

When does the tail wag the dog?

Curvature and market making

Guillermo Angeris

angeris@stanford.edu

Alex Evans

alex@placeholder.vc

Tarun Chitra

tarun@gauntlet.network

December 2020

Abstract

Liquidity and trading activity on constant function market makers (CFMMs) such as Uniswap, Curve, and Balancer has grown significantly in the second half of 2020. Much of the growth of these protocols has been driven by incentivized pools or ‘yield farming’, which reward participants in crypto assets for providing liquidity to CFMMs. As a result, CFMMs and associated protocols, which were historically very small markets, now constitute the most liquid trading venues for a large number of crypto assets. But what does it mean for a CFMM to be the most liquid market? In this paper, we propose a basic definition of price sensitivity and liquidity. We show that this definition is tightly related to the curvature of a CFMM’s trading function and can be used to explain a number of heuristic results. For example, we show that low-curvature markets are good for coins whose market value is approximately fixed and that high-curvature markets are better for liquidity providers when traders have an informational edge. Additionally, the results can also be used to model interacting markets and explain the rise of incentivized liquidity provision, also known as ‘yield farming.’

Introduction

With the advent of Bitcoin and, more generally, the blockchain, there has been a strong desire for automated censorship-resistant decentralized exchanges (DEXs). As blockchains provide a censorship-resistant means for executing programs in replicated state machines, many early DEX designs focused on emulating traditional market structure. These early attempts implemented data structures from conventional markets, such as limit order books, in smart contracts. However, due to both computational and latency constraints, blockchains often end up being suboptimal for order books. Additional notions for

Constant function market makers. One solution to this problem is the family of constant function market makers (CFMM) [4], starting with Uniswap [2, 5] which were invented as blockchain-native mechanisms for decentralized trading. CFMMs require constant space

and time interactions, unlike limit order books which have $O(n)$ space and, in many practical applications, $O(n^2)$ time complexity to process n trades. (There are some implementations achieving asymptotically better results in space, such as [34], but these designs are not, to our knowledge, used in practice.) This constant space and time requirement is ideal for blockchain environments where storage is expensive, as any data stored needs to be replicated and available at all consensus nodes, while compute can be costly for end users due to transaction fees.

CFMM agents and interactions. CFMMs are a special type of market maker that mediates the interactions between two principal agents: liquidity providers (LPs) and traders. LPs provide capital to the CFMM by locking assets in a smart contract that implements the CFMM. LPs are then incentivized to not withdraw capital in order to earn trading fees from traders who trade against liquidity providers’ locked capital. This is implemented using the following mechanism: when a liquidity provider locks their assets into a CFMM smart contract, the smart contract creates LP shares and sends them back to the LP agent. These LP shares are tokens that are essentially vouchers for the cash flows of the CFMM—they can be redeemed at a future time for the LP’s share of the CFMM’s assets and a pro-rata share of fees. For instance, if an LP provides 10% of the liquidity to a pool, then upon redemption they receive 10% of the assets held in the pool, which includes the fees accrued by the pool.

Relative liquidity of CFMMs. Prior work on CFMMs [5, 4] has analyzed necessary and sufficient conditions for CFMMs to track an external reference market with infinite liquidity. In this setting, CFMMs are modeled as secondary markets with finite liquidity whose price is adjusted by arbitrageurs to match that of the reference market. This model allows one to answer the question of whether CFMMs can serve as price oracles; *i.e.*, difficult-to-manipulate on-chain price feeds that other smart contracts can use. However, the recent advent of incentivized CFMMs (referred to as “yield farming”) has led to a number of hundred million dollar markets whose *reference* market is a CFMM with finite liquidity. In fact, Uniswap’s volume of \$440m on August 30, 2020 surpassed that of Coinbase Pro (\$380m), the largest US cryptocurrency market, making a number of markets significantly more liquid on CFMMs than on centralized order books [15, 29]. (See figure 1.) The natural next question to explore is: what happens when a CFMM becomes the *most liquid* market; *i.e.*, when the CFMM is more liquid than the external market?

In order to answer this question, we analyze how a CFMM market interacts with external markets that have finite liquidity. Our framework (§1 and §3.1) is general enough to include external markets that are limit order books or other CFMMs. For instance, the secondary market could be Uniswap and the external market could be a Balancer pool with the same assets. With this framework, we find that an analogue of the Gaussian curvature of a CFMM’s trading function dictates how much prices differ between a CFMM market and a less-liquid secondary market, after no-arbitrage is enforced. Given that CFMM trading functions need only be convex and not smooth, we construct an analogue of standard Gaussian curvature that works in the convex setting.

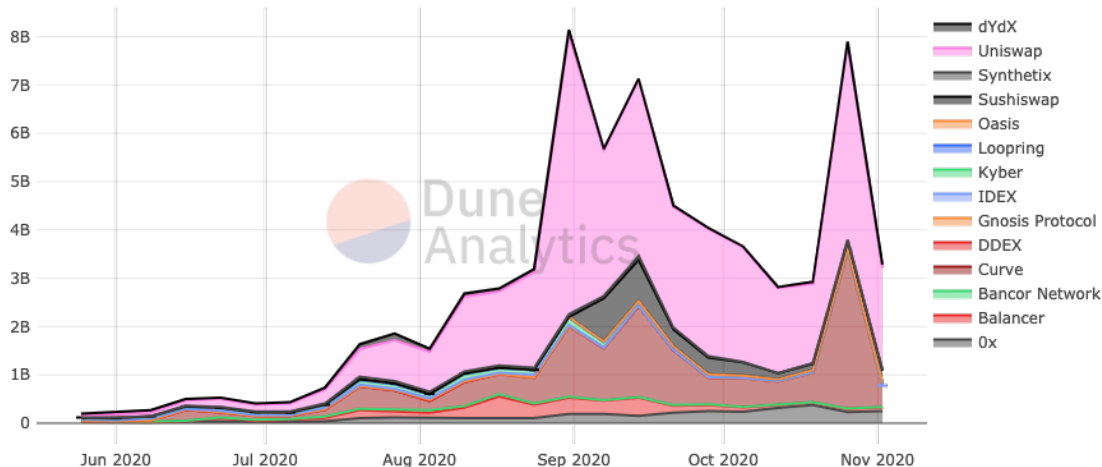


Figure 1: Decentralized exchange volumes in 2020. Uniswap has had dominant market share over the year, although it was challenged by Sushiswap in late-August. Source: Dune Analytics

Price stability. In practice, highly liquid CFMMs appear to have local price stability between pairs of reserve assets. Issuers of on-chain assets have used this property to incentivize additional liquidity in low-curvature CFMMs in order to reduce volatility in the targeted assets. A well-known example involves sUSD, a dollar-pegged ‘stablecoin’ issued through the Synthetix protocol [45]. While sUSD is intended to track the price of \$1, in practice it was quite volatile around this peg, making it less useful to users seeking a stable coin price. In response, in March 2020, Synthetix incentivized the creation of a low-curvature CFMM through Curve [22] that provided trading pairs between sUSD and other, more liquid, stablecoins. Shortly after the deployment of the CFMM, sUSD appeared significantly less volatile. Our framework provides a plausible explanation for this apparent price-stabilizing effect of low-curvature CFMMs. Using the definitions provided in §1 and a basic inequality in §3, we relate the price stability to the liquidity differences between the CFMM and the external market. While this model provides a plausible explanation for the stability in sUSD price observed between March and June 2020 (green shaded region in figure 2), it does not explain the subsequent volatility experienced between June and September 2020 (purple region). This latter period corresponds to the rise of ‘yield farming’ in the summer of 2020, wherein a number of stablecoins traded above their dollar pegs. We extend our model in §3.3 to incorporate the interaction between yield farming and curvature in CFMMs.

Liquidity provider returns. A central question, answered in [4], is to determine the expected value or payoff of holding LP shares under the assumption of no-arbitrage. This payoff can be used to compare the returns that a CFMM liquidity provider earns to those of a market maker on a traditional exchange (*e.g.*, a limit order book). Unlike limit order books, CFMMs only offer market orders and have deterministic slippage costs. The deterministic nature of slippage in CFMMs often leads to a variety of front-running attacks and deadweight

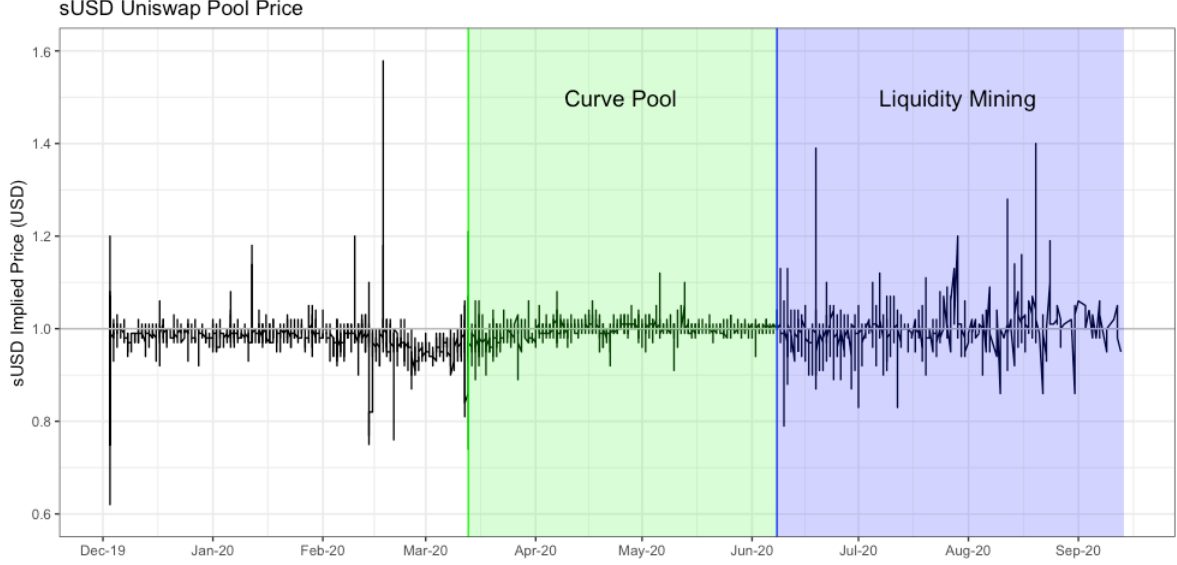


Figure 2: Implied price of sUSD on Uniswap. Shaded regions indicate the initial period after the Curve pool was deployed and the period of heavy incentivization of liquidity in on-chain CFMMs (“yield framing”), respectively. Source: Uniswap.

loss that miners and validators can capture at the expense of users [19]. However, this deterministic slippage also makes it easier for LPs to compute their expected payoff or expected loss.

