(q)) &= \lim_{q \rightarrow \infty} \left[\frac{\alpha(\alpha - 1)r^2K^2}{(r + \sigma^2/2)} + \mathcal{O}(q^{-1}) \right] \\ &= \frac{\sigma^2K^2}{2} \end{aligned}$$

Proof of Lemma 4.

Point-wise convergence. Both $\tilde{\Pi}_q$ and Π *vanish for all* $S \geq b$. Moreover, the CI-put delta satisfies $0 \leq |X_q(S; K)| \leq 1$. Consequently $|\partial_S \tilde{\Pi}_q(S)| = |\sum_i w_i X_q(S; K_i)| \leq \sum_i |w_i| = 1$ for every q and S . Let $i^* = i^*(S, q)$ be the (unique) index with $S \in [S_\ell(q; K_{i^*}), S_u(q; K_{i^*})]$. By Lemma (1), $X_q(S; K_{i^*}) \rightarrow -\mathbf{1}_{\{S < K_{i^*}\}}$, while $X_q(S; K_i) \rightarrow 0$ for $i \neq i^*$. Dominated convergence therefore gives

$$\lim_{q \rightarrow \infty} \partial_S \tilde{\Pi}_q(S) = - \int_a^b w(K) \mathbf{1}_{\{S < K\}} dK = X(S).$$

For any $S \leq b$

$$\tilde{\Pi}_q(S) = \tilde{\Pi}_q(b) - \int_S^b \partial_S \tilde{\Pi}_q(u) du.$$

Because $\tilde{\Pi}_q(b) = 0$ for all q and the integrands converge point-wise while being uniformly bounded by 1, dominated convergence implies $\tilde{\Pi}_q(S) \rightarrow \Pi(S)$.

Limit of $w_j q$. From the statement of the Lemma,

$$w_j = X(S_u(q; K_j)) - X(S_\ell(q; K_j))$$

Therefore,

$$\begin{aligned}
 \lim_{q \rightarrow \infty} w_j q &= \lim_{q \rightarrow \infty} q [X(S_u(q; K_j)) - X(S_\ell(q; K_j))] \\
 &= \lim_{q \rightarrow \infty} q(S_u(q; K_j) - S_\ell(q; K_j)) \frac{X(S_u(q; K_j)) - X(S_\ell(q; K_j))}{(S_u(q; K_j) - S_\ell(q; K_j))} \\
 &= \lim_{q \rightarrow \infty} q(S_u(q; K_j) - S_\ell(q; K_j)) \lim_{q \rightarrow \infty} \frac{X(S_u(q; K_j)) - X(S_\ell(q; K_j))}{(S_u(q; K_j) - S_\ell(q; K_j))} \\
 &= \lim_{q \rightarrow \infty} q(S_u(q; K_j) - S_\ell(q; K_j)) X'(K_j) \\
 &= \frac{\sigma^2 S_t^2}{2} X'(S_t)
 \end{aligned}$$

Proof of Theorem 5. From Lemma 4, we can approximate the continuous portfolio by a discrete portfolio of CI puts. For a specific funding fee q , at each time t , there is only one active option j with weight w_j , price $P_q(S_t, K_j)$. The total funding fee accrued over $[t, t + dt]$ is:

$$d\text{Fee}_t^q = w_j q dt$$

Again from Lemma 4, we have $\lim_{q \rightarrow \infty} w_j q = \frac{\sigma^2 S_t^2}{2} X'(S_t)$. This implies that

$$\lim_{q \rightarrow \infty} d\text{Fee}_t^q = \lim_{q \rightarrow \infty} w_j q dt = \frac{\sigma^2 S_t^2}{2} X'(S_t) dt$$

This exactly matches Eq. (6). Hence,

$$d\text{Fee}_t = d\text{LVR}_t$$

Integrating over $t \in [0, T]$, we get:

$$\text{Fee}_t|_0^T = \int_0^T d\text{Fee}_t = \int_0^T d\text{LVR}_t = \text{LVR}|_0^T$$

Proof of Theorem 6. Equation (7) gives

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^2 P_q}{\partial S^2} + r S \frac{\partial P_q}{\partial S} - r P_q = q, \quad S \in (S_\ell, S_u).$$

and

$$\begin{aligned}
 d\text{LVR}_t &= \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 P_q}{\partial S^2} dt \\
 &= q dt - r(S \frac{\partial P_q}{\partial S} - P_q) dt
 \end{aligned}$$

Therefore, the residual term $\epsilon(t) = -r(S \frac{\partial P_q}{\partial S} - P_q)$. Note that we use S instead of S_t for brevity. $\epsilon(\cdot)$ is a function of t .

Bounding $|\epsilon(t)|$. In the region $S \in (S_\ell, S_u)$,

$$\frac{\partial \epsilon(t)}{\partial S} = -r S \frac{\partial^2 P_q}{\partial S^2}$$

Since $\frac{\partial^2 P_q}{\partial S^2}$ is always non-negative,

$$\frac{\partial \epsilon(t)}{\partial S} = -rS \frac{\partial^2 P_q}{\partial S^2} \leq 0$$

Therefore, $\epsilon(t)$ is monotonically decreasing with S . Moreover, when evaluated at $S = S_\ell$,

$$\begin{aligned} \epsilon(t) &= -r((-1 \cdot S_\ell) - (K_* - S_\ell)) \\ &= rK_*. \end{aligned}$$

and when $S = S_u$,

$$\epsilon(t) = 0.$$

Hence, $|\epsilon(t)| \leq rK_*$. ◀

C Effective Time Horizon $\bar{\tau}(q)$

The proof of Theorem 7 can be found in the extended version [17].

D Volatility Estimation Error Bounds

Proof of Theorem 8. The root mean squared error of the estimate $\tilde{\sigma}_{\text{eff}}^2 \approx \frac{w(\bar{\tau})}{\bar{\tau}}$ is

$$RMSE = \sqrt{\mathbb{E} \left[\left(\tilde{\sigma}_{\text{eff}}^2 - \frac{w(\bar{\tau})}{\bar{\tau}} \right)^2 \right]}$$

Let $f(\tau) = \tilde{\sigma}_{\text{eff}}^2 = \frac{w(\tau)}{\tau}$, which is differentiable almost everywhere, with

$$f'(\tau) = \frac{m_i \tau - w(\tau)}{\tau^2} \quad m_i \equiv \frac{\hat{\sigma}^2(T_{i+1})T_{i+1} - \hat{\sigma}^2(T_i)T_i}{T_{i+1} - T_i} \quad f(0) = m_0$$

f is Lipschitz continuous, so let $M = \max_i \sup_{\tau \in [T_i, T_{i+1})} |f'(\tau)|$. Then, $|f(\tau) - f(\bar{\tau})| \leq M|\tau - \bar{\tau}|$ and

$$RMSE = \sqrt{\mathbb{E}[(f(\tau) - f(\bar{\tau}))^2]} \leq \sqrt{\mathbb{E}[M^2(\tau - \bar{\tau})^2]} = \sqrt{M^2 \text{Var}(\tau)} = M\sqrt{\text{Var}(\tau)} \quad (28)$$

Similarly, the mean absolute deviation is

$$MAD = \mathbb{E}[|f(\tau) - f(\bar{\tau})|] \leq M\mathbb{E}[|\tau - \bar{\tau}|] \quad (29)$$
◀