Using some results from mathematical finance [13], it can be shown that Uniswap, the first CFMM to have over \$100 million in digital assets, has LP shares that should be thought of as closer to a perpetual options underwriter than a spot trading venue [18]. Combined, these properties suggest that CFMM LPs own a combination of derivative securities on the underlying tokens. With CFMMs that can dynamically adjust, such as Balancer [43], one can arbitrarily tune the payoff function of these derivative securities by adjusting the shape of the CFMM pricing curve [25]. Given that these characteristics deviate from those of conventional spot financial exchanges, it is a natural to ask what happens when a CFMM is the dominant market for a set of assets.

Practical observations. The intuition for why curvature affects actual market performance comes from empirical observations that CFMMs with lower curvature can increase LP profits for mean-reverting assets. Curve [22], which was designed for trading highly correlated assets by offering lower curvature, was able to attract \$1 billion in liquidity and reach \$350 million in daily trading volume because LPs were able to extract more fees than in Uniswap. On the other hand, Balancer provides an adjustable CFMM curve, whose curvature controls the amount of loss engendered by LPs when large price moves occur.

Uninformed trading. Using some basic results and definitions, we show in §2.1 that assets in which trades are mostly uninformed provide LPs with positive payoff whenever the curvature is low. These results also suggest that the curvature of a CFMM trading function controls how much rebalancing a given CFMM will perform. This provides a clear reason for why certain curves appear to be better in practice for mean-reverting assets (*e.g.*, Curve for stablecoin–stablecoin trades).

Informed trading. Next, we show that, in order for an LP to be profitable, the curvature of a CFMM should be large to account for the amount of information that an informed trader brings to the market. This result is an analogue of classical market microstructure results involving multiplayer games between informed traders and market makers [36, 28, 27]. Our formulation extends the informed trader framework for Uniswap of [6] to general CFMMs through a two-player game where the informed trader makes a maximum bet [12] on the next market price update given an informational advantage. Using the curvature framework of §1, we illustrate a condition (similar to Glosten’s classical bound [27]) that connects the informational edge of the informed trader and the curvature and fees of a CFMM to the payoffs of the informed trader and the LP. We extend this single period model to a multiperiod model and make conjectures about the optimal information flow in a CFMM in appendix G.

Yield farming. We also show lower bounds to so-called ‘yield farming’ payoffs that are sufficient to compensate LPs, based on the curvature difference between two markets. Yield farming can be seen as analogous to market-maker subsidies in conventional markets [11]. In yield farming, an LP first supplies reserves to a CFMM containing some token T and a numéraire such as Ethereum (ETH). The LP then locks the corresponding LP shares they receive from the CFMM and receives newly-minted tokens of T over time in return. By locking LP shares in the smart contract, users that provably provide liquidity for the T –ETH trading pair are subsidized. In practice, this incentive bootstraps liquidity for token T by incentivizing users to be LPs [20]. We show that these incentives can be bounded from below by the curvature of a CFMM.

Portfolio Greeks. Finally, we apply our results regarding curvature to analyze financial properties of CFMM LP shares. First, we extend the frameworks of [4, 18] to compute the first and second order Greeks (*e.g.*, Δ, Γ) for CFMM LP share payoffs in the two coin case, and then use the basic liquidity framework provided in §1 to give bounds on the portfolio Greeks. While we note that these results, due to the assumptions, may be of limited practical importance, they can be used to interpret the curvature bounds in several different ways. We leave a suitable and practical version of the result as an open conjecture.

Summary. These results show that curvature is a crucial design parameter to tune when designing CFMMs. A number of CFMM designs have been proposed for prediction markets [39], derivatives trading [1, 35], and self-balancing ETFs [43]. In each of these appli-

cations, a CFMM LP share represents a complex payoff function that changes dramatically based on the expected trading behavior of the assets held within the CFMM. Our results show that the returns and payoffs realized by holders of CFMM LP shares are intrinsically tied to the curvature of the trading function. Moreover, they explain a number of empirical outcomes that have happened throughout the large number of different CFMMs available on Ethereum. CFMM designers need to be cognizant of trade-offs that are made by adjusting curvature, especially as payoff functions become increasingly complex. These results also provide guidance on how to parameterize yield farming incentives to achieve certain liquidity targets. As a number of yield farming assets have failed due to over-incentivization of liquidity [44, 33], it is increasingly important to understand how to efficiently incentivize on-chain liquidity. These results provide ways to sensibly optimize incentives to meet liquidity goals.

1 Two asset market model

In this section, we define basic terms and models used throughout the remainder of the paper. In particular, we will define what it means for a market to be ‘stable’ or ‘liquid.’ We note that, while some markets in practice can simultaneously trade n coins for n coins, we will focus on the case where the market only trades two coins, and this is the model we will use in the majority of the paper. (We will sometimes present the n coin generalizations, as in, *e.g.*, appendices E and I, when they are simple, but this is the exception rather than the rule.) In this case, we will call one asset the *traded coin*, and this asset is distinct from the *numéraire*, which is the base currency used to measure prices. Unless otherwise specified, all trades (which we will often simply call Δ) are positive when they are buying Δ amount of the traded asset, and negative when they are selling it.

Price impact function. We will define the *price impact* function $g : \mathbf{R} \rightarrow \mathbf{R}_{++}$ of a market to be the function that connects the market’s marginal price before the trade, which we will call m_0 , to the price after the trade. More specifically, we have that $g(0) = m_0$, while $g(-\Delta)$ specifies the CFMM’s marginal price immediately after the trade, which sells Δ of the traded coin to the market, is performed.

We assume two basic facts about g : one, that g is continuous, and, two, that g is nondecreasing. In other words, g is a continuous function that expresses how the market’s price changes after a given trade, with the assumption that trades of size 0 (*i.e.*, null trades) don’t change the market price, and that larger trades lead to higher marginal prices. Note that these assumptions are common in the order-book literature (see, *e.g.*, [7]) and true for all CFMMs (see, *e.g.*, [4]). Additionally, this assumption is equivalent to the convexity of the quantity function $q(\Delta) = \int_0^\Delta g(t) dt$, which is also a common assumption in the classical economics literature (see [26]).

No-arbitrage. A common way to model interactions between different markets is the assumption of no-arbitrage. One way of stating this assumption is through the existence

of an agent, called an *arbitrageur*. This agent is allowed to borrow any amount of coin Δ , trade it between the available markets (to receive some amount of coin, say, Δ^a), and then pay back the borrowed amount Δ , to receive $\Delta^a - \Delta$ profit. If there exists a trade which guarantees that $\Delta^a > \Delta$, we then say that there is an *arbitrage opportunity* which we assume the arbitrageur will execute to receive strictly positive profit.

In our presentation, we will assume that there is an (infinitely-liquid) reference market with fixed price m_a . (We provide a generalization to reference markets which do not have infinite liquidity in §3.) An arbitrageur would then attempt to maximize its profits by trading some amount of coin Δ between both markets, which would give a total payoff of

$$\int_0^\Delta (m_a - g(-t)) dt.$$

A necessary condition for this quantity to be maximized is that the marginal price of both markets after the no-arbitrage trade, which we will call Δ^* , must be equal; *i.e.*, that Δ^* must satisfy

$$m_a = g(-\Delta^*), \tag{1}$$

which follows from the first-order optimality conditions applied to the arbitrageur's payoff.

Without loss of generality, we will assume that $m_a \leq m_0 = g(0)$. Note that this is always possible in a two-coin economy by changing the choice of the coins to be traded; *i.e.*, by swapping the places of the traded coin and the numéraire, the resulting market prices are $1/m_a \leq 1/m_0$, which satisfy the above inequality. Because of this assumption and the fact that g is a nondecreasing function, we will have that $\Delta^* \geq 0$. (We discuss this further in the case where the market is not infinitely liquid in §3.1.)

Price stability. We will say that the price impact function g is μ -stable (with $\mu \geq 0$) if it satisfies:

$$g(0) - g(-\Delta) \leq \mu\Delta. \tag{2}$$

In other words, we say that the price impact function g for some market is μ -stable whenever a (nonnegative) trade of size Δ does not change the market's price by more than $\mu\Delta$. Because g is increasing by assumption, when $\Delta \geq 0$, inequality (2) is equivalent to the sensitivity-like bounds,

$$|g(0) - g(-\Delta)| \leq \mu\Delta,$$

as both sides are nonnegative.

There are several useful sufficient conditions for (2) to hold. For example, it suffices that its first derivative is bounded from above by μ for all $\Delta \geq 0$; *i.e.*,

$$\frac{\partial g(-\Delta)}{\partial \Delta} \leq \mu,$$

but this is not a necessary condition, as the function g need not be differentiable in its second argument. We show a connection between this sufficient condition and the curvature of a trading function for a given CFMM in appendix C; we also discuss explicit bounds for μ for some CFMMs used in practice, and how they relate to common intuition, in the following section.

Liquidity bounds. We will say that a price impact function for a market is κ -*liquid* if it satisfies

$$g(0) - g(-\Delta) \geq \kappa\Delta. \quad (3)$$

In other words, selling Δ coin to the market decreases the reported price by at least $\kappa\Delta$, which implies that there is some amount of price slippage that is linearly bounded from below by a factor of at least κ . Additionally, we note that it is possible (and, in fact, common for many markets) that a trading function is both μ -stable and κ -liquid, with $\kappa \leq \mu$.

Trade sizes. It is often the case that bounds of the form of (3) (and, sometimes, bounds of the form of (2), as we will see later in this section) are not global, but, instead, hold over some interval of size L ; *i.e.*, such bounds only hold for $0 \leq \Delta \leq L$. In these specific cases, we will mention the corresponding conditions in the statements, as required for the results to hold for CFMMs in practice. On the other hand, we note that all of the proofs presented have results which immediately carry over in this case, even when the interval is not explicitly mentioned. (We leave such extensions as simple exercises for the reader.)

1.1 CFMM curvature

In decentralized finance, the market we are studying is almost always a *constant function market maker* or CFMM. (See, *e.g.*, [4] for an introduction.) In this case, the markets' behaviors are specified by (often simple) mathematical formulae, and often have closed-form solutions for the constants μ and κ . In this section, we will show how to compute μ and κ for some of the CFMMs used in practice.

Constant function market makers. A CFMM is an algorithmic market maker [48, 47, 32] defined by its *reserves*, specifying how much of each coin is available for trading, and its *trading function*, which controls whether the market maker will accept or reject a proposed trade. The reserves are given by $R \in \mathbf{R}_+$ for the coin to be traded, and $R' \in \mathbf{R}_+$ for the numéraire coin, while its trading function is given by $\psi : \mathbf{R}_+^2 \times \mathbf{R}^2 \rightarrow \mathbf{R}$. The trading function maps the pair of reserves $(R, R') \in \mathbf{R}_+^2$ and a trade, purchasing Δ of the coin to be traded and Δ' of the numéraire, $(\Delta, \Delta') \in \mathbf{R}^2$, to a scalar value.

When a CFMM has reserves (R, R') , any agent may propose a trade (Δ, Δ') . By definition, the CFMM accepts the trade whenever

$$\psi(R, R', \Delta, \Delta') = \psi(R, R', 0, 0).$$

The CFMM then pays out Δ of the traded coin to the trader and receives Δ' of the numéraire. This results in the following update to the reserves: $R \leftarrow R - \Delta$ and $R' \leftarrow R' + \Delta'$. (Negative values of Δ and Δ' reverse the flow of the coin.)

For notational convenience, we will abuse notation slightly by writing $\psi(\Delta, \Delta')$ for $\psi(R, R', \Delta, \Delta')$, such that the function $\psi(\cdot, \cdot)$ implicitly depends on the reserve values (R, R') in the remainder of the paper. Additionally, because our focus is on the two coin case and

not the general n coin case, we use different notation than [4] to prevent overly-cumbersome proofs and results. We show the exact connection between both forms in appendix A.

Marginal prices. Given a CFMM with trading function ψ , the marginal price at these reserves is given by [4, §2.4]:

$$g(\Delta) = -\frac{\partial_1 \psi(\Delta, \Delta')}{\partial_2 \psi(\Delta, \Delta')}, \quad (4)$$

where $\partial_i \psi$ denotes the partial derivative of ψ with respect to the i th argument and $\Delta' \in \mathbf{R}$ is the (usually unique) solution to

$$\psi(\Delta, \Delta') = \psi(0, 0),$$

for a given $\Delta \in \mathbf{R}$. Note that this is only defined whenever $\Delta \leq R$ and $\Delta' \geq -R'$; *i.e.*, when there are enough reserves to complete the trade. Because it is often the case that ψ satisfies the condition $\Delta' = \int_0^\Delta g(-t) dt \leq R'$, even as $\Delta \uparrow \infty$, we will consider these bounds to be implicit, unless otherwise stated.

Marginal prices with fees. While it is possible to implicitly include fees in the definition of the CFMM's trading function, it is often simpler to include the fee explicitly. In many cases such as Uniswap, Balancer, and Curve, the fee is given by some number $0 < \gamma \leq 1$ such that $(1 - \gamma)$ is the percentage fee taken for each trade, and the fee-less CFMM, with trading function ψ , is modified in the following way [4, §3.1]:

$$\psi^f(\Delta, \Delta') = \psi(\gamma\Delta, \Delta') = \psi(0, 0),$$

where ψ^f is the CFMM trading function with fees, for trades which sell some amount of coin $\Delta \leq 0$ to the CFMM. The reserves are updated in a similar way as the original CFMM. We note that the case where $\Delta \geq 0$ can be derived by appropriately exchanging the traded coin and the numéraire. The directionality here comes from the fact that fees are usually charged ‘on the way in,’ or, in other words, asymmetrically charged to the coin being sold to the CFMM. (See, *e.g.*, appendix A.)

In this case, we can write the marginal price of a given trade of size $\Delta \leq 0$ after fees in terms of the marginal price of the original CFMM, since

$$g^f(\Delta) = -\frac{\partial_1 \psi^f(\Delta, \Delta')}{\partial_2 \psi^f(\Delta, \Delta')} = -\frac{\partial_1 \psi(\gamma\Delta, \Delta')}{\partial_2 \psi(\gamma\Delta, \Delta')} = \gamma g(\gamma\Delta),$$

where Δ' is the (usually unique) solution to $\psi(\gamma\Delta, \Delta') = \psi(0, 0)$, and, as before, $\Delta \leq 0$; *i.e.*, we are selling the coin to be traded to the CFMM.

Given that we can express the price impact function g^f with fees in term of the fee-less price impact function g , the next problem is to find bounds of the form of (2) for fee-less CFMMs, which we show for a few special cases.

Constant sum market maker. The simplest example of a CFMM is the constant sum market maker, whose trading function is the *constant sum* trading function:

$$\psi(\Delta, \Delta') = (R - \Delta) + (R' + \Delta').$$

In this case, the marginal price is, from (4):

$$g(\Delta) = 1,$$

whenever $-R' \leq \Delta \leq R$ and is otherwise undefined. In this case, we have the following bound, whenever $0 \leq \Delta \leq R'$:

$$g(0) - g(-\Delta) = 0 \leq \mu\Delta,$$

with curvature bound $\mu = 0$. (This is simply due to the fact that a constant sum market maker always reports a fixed price, so long as it has nonzero reserves.) Similarly, we also have $\kappa = 0$, on the same interval, by the same argument.

Global bounds on μ for convex impact. While the constant sum market maker has a simple enough trading function that it can be analyzed directly, analyzing other trading functions in the same way can quickly lead to very complicated results. A useful and simple condition applies in the common case that the price impact function, g , is a differentiable convex function. In this case, we have that, for all valid Δ ,

$$g(-\Delta) \geq g(0) - g'(0)\Delta,$$

where $g'(0)$ is the derivative of g evaluated at zero, by the first-order condition for convexity [8, §3.1.3]. This can be written as

$$g(0) - g(-\Delta) \leq g'(0)\Delta. \tag{5}$$

Setting $\mu = g'(0)$ yields the desired result. When g is differentiable, taking the limit as $\Delta \downarrow 0$, shows that this is the tightest possible bound on μ over any nonzero interval size. (If g is not differentiable, we can take μ to be the largest subgradient of g at 0, which is the tightest possible bound by the same argument.)

Local bounds on κ for convex impact. On the other hand, given any interval of size $L \geq 0$, we can give a κ -liquidity bound for g , by noting that, because g is convex, the definition of convexity gives the following inequality:

$$g(-\Delta) = g\left(\left(1 - \frac{\Delta}{L}\right)0 + \frac{\Delta}{L}(-L)\right) \leq \left(1 - \frac{\Delta}{L}\right)g(0) + \frac{\Delta}{L}g(-L).$$

A basic rearrangement shows:

$$g(0) - g(-\Delta) \geq \left(\frac{g(0) - g(-L)}{L}\right)\Delta,$$

and setting $\kappa = (g(0) - g(-L))/L$ gives the result. Since the bound is tight at $\Delta = 0$ and $\Delta = L$, this κ yields the tightest possible bound along this interval.

Uniswap. One of the simplest nontrivial convex bounds is the bound for Uniswap (or constant product market maker) with no fees, where we can write

$$g(\Delta) = \frac{k}{(R - \Delta)^2}, \quad (6)$$

where $k = RR'$ is the product constant [5, App. A].

If $R > \Delta$ (*i.e.*, there are enough reserves to carry out a trade of size Δ) this function is a convex function, as $x \mapsto 1/x^2$ is convex over the positive reals. Using inequality (5), we have, for all $\Delta \geq 0$,

$$g(0) - g(-\Delta) \leq 2 \frac{k}{R^3} \Delta = 2 \frac{g(0)}{R} \Delta = \mu \Delta.$$

In the special case where the marginal price at the zero trade is $g(0) = 1$ (as is common with stablecoin–stablecoin markets), which happens when $R = R'$, we can re-write μ in terms of the portfolio value of the reserves as

$$\mu = \frac{4}{P_V},$$

where the portfolio value is given by $P_V = g(0)R + R' = 2R$.

Similarly, for any interval size $L \geq 0$, we have

$$\kappa = \frac{k}{L} \left(\frac{1}{R^2} - \frac{1}{(R + L)^2} \right) = \frac{g(0)}{L} \left(1 - \frac{1}{(1 + L/R)^2} \right).$$

Note that both μ and κ both decrease when R increases as liquidity is added (*e.g.*, via Uniswap’s `addLiquidity` function) for a fixed price $g(0)$ and interval size L . In other words, Uniswap’s effective curvature *decreases* as the reserves increase, as one might intuitively expect. We can interpret this as, for a fixed trade cost, larger trades can take place in Uniswap with higher reserves. An alternative interpretation is that the cost of manipulation also increases in the reserve size R . These results were proven in [5, §2.3] using different techniques which do not easily generalize.

Two asset Balancer. For Balancer (which is also sometimes called the constant mean market maker [5, §3], or the geometric mean market maker [25]) with two assets and weight $\tau \in (0, 1)$, we have the trading function

$$\psi(\Delta, \Delta') = (R - \Delta)^\tau (R' + \Delta')^{1-\tau}.$$

Let $\xi = \frac{\tau}{1-\tau}$ for notational convenience. We then have:

$$g(\Delta) = \frac{d\Delta'}{d\Delta} = \left(\frac{\tau}{1-\tau} \right) \frac{k^{1/(1-\tau)}}{(R - \Delta)^{1+\tau/(1-\tau)}} = \frac{\xi k^{\xi/\tau}}{(R - \Delta)^{1+\xi}}.$$

Note that when $\tau = \frac{1}{2}$ then $\xi = 1$ and this price impact function is equal to Uniswap’s, generalizing result (6). This function is convex since $x \mapsto x^{-(1+\xi)}$ is convex over the positive reals for any $\xi > -1$. This implies

$$g(0) - g(-\Delta) \leq \xi(1 + \xi) \frac{k^{\xi/\tau}}{R^{2+\xi}} \Delta = (1 + \xi) \frac{g(0)}{R} \Delta = \mu \Delta.$$

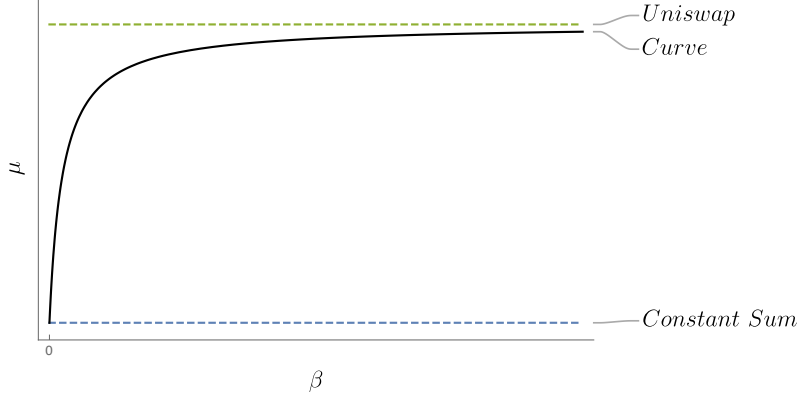


Figure 3: Curvature constant μ for Curve for different values of β , where $g(0) = 1$.

As with Uniswap, we can give a simple expression for μ in terms of the portfolio value in the special case where $g(0) = 1$:

$$\mu = \frac{1}{\tau(1 - \tau)P_V}$$

where the portfolio value is given by $P_V = g(0)R + R' = R/\tau$. (Note that the expression for μ is symmetric about $\tau = 1/2$, even though the portfolio value P_V is not.)

Similarly, for any interval of size L , we have that

$$\kappa = \frac{\xi k^{\xi/\tau}}{L} \left(\frac{1}{R^{1+\xi}} - \frac{1}{(R + L)^{1+\xi}} \right) = \frac{g(0)}{L} \left(1 - \frac{1}{(1 + L/R)^{1+\xi}} \right)$$

As before, we have that both μ and κ are decreasing functions in the reserves R for a fixed marginal price $g(0)$ and interval length L . We can additionally recover the Uniswap bounds on κ and μ by setting $\tau = 1/2$, or, equivalently, $\xi = 1$.

Two asset Curve. Another very popular CFMM is Curve [21], with trading function

$$\psi(\Delta, \Delta') = \alpha((R - \Delta) + (R' + \Delta')) - \beta((R - \Delta)(R' + \Delta'))^{-1}.$$

It is worth noting that, as α/β becomes large, then ψ is approximately close to the trading function for the constant sum market maker. Similarly, as α/β becomes small, ψ becomes similar to the Uniswap trading function. (See, *e.g.*, figure 3, which shows that μ for Curve converges to the curvature constants for constant sum as $\beta \downarrow 0$ and Uniswap as $\beta \uparrow \infty$.) One can imagine variations on Curve where the product term is replaced by the reciprocal of the weighted geometric mean, as with Balancer, or a number of other functions.

The marginal price function for Curve is relatively complicated, but can be derived with some work (see [4, §2.4]):

$$g(\Delta) = \frac{\alpha(R - \Delta)(R' + \Delta') + \beta(R - \Delta)}{\alpha(R - \Delta)(R' + \Delta') + \beta(R' + \Delta')},$$

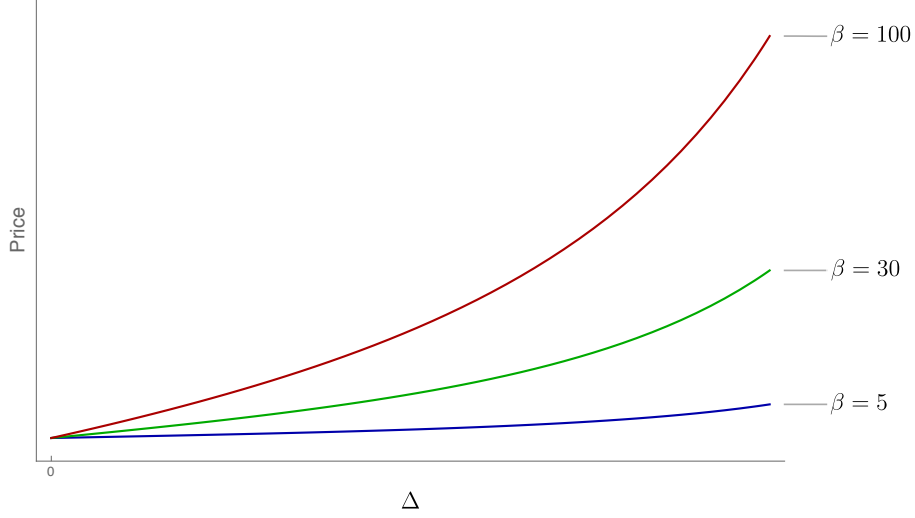


Figure 4: Curve price impact function plotted for different values of β . Here, $g(0) = 1$.

where

$$\Delta' = R' + \frac{1}{2\alpha} \left(\sqrt{(\alpha(R - \Delta) - k)^2 - 4\alpha\beta(R - \Delta)^{-1}} - (\alpha(R - \Delta) - k) \right),$$

and we have defined $k = \psi(0, 0) = \alpha(R + R') - \beta(RR')^{-1}$ for notational convenience. From figure 4 we see that g is indeed convex for a number of parameters β (setting $\alpha = 1$ without loss of generality, as g is homogeneous of degree zero with respect to (α, β)). Showing that g is convex is rather more involved; we provide a proof in appendix D.

While, in general, $\mu = g'(0)$ can yield complicated expressions for Curve, the special case where $g(0) = 1$ can be written as a simple function of the portfolio value:

$$\mu = \frac{32\beta}{8\beta P_V + \alpha P_V^4},$$

where $P_V = g(0)R + R' = 2R$. This provides a simple expression for the maximum slippage a trader can expect for a given trade size when assets on Curve are trading at their peg.

1.2 Examples in practice

One of the simplest examples of curvature impacting trading performance is in stablecoin trading venues. Stablecoins, which are assets that aim to be approximately pegged to a fiat numéraire such as the US dollar, are assets with an approximately constant price due to their peg. However, their prices and outstanding supply often differ for systematic reasons. For instance, one type of asset might only be centralized and only allowed to be created by non-US entities (*e.g.*, Tether), whereas another asset is more decentralized (*e.g.*, MakerDAO). The former stablecoin might have a small number of large institutions using the coin whereas the latter is likely to have more small participants. If this is the case, it will naturally be

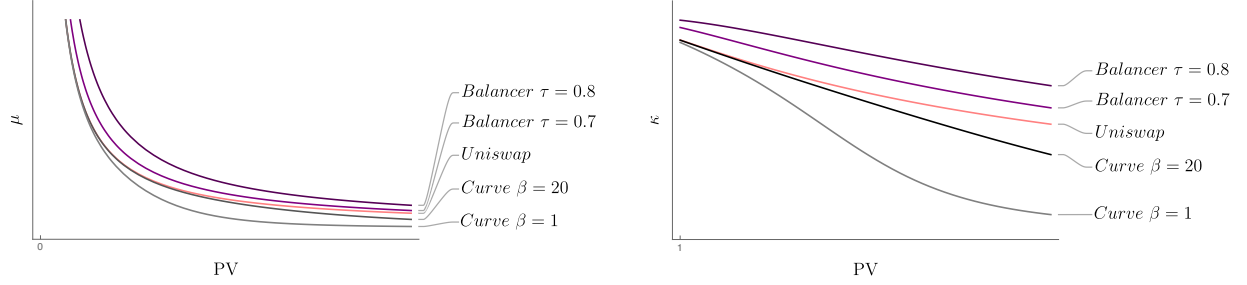


Figure 5: (Left) Plots of μ for Uniswap, Balancer, and Curve for different parameter choices as a function of the total reserve value for $g(0) = 1$. (Right) Plots of κ for Uniswap, Balancer, and Curve as a function of the total reserve quantity for $g(0) = 1$ fixed and $L = 1$.

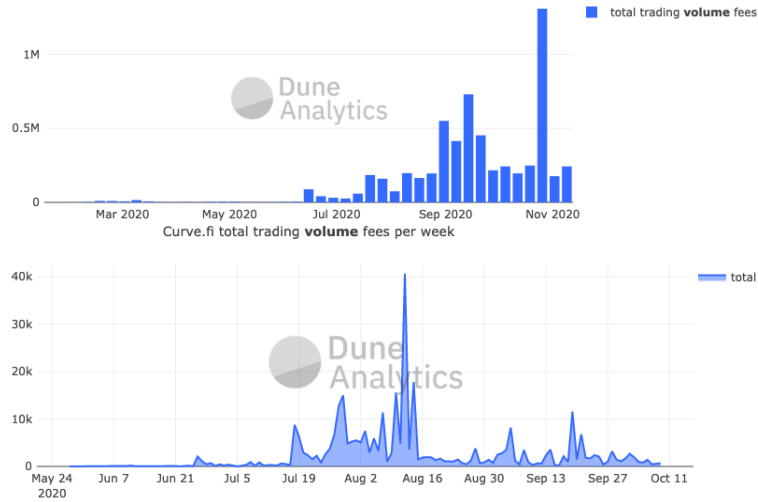


Figure 6: Aggregated trading fees for Curve (top) and mStable (bottom) in USD. Source: Dune Analytics.

easier to perform bigger trades in the former currency. A CFMM designed for stablecoin–stablecoin trading, such as Curve, needs to have curvature that is adapted for these types of trading.

Uninformed trading. Generally speaking, the trading of stablecoins for stablecoins tends to be uninformed. That is, users often trade these stablecoin pairs because a specific smart contract or exchange that they want to interact with only allows the use of a particular stablecoin. There is no information about the direction of trading and, given the ease of creation-redemption arbitrage for stablecoins, there is little use in trying to predict stablecoin prices [40, 31].

Curvature. Intuitively, then, venues for stablecoin–stablecoin trades should have relatively low-curvature price impact functions, which would entice traders due to the small price slippage, and entice liquidity providers due to the small opportunity cost. Taking this idea to its extreme, one might argue that assets that are supposed to be the same value should be traded on a CFMM with zero curvature. An example of such a market is mStable, which uses the constant sum trading function presented in §1.1. These curvature-less markets have trouble responding to price, as they effectively quote a fixed price for any trade performed. In practice, as illustrated in figure 6, we see that Curve generates almost an order of magnitude more in trading fees for LPs than mStable. This is driven by two phenomena. First, a zero-curvature AMM will end up less liquid in practice because it quickly runs out of reserves as price fluctuates. In the case of mStable, there is a chronic shortage of Dai for trades as Dai is frequently trading above its peg. Second, LPs face maximal opportunity cost relative to a low-curvature CFMM with the same assets and fees. Most of the trading volume and liquidity provision on mStable appears instead to be driven by *yield farming*. We discuss such incentives in more detail later in §3.3.

2 LP returns and curvature

Empirically, it has been observed that returns to LPs in CFMMs are closely tied to both the shape of the CFMM trading set and the properties of the price process of the two assets. Curve’s advantage over Uniswap for mean-reverting, low volatility assets led to it attracting significantly more trading volume for certain assets. As shown in §1.1, Curve has a low curvature regime around a particular price and a high curvature region far away from this price. This design was chosen to optimize profits earned by LPs for mean reverting assets while allowing traders to place large sized orders when assets are near their mean. A natural question to ask is: how much does adjusting the curvature of a CFMM for such assets affect LP returns?

Liquidity provider portfolio value. The LP portfolio value is defined as the value of coins that an LP has locked in a CFMM. If a user owns $b \in [0, 1]$ percent of the LP shares in a CFMM, then they own the right to claim bR of asset 1 and bR' of asset 2, where

R and R' are, as before, the reserves of the CFMM. In this subsection, we will formalize heuristics arguments of [21, 23] which show that LP portfolio values are directly affected by the curvature of a trading function. We will also generalize these claims to generic price processes interacting with CFMMs by considering adverse selection towards LPs. In particular, we will consider how LP returns are affected by *informed traders*, who have an estimate for the probability distribution of future prices. These results will illustrate that the design of efficient CFMMs for a variety of markets depends on how the curvature is adjusted to ensure that LPs can be profitable. In particular, we will see that, for stablecoins, where most trades are uninformed, low curvature improves performance and LPs still come out with positive profit even for large trades, while markets where there exist informed traders with a bigger edge, higher curvature is preferable to prevent downside LP losses.

2.1 Uninformed trading and low curvature CFMMs

In this scenario, we consider a trader who wishes to buy some amount of coin $\Delta \geq 0$ from the market. Such purchases will cause some amount of slippage in the reported price g , causing some loss to the LP, which may be recouped with fees. The question is, assuming that this CFMM is the only available market, what is the largest trade that a trader can perform such that a liquidity provider still has positive payoff from the trade?

Curvature and profits. Formally, suppose that an LP provides all of the assets $R, R' \in \mathbf{R}_+$ to a CFMM with μ -stable price impact function g . We assume that the CFMM charges some fee $(1 - \gamma)$, but this fee is not included in g . The no-fee portfolio value of this CFMM is:

$$g(0)R + R'.$$

After a price change to $m_a = g(-\Delta) \leq g(0)$, the opportunity cost (sometimes called the ‘impermanant loss’) of this portfolio is given by

$$\underbrace{(g(-\Delta)(R + \Delta) + R' - \Delta')}_{\text{LP portfolio value}} - \underbrace{(g(-\Delta)R + R')}_{\text{Equivalent portfolio value}} = g(-\Delta)\Delta - \Delta',$$

and, by definition of marginal price

$$\Delta' = \int_0^{-\Delta} g(t) dt.$$

Since g is a nonincreasing function, we have

$$\Delta' = \int_0^{-\Delta} g(t) dt \leq g(0)\Delta,$$

so,

$$g(-\Delta)\Delta - \Delta' \geq (g(-\Delta) - g(0))\Delta \geq -\mu\Delta^2,$$

is a lower bound on the opportunity cost. Here, the second inequality follows from the definition of μ -stability.

On the other hand, if g is a marginal price function with some fee $0 < \gamma \leq 1$, the value of fees earned is at least $(1 - \gamma)\Delta g(-\Delta) = (1 - \gamma)\Delta m_a$, where m_a is the new price (see appendix B for a general statement and proof), so LPs are guaranteed to make a profit whenever

$$(1 - \gamma)\Delta m_a > \mu\Delta^2,$$

which happens when

$$\Delta < \frac{(1 - \gamma)m_a}{\mu}. \quad (7)$$

Discussion. This inequality shows that a sufficient condition on the trade size for which LPs still make a profit is inversely proportional to the curvature bound on the CFMM. Additionally, using this formula, we can compute a lower bound on the rate of growth of profits for LPs as a function of fees and curvature, given a distribution of trades, $\mathbf{Prob}[\Delta \leq x]$. An extension of this result can provide a discrete time, curvature-based analogue of [51] that generalizes to a number of CFMMs other than Uniswap.

Another interpretation of inequality (7) is that, as the effective curvature decreases, traders can perform large trades, while liquidity providers still come out ahead, relative to an equivalent portfolio which simply holds R of the traded coin and R' of the numéraire. Note that this inequality comes from the fact that the trade does not depend on the future price of the coin, which we will call an ‘uninformed’ trade, and such trades are, as discussed previously, very common in stablecoin–stablecoin trades.

2.2 Informed trading

On the other hand, liquidity provider losses change drastically when we have an agent who attempts to maximize their profits given information about future prices. In order to model this phenomenon, we need to describe a participant other than the LP, who has some amount of knowledge of future prices. Analogous to [6] and the classical market microstructure models [36, 28], we will consider a market with an LP and an informed trader under the assumption of no-arbitrage. We will construct a two-player game between an informed trader who can predict the next price update of some external market with non-trivial edge, and a liquidity provider whose funds are locked in the CFMM. Using this game, we will show a profit (or loss) lower bound for both LPs and informed traders, akin to those used to describe market maker profits in open limit order books [28, 27]. From this lower bound, we will show that informed traders need less of an informational edge to guarantee that trading with a lower curvature CFMM has profits as large as trading with a higher curvature CFMM.

Problem set up. In this game, as before, we have two agents: a liquidity provider and an informed trader, where the informed trader is allowed to trade with the CFMM.

We will assume that the CFMM, with fee-less marginal price function g (such that the marginal price with fees is $g^f(-\Delta) = \gamma g(-\gamma\Delta)$ with $\Delta \geq 0$) and the reference market both start at some fixed price $m_0 = \gamma g(0)$. We will assume the function g is μ -stable and κ -liquid in some interval $0 \leq \Delta \leq L$. The informed trader then knows that the reference market price will decrease to some amount $m_1 \leq m_0$ with probability α or stay at m_0 with probability $(1 - \alpha)$. By no-arbitrage, any price discrepancies between the CFMM and the reference market are immediately removed, so the informed trader must make a trade which maximizes the expected profit, before the trader is able to see the new price.

Informed trader edge. The expected edge of an informed trader under this framework is given by

$$E_V(\Delta) = \underbrace{\int_0^\Delta \gamma g(-\gamma t) dt}_{\text{total profit for trading } \Delta} - \underbrace{(\alpha m_1 \Delta + (1 - \alpha) m_0 \Delta)}_{\text{expected profit for holding } \Delta \text{ asset}}.$$

We can rewrite this in a slightly simpler form by noting that, by assumption, $m_0 = \gamma g(0)$, so

$$E_V(\Delta) = \gamma \int_0^\Delta (g(-\gamma t) - g(0)) dt + \alpha(m_0 - m_1)\Delta.$$

Using the fact that g is μ -stable, we have $g(-\gamma t) - g(0) \geq -\mu\gamma t$, and that

$$E_V(\Delta) \geq \alpha(m_0 - m_1)\Delta - \frac{1}{2}\mu\gamma^2\Delta^2.$$

Taking the supremum of both sides over $\Delta \geq 0$ (since an informed traders seek to maximize their profit) gives that

$$E_V^* \geq \frac{\alpha^2(m_0 - m_1)^2}{2\mu\gamma^2}.$$

Here, the surprising fact that the fee γ appears in the denominator (versus, as one might expect, in the numerator) happens because the price $m_0 = \gamma g(0)$ depends implicitly on the fee. This result holds independent of the interval size L if g is μ -stable for all $\Delta \geq 0$. Note that, if μ is very small (*i.e.*, the CFMM has low curvature) then E_V^* is large; similarly, given two CFMMs, one with lower and one with higher curvature, α , the edge, needs to be larger in the CFMM with higher curvature to achieve the same lower bound for the payoff.

Liquidity provider loss. We can similarly get a lower bound on the expected loss of an LP since it is equal to $-E_V(\Delta)$:

$$-E_V(\Delta) = \gamma \int_0^\Delta (g(0) - g(-\gamma t)) dt - \alpha(m_0 - m_1)\Delta,$$

but, if $\Delta \leq L$, we have that $g(0) - g(-\gamma t) \geq \gamma\kappa t$, and so

$$-E_V(\Delta) \geq \frac{1}{2}\kappa\gamma^2\Delta^2 - \alpha(m_0 - m_1)\Delta.$$

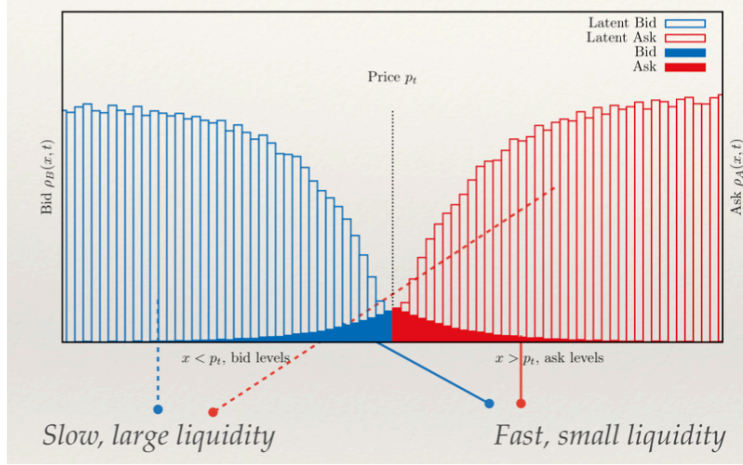


Figure 7: Order book shapes for different liquidity levels. Glosten and others have show that the concave shape is better for large liquidity and low trading velocity (*e.g.*, trades per second) whereas the convex shape is better for many small orders and high trading velocity. Figure from [37].

Minimizing both sides gives

$$-E_V^* \geq -\frac{\alpha^2(m_0 - m_1)^2}{2\kappa\gamma^2} \quad (8)$$

whenever $\alpha(m_0 - m_1) \leq L\kappa\gamma^2$ (*i.e.*, when the unconstrained minimum lies in the interior of the interval $[0, L]$) and is otherwise bounded by $-E_V^* \geq \kappa\gamma^2 L^2/2 - \alpha(m_0 - m_1)L$. We will mostly consider the first case in the following discussion, since we can often expand the interval L to be large enough to contain this bound. (Note also that (8) is equivalent to giving an upper bound on the expected value of an informed trader.)

Discussion. This matches the empirical observation that lower curvature CFMMs tend to have higher liquidity for assets that do not require much information to trade whereas LPs of higher-curvature CFMMs lose less to informed traders. This result illustrates that unlike common wisdom in the CFMM design space, one need not only have an optimal fee to maximize LP returns, but one needs to adjust the curvature as well.

Moreover, this result represents an analogue of classical microstructure results that show that the shape of an order book gives bounds on adverse selection. Glosten [27] showed that when you consider market makers who have to quote prices on multiple markets, then the shape of the order book impacts how liquidity changes in response to adverse selection. In figure 7, we see two different shapes for an order book, one approximately concave (the unshaded bars) and one convex (the filled in region). When a market maker observes or realizes adverse selection costs, they make a market more illiquid (*e.g.*, by canceling orders) to force active traders to pay a higher impact cost to market makers. This leads to the higher curvature, concave shape seen in figure 7. Equation (8) then suggests that market makers in CFMMs can replicate the same effect by increasing the curvature, therefore increasing the curvature lower bound, κ .

3 Price stability and yield farming

In this section, we will describe price stability when arbitrageurs trade between two markets, each with different curvature bounds. This stability result provides a quantitative explanation for the stability phenomenon of figure 2. Given that the empirically observed sUSD price instability was due to liquidity incentives (yield farming), it is natural to expect that there is a relationship between curvature and the precise costs of a stability incentive. In other words, the question we seek to answer is: how much do we need to pay LPs for providing liquidity? We describe this connection precisely in §3.3, which shows that there is an optimal liquidity subsidy for a market that interacts with an external market with different curvature.

3.1 Model description

Here, we will define the market model used throughout the remainder of this section. In our model, we have two available markets: the external market (whose price fluctuates due to extrinsic demand) and the secondary market (which we will assume is a CFMM, though the model holds more generally), along with an arbitrageur agent which seeks to maximize their profit by exploiting the difference in price between these two markets.

Market model. We describe a relatively simple, but very general, model of the *external market* and how it interacts with the given CFMM. In particular, the external market reports some strictly positive price $m_0 \in \mathbf{R}_{++}$ at the start of the round. The basic model of interactions between the markets and the arbitrageur proceeds as follows:

1. At the round start, the quoted external market price is m_0^e , while the secondary market price is m_0^s .
2. An arbitrageur then trades with the external market and the secondary market (which will usually be a CFMM). This results in a new external *and* secondary market price m_a which are equal since no-arbitrage has been enforced.
3. The external market price then changes from the no-arbitrage price m_a to a new price by some process modeling external influences. Step 1 is repeated with the new external and secondary market prices.

In fact, in our presentation, we will not assume anything about the dependence of the new price on m_a or even on m_0^e or m_0^s , which means that the results here hold for essentially all (reasonable) models of exogenous price changes for the external market price.

Main goal. We will show that, even when the external market price m_0^e differs widely from the secondary market price m_0^s , the (new) arbitrated market price m_a does not differ too much from the previous secondary market price m_0^s . Written out, we wish to find conditions such that, even when the price difference between both markets before no-arbitrage, $m_0^e - m_0^s$,

is large, the difference between the no-arbitrage price and the secondary market's price $m_a - m_0^s$ is small, in a precise sense. This would imply that even though the external market price deviates from the secondary market price $m_0^e - m_0^s$, the secondary market is able to force the new no-arbitrage price, m_a , back to a price that was close to its previous value, m_0^s .

Assumptions. In this set up, the external market will have a price impact function f which is, as before, continuous and nondecreasing. We will define the initial price of the external market as $m_0^e = f(0)$. Additionally, we will assume the external market is κ -liquid, with a slightly different definition than the one given in §1.1: we will say an external market is κ -liquid if it satisfies, for $\Delta \geq 0$,

$$f(\Delta) - f(0) \geq \kappa \Delta.$$

This differs from the original definition given in §1.1 because, here, Δ is the amount *purchased* from the external market, rather than the amount sold to it. In this case, if f is a differentiable convex function, we have that $\kappa = f'(0)$ is the tightest possible constant κ satisfying this condition and is a global bound (holding for all $\Delta \geq 0$) which follows immediately from the first-order convexity conditions.

As before, we will simply assume that the secondary market, with continuous, nondecreasing price impact function g , is μ -stable with the usual definition given in §1.1. We will similarly define $m_0^s = g(0)$.

3.2 Stability and Curvature

In this section, we derive the main result for general, continuous price impact functions, satisfying the conditions outlined previously.

Main result. Assume the price impact function of the primary market, f , is κ -liquid (in the sense above) and that the price price impact function of the secondary market, g , is μ -stable. We will show that the secondary-market's price change is bounded in the following way:

$$m_0^s - m_a \leq \frac{\mu}{\kappa}(m_0^s - m_0^e). \quad (9)$$

Note that because both sides of the inequality are nonnegative, inequality (9) can also be written as

$$|m_0^s - m_a| \leq \frac{\mu}{\kappa}|m_0^s - m_0^e|.$$

In other words, the no-arbitrage price change is at most a factor of μ/κ from the difference between the primary and secondary markets. This quantity (and therefore the price change after arbitrage) is small whenever the secondary market is very liquid (μ is small), or when the external market is very illiquid (κ is large). While apparently simple, we show in §3.3 that this result can be applied to many useful circumstances.

Proof of main result. By assumption, we have

$$f(0) = m_0^e \leq m_0^s = g(0),$$

so $f(0) \leq g(0)$. Then, if we can find any $\Delta \geq 0$ that satisfies

$$f(\Delta) \geq g(-\Delta), \tag{10}$$

then there exists some $0 \leq \Delta^* \leq \Delta$ such that

$$f(\Delta^*) = g(-\Delta^*),$$

by continuity and monotonicity of f and g . As before, the no arbitrage price m_a is defined as $m_a = f(\Delta^*) = g(-\Delta^*)$.

To show that there exists a Δ satisfying (10), note that any $\Delta \geq 0$ which satisfies

$$m_0^e + \kappa\Delta \geq m_0^s, \tag{11}$$

automatically satisfies (10) since

$$f(\Delta) \geq f(0) + \kappa\Delta = m_0^e + \kappa\Delta \geq m_0^s = g(0) \geq g(-\Delta)$$

where the first inequality follows from (3), the second inequality follows from (11), while the last inequality follows from the fact that g is a nondecreasing function.

In order to satisfy (11), we can easily choose

$$\Delta = \frac{m_0^s - m_0^e}{\kappa}.$$

Such a Δ will then satisfy (10). This, in turn, implies that a no-arbitrage trade Δ^* satisfies $0 \leq \Delta^* \leq \Delta$, and

$$m_0^s - m_a = g(0) - g(-\Delta^*) \leq g(0) - g(-\Delta) \leq \mu\Delta = \frac{\mu}{\kappa}(m_0^s - m_0^e).$$

Here, the first equality follows from the definition of m_0^s and g , while the first inequality follows from the monotonicity of g and the second inequality follows from (2). The resulting inequality is the one given in (9).

Assumptions. While the assumption that $m_0^e \leq m_0^s$ might appear restrictive, it is actually fully general. For example, if $f(\Delta)$ and $g(-\Delta)$ specify the price of an asset A with respect to a tradeable asset B , after buying Δ amount of asset A from the primary (or secondary) market, then the amount of asset B traded with each market is given by the quantity functions

$$p(\Delta) = \int_0^\Delta f(t) dt, \quad q(-\Delta) = \int_0^\Delta g(-t) dt.$$

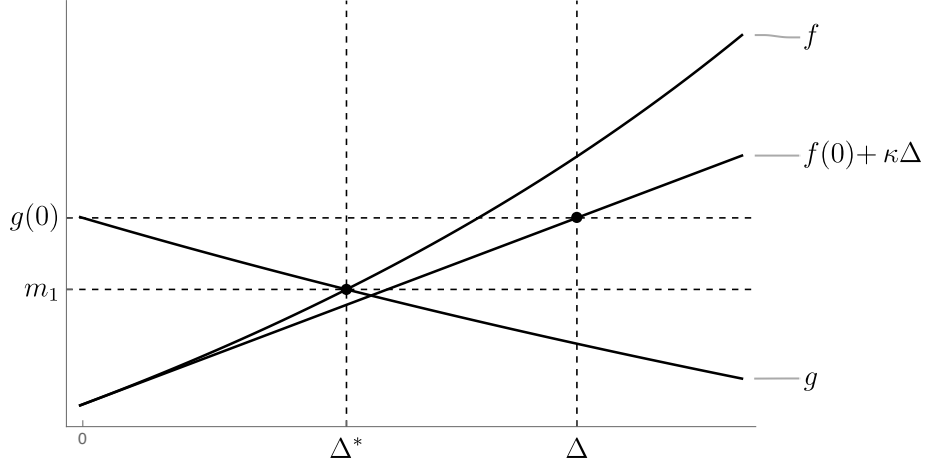


Figure 8: Visual proof of the main result. Given $g(0) > f(0)$ and f and g monotonic and continuous, we find Δ such that $f(\Delta) \geq g(-\Delta)$. The fact that there exists $0 \leq \Delta^* \leq \Delta$ such that $f(\Delta^*) = g(-\Delta^*)$ is then obvious from the figure. We also note that $\Delta = \frac{g(0)-f(0)}{\kappa}$ will be a worse overestimator for Δ^* if the curvature of the two impact functions is large.

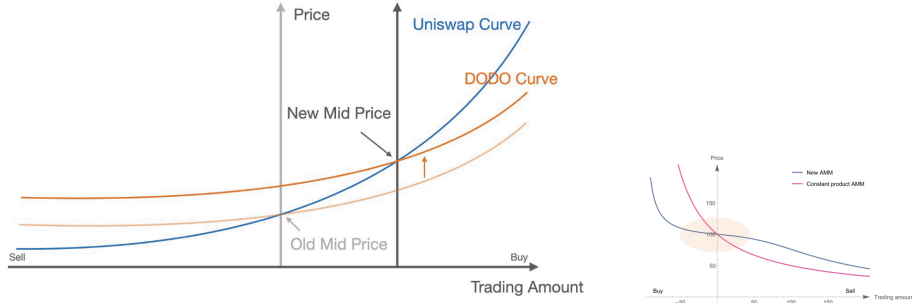


Figure 9: Marginal price curves for Dodo [9] (left) and McDEX [35]. Note that Dodo's curve flattens in response to a trade, whereas McDEX's curve has asymmetric slippage. Both of these curves aim to provide leverage, so they have higher curvature when the protocol takes more risk (*e.g.*, McDEX when the protocol's net position is negative). For the McDEX curve, we can view this as $\mu^- > \mu^+$ for some interval $[-50, 50]$

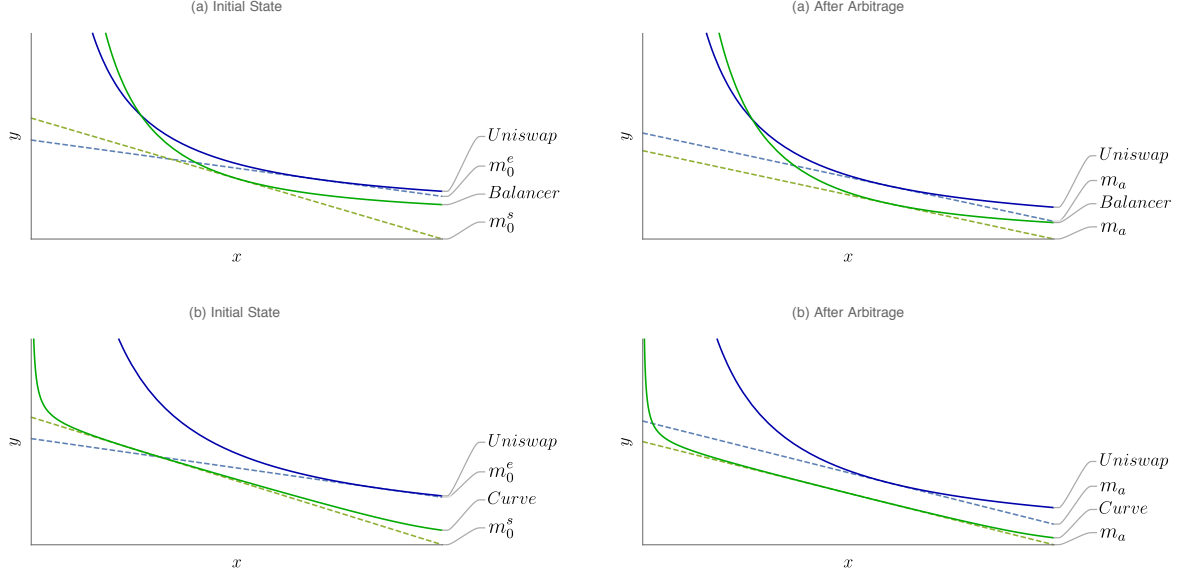


Figure 10: Visual illustration of arbitrage between two CFMMs with different curvature. Note the changes in slope, corresponding to the changes in market price after no-arbitrage is enforced.

On the other hand, we may ask what the price of asset B is with respect to asset A after buying some amount Δ' of asset B . The quantity of asset A received for Δ' of asset B is easily seen to be $p^{-1}(\Delta')$ for the primary market and $q^{-1}(\Delta')$ for the secondary market. Both exist because f and g are strictly positive which implies that p and q are strictly monotonic. This implies that the respective price (using implicit differentiation) is given by

$$(p^{-1}(\Delta'))' = \frac{1}{p'(p^{-1}(\Delta'))} = \frac{1}{f(p^{-1}(\Delta'))} = \frac{1}{f(\Delta)},$$

where we have defined $\Delta = p^{-1}(\Delta')$, and similarly for q . So, if $m_0^e \geq m_0^s$, we may always ‘swap’ asset A for asset B in this sense, such that the resulting marginal prices are given by $1/m_0^e \leq 1/m_0^s$, and enforce no-arbitrage conditions over coin B , instead.

Extensions. As with §1.1, may not be the case that constants μ or κ exist for trades of all possible sizes. In general, such constants do exist for trades of bounded size, say $0 \leq \Delta \leq L$. In this case, the main result extends immediately in the following way: for any prices satisfying

$$m_0^s - m_0^e \leq \kappa L, \tag{12}$$

we have

$$m_0^s - m_a \leq \frac{\mu}{\kappa}(m_0^s - m_0^e).$$

The proof of this statement is identical to the one above, with the additional condition that the primary and secondary market prices differ by no more than κL .

3.3 Yield farming subsidy

One of the main drivers of the growth in CFMM usage in 2020 was *yield farming*. Yield farming, which is similar to maker–taker rebates in traditional trading [11], involves subsidizing the provision of liquidity for a new issued crypto asset. Suppose that some asset is issued at time t_0 by a smart contract and that it has an inflation schedule $i_t \in \mathbf{R}_+$, where i_t is the number of units of X produced at time t . In order to incentivize liquidity between the new asset and a numéraire, the smart contract reserves some percent (say, ℓ_t) of inflation for liquidity provision. If a user creates LP shares for a CFMM, which trades this coin and then stakes (or, in other words, locks) these shares into a smart contract, then they receive some amount of this coin from the smart contract for providing liquidity. For instance, if there are 1000 LP shares locked into the contract and a single user created 100 of these locked shares, they might receive $\frac{100}{1000}i_t\ell_t$ units of the new asset at time t . By subsidizing liquidity, the smart contract issuing the new coin can ensure that users can trade the new asset while also ensuring that liquidity providers have lower losses. The main loss that CFMM liquidity providers face is ‘impermanent loss’ or losses due to the concavity of the portfolio value of an LP share; see, *e.g.*, [18] for the specific case of Uniswap. One can also directly show these losses occur by using the definition of the portfolio value in [4].

Sufficient subsidy. A protocol designer would then want to ensure that LPs are compensated enough, say, by some amount R_ℓ in the traded coin, to have nonnegative profit after a no-arbitrage trade with the external market; *i.e.*, we need to guarantee that the portfolio value of an LP, after arbitrage and subsidy, is nonnegative. To do this, note that the opportunity cost or ‘impermanent loss’ of being an LP in the secondary market is given by (in a similar way to §2.1):

$$m_a\Delta - \int_0^{-\Delta} g(-t) dt \geq (m_a - g(0))\Delta = (m_a - m_0^s)\Delta,$$

where we have used the fact that g is nondecreasing, and $g(0) = m_0^s$ by definition. Applying (9) gives

$$(m_a - m_0^s)\Delta \geq -\frac{\mu}{\kappa}(m_0^s - m_0^e). \quad (13)$$

Therefore, to incentivize LPs to continue adding liquidity to the μ -stable secondary market, assuming an external market that is κ -liquid, it suffices to subsidize the LPs by some amount

$$R'_\ell = \frac{\mu}{\kappa}(m_0^s - m_0^e),$$

in the numéraire. Alternatively, they can be subsidized by at least

$$\frac{R'_\ell}{m_a} = \frac{\mu}{\kappa m_a}(m_0^s - m_0^e) \leq \frac{\mu}{\kappa} \left(1 - \frac{m_0^e}{m_0^s}\right)$$

in the traded coin, since $m_0^e \leq m_a \leq m_0^s$. In other words, the total quantity of subsidy is proportional to the curvature of the primary market, and inversely proportional to the

curvature of the secondary market. If we define $h = m_0^e/m_0^s$ to be the percentage growth of the asset (note that $h \leq 1$ since $m_0^e \leq m_0^s$) then we have simple expression for the amount of subsidy that is sufficient, in the traded asset, which we will define as:

$$R_\ell = \frac{\mu}{\kappa}(1 - h). \quad (14)$$

Discussion. This gives a simple condition which guarantees that liquidity providers have nonnegative returns for providing liquidity. In particular, we note that more subsidy has to be provided as h becomes small (*i.e.*, the price is changing with large drift) or when μ/κ is large (*i.e.*, the secondary market, for which the LPs are providing liquidity for, is very illiquid when compared to the external market). In general, this means that how much subsidy one might need to provide depends not just on the drift of the asset, but also the relative curvature of the two markets.

4 Dynamic hedging of CFMMs

We describe how dynamic hedging quantities for contingent claims (*e.g.*, an asset’s Delta and Gamma) are constructed for CFMMs, providing a means for LPs to hedge risk such as so-called ‘impermanent loss.’ In order to simplify notation, we will refer to the Delta and Gamma of a portfolio as P_Δ and P_Γ , respectively. To connect portfolio value to curvature, we will show that analogues of dynamic hedging Greeks are closely related to curvature. These quantities are important for liquidity providers as they represent the net exposure to the underlying collateral the LPs have, as well as how to hedge loss due to drift (known colloquially in decentralized finance as ‘impermanent loss’). While Uniswap can be statically replicated by a portfolio of options [18], it is unclear precisely how to do this for generic CFMMs. We will compute dynamic hedging quantities for LPs here and show how they behave under trades. This behavior, which will be connected to curvature, represents extra convexity in the replicating portfolio that is needed to compensate for the LP share of a high curvature CFMM.

Approximate hedges. We apply these hedging results to try to construct approximate hedges for impermanent loss. If the price of the traded coin increases while volatility is mollified, the LP realizes impermanent loss [5, 51]. As the price drifts up, if the LP is able to increasingly sell put options for the tradable asset, then they can lower their realized impermanent loss when they exit the pool. The question then is: how many put options does an LP need to sell as a function of the price upon their entry into the pool?

We illustrate that if options on realized impact costs exist, then an LP can enter a short put option portfolio on impact costs to hedge their impermanent loss. Such options correspond to options on the change in the market price when a trade of size Δ is made. Owners of such an option effectively have insurance on price changes in the underlying due to a single large trade.

Delta hedging in practice. While unnatural relative to conventional options on stocks, options on impact cost exist in traditional finance when considering American Depositary Receipts (ADRs) in the US equity market. Suppose that there is a stock that exists both in a non-US market and on a US exchange as an ADR. ADRs allow for the non-US stock to be traded in the US via a synthetic asset that has a creation and redemption mechanism. When the foreign exchange rate between the non-US currency and US dollars drifts wildly, there is an impact of foreign exchange rate on option pricing. The difference in price between options on the non-US stock and options on US ADRs replicates an option on the impact cost of the foreign exchange rate on the stock price [41, 42, 3].

4.1 Computing the Greeks

The portfolio value of the reserves, as before, is defined as

$$P_V = mR + R',$$

where R and R' depend implicitly on the price m and the value of $\psi(0, 0)$. We then have:

$$P_\Delta = \frac{dP_V}{dm} = R, \quad P_\Gamma = \frac{d^2 P_V}{dm^2} = \frac{dR}{dm}, \quad (15)$$

which we show below. (Here P_Δ and P_Γ represent the corresponding Greeks of P_V .)

Computation. Because it is often the case that the function ψ is of the following form (when there are no fees)

$$\psi(R, R', \Delta, \Delta') = \Psi(R - \Delta, R' + \Delta'),$$

for some function $\Psi : \mathbf{R}^2 \rightarrow \mathbf{R}$ (as is the case with all examples presented in §1.1), it suffices to consider the function Ψ in terms of only the reserves R and R' . We will assume this is true in the following derivation.

As before, the portfolio value is given by

$$P_V = mR + R',$$

where m is the market price, which, by no arbitrage, must satisfy (at the reserve values R and R' , after no arbitrage)

$$m = \frac{\partial_1 \Psi(R, R')}{\partial_2 \Psi(R, R')}.$$

(Note that there is no negative sign here, due to the definition of Ψ .) By implicitly differentiating $\Psi(R, R')$ over any level set, we have that

$$\frac{d\Psi(R, R')}{dm} = (\partial_1 \Psi(R, R')) \frac{dR}{dm} + (\partial_2 \Psi(R, R')) \frac{dR'}{dm} = 0,$$

so,

$$m \frac{dR}{dm} = - \frac{dR'}{dm}.$$

Using this final condition:

$$P_\Delta = \frac{dP_V}{dm} = R + m \frac{dR}{dm} + \frac{dR'}{dm} = R,$$

as required, while the expression for P_Γ follows from the fact that

$$P_\Gamma = \frac{dP_\Delta}{dm} = \frac{dR}{dm}.$$

Approximate hedging. On the other hand, the deterministic price schedule of CFMMs also implies that directly hedging trade quantities is equivalent to hedging prices if the CFMM is the primary market. In the case that g is convex, which we assume here, we can give bounds that can be easily replicated with options. We note that this result only holds for decreasing prices, due to the assumptions on μ and κ , but we suspect there exist more general bounds, holding over any change in price, which are very similar in spirit.

Suppose that an LP has reserves R and R' in a CFMM and the CFMM is the primary market. First, note that,

$$R = P_\Delta = \frac{dP_V}{dm} = \frac{dP_V}{d\Delta} \frac{d\Delta}{dm}.$$

After a trade of size Δ , we have

$$P_V(-\Delta) = g(-\Delta)(R + \Delta) + (R' - \Delta'),$$

so that if g is differentiable, we have

$$\frac{dP_V}{d\Delta} = g'(-\Delta)(R + \Delta) + g(\Delta) - \frac{d\Delta'}{d\Delta} = g'(-\Delta)(R + \Delta), \quad (16)$$

where we used (4) in the last equality. Therefore, if our CFMM is μ -stable in the sense of (2), then, because g is convex, nondecreasing, we have that $g'(-\Delta) \leq \mu$, so

$$\frac{dP_V}{d\Delta} = g'(-\Delta)(R + \Delta) \leq \mu(R + \Delta).$$

Finally, because $\frac{d\Delta}{dm} \leq 0$, then we have that

$$P_\Delta = \frac{dP_V}{d\Delta} \frac{d\Delta}{dm} \geq \mu(R + \Delta) \frac{d\Delta}{dm}.$$

If the CFMM is both μ -stable and κ -liquid, then we have the following expansion:

$$\mu(R + \Delta) \frac{d\Delta}{dm} \leq P_\Delta \leq \kappa(R + \Delta) \frac{d\Delta}{dm}. \quad (17)$$

These inequalities show that the curvature constants provide means for super and subhedging of LP share risk. In particular, these linear bounds allow for an LP to hedge their risk using simpler instruments, as described in the sequel. Finally, we note that using these calculations, one can show that hedging a portfolio of n stablecoins is approximately equivalent to hedging two stablecoins via an n -dimensional generalization presented in Appendix E.

5 Conclusion and future work

In this paper, we explored how the shape of a constant function market maker affects its ability to serve as the primary market for digital assets. To do this, we defined a notion of curvature for a given market and then showed its implications for liquidity providers in the market, showing that this notion of curvature is very closely related to the intuitive notion of curvature in the case of constant function market makers.

Liquidity provider returns. There has been much empirical evidence suggesting that the return profile of a CFMM liquidity provider depends on the shape of the CFMM’s trading function. We studied the LP returns under three different scenarios: uninformed trading, informed trading, and yield farming. Some of the results presented take inspiration from the traditional market microstructure literature and consider the return profile of informed traders. These traders can be viewed as bringing information to the market by placing a Kelly-style bet on the next price update that will take place in the CFMM. We were able to use no-arbitrage to derive lower bounds on the liquidity provider loss (and, conversely, lower bounds on the edge trader’s expected profit) akin to that of Glosten and Milgrom which connect curvature to the informed trader’s informational edge. The results from this section provide a simple economic interpretation of curvature as the amount of information an informed trader needs to achieve a certain profit (given a fixed edge and market price).

Yield farming. We then extended the definition to also include interacting markets with finite liquidity. These notions of shape or curvature could then be used to capture how a single trade on a market with finite liquidity affects prices on another market with finite liquidity. Using these definitions for curvature, we were able to bound the tracking error when an arbitrageur trades between a pair of markets with different curvatures. When specialized to CFMMs, the results presented generalize some of the results of [5, 4] to the case where the market is not infinitely liquid.

We then used this to analyze the yield farming phenomenon, where protocols began providing subsidies to liquidity providers. Here, we showed a lower bound to the amount of subsidy needed to pay liquidity providers to account for their ‘impermanent loss’ when compared to a market with bounded liquidity, which depended on the curvature of both markets and the rate of growth of the asset. Combined, these results suggest that the curvature of a CFMM needs to be optimized to avoid adverse selection while also capturing trading volume and fees related to asset price growth.

Future work. This work can be extended in a number of ways. On the practical side, many of the results presented here only work in two dimensions (*e.g.*, two asset trading). Generalizing our results to n dimensions would be a useful but potentially difficult problem. For example, it is not clear how to define μ -stability in higher dimensions for general CFMMs, without being overly restrictive. Moreover, even though we give some sufficient conditions on the curvature of a ‘good’ CFMM for certain applications, it is still an open question for how

to take a given price process, represented, say, as an Itô process or a jump-diffusion process, and then construct a ‘good’ CFMM. If this were found, then one could take historical data for a crypto asset and construct an optimized CFMM for trading this asset. Finally, it is clear that dynamic CFMMs [25, 54, 10, 46, 52], *i.e.*, CFMMs whose trading functions vary in time according to either a stochastic or control mechanism, continuously affect the curvature of the trading function. Given the results of this paper, a natural extension to inquire about is: how should one design an optimal control mechanism to replicate a desired payoff or behavior? The results of this paper suggest that the trade-off between adverse selection and payoff growth are extremely important to such designs, especially for products with sharp payoffs or time decay (*e.g.*, barrier options).

The results of §3.3 intimate that there is such a mapping, akin to super replication results from traditional mathematical finance. Curiously, results from the mathematical finance literature also heavily use convex duality and it is likely that there are a number of fruitful translational results can be found. These results will likely give more power to the dynamic hedging results of §4 and will be useful for comparing CFMM payoffs to traditional derivatives pricing. We suspect that some of our conjectures in appendix I, regarding the superhedging of contingent claims specified by CFMM portfolio values, is likely to be a problem with deep connections to such results.

Acknowledgments

The authors would like to thank Yi Sun, Ciamac Moallemi, Hsien-Tang Kao, Victor Xu, Rei Chiang, and Adam Lerer for helpful comments and suggestions. We also would like to thank Matteo Liebowitz of Uniswap for providing data.

References

- [1] *Perpetual protocol documentation*, <https://docs.perp.fi>.
- [2] Hayden Adams, *Uniswap birthday blog - v0*, Nov 2019.
- [3] Jasper Anderluh and Hans van der Weide, *Adr option trading*, https://www.swi-wiskunde.nl/swi2004/wp-content/uploads/sites/17/2018/08/AOT_scproc_2004.pdf, p. 1.
- [4] Guillermo Angeris and Tarun Chitra, *Improved price oracles: Constant function market makers*, Proceedings of the 2nd ACM Conference on Advances in Financial Technologies (New York, NY, USA), AFT ’20, Association for Computing Machinery, 2020, p. 80–91.
- [5] Guillermo Angeris, Hsien-Tang Kao, Rei Chiang, Charlie Noyes, and Tarun Chitra, *An analysis of Uniswap markets*, Cryptoeconomic Systems (to appear).
- [6] Jun Aoyagi, *Lazy liquidity in automated market making*, Available at SSRN 3674178 (2020).

- [7] Jean-Philippe Bouchaud, Yuval Gefen, Marc Potters, and Matthieu Wyart, *Fluctuations and response in financial markets: the subtle nature of ‘random’ price changes*, Quantitative finance **4** (2004), no. 2, 176–190.
- [8] Stephen P. Boyd and Lieven Vandenbergh, *Convex optimization*, Cambridge University Press, Cambridge, UK ; New York, 2004.
- [9] Dodo Breeder, *Introducing dodo: 10x better liquidity than uniswap*, <https://dodo-in-the-zoo.medium.com/introducing-dodo-10x-better-liquidity-than-uniswap-852ce2137c57>, June 2020.
- [10] DODO Breeder, *Why there is no impermanent loss on dodo*, <https://medium.com/dodoex/why-there-is-no-impermanent-loss-on-dodo-4d82b2b532bc>, Sep 2020.
- [11] Michael Brolley and Katya Malinova, *Informed trading and maker-taker fees in a low-latency limit order market*, Available at SSRN 2178102 (2013).
- [12] Enzo Busseti, Ernest K Ryu, and Stephen Boyd, *Risk-constrained kelly gambling*, The Journal of Investing **25** (2016), no. 3, 118–134.
- [13] Peter Carr and Dilip Madan, *Towards a theory of volatility trading*, Option Pricing, Interest Rates and Risk Management, Handbooks in Mathematical Finance (2001), 458–476.
- [14] Peter Carr and Qiji Jim Zhu, *Convex duality and financial mathematics*, Springer, 2018.
- [15] Larry Cermak, *Uniswap’s monthly trade volume exceeded coinbase’s in september*, <https://www.theblockcrypto.com/linkedin/79775/uniswap-coinbase-monthly-volume-september>, October 2020.
- [16] George HG Chen and R Tyrrell Rockafellar, *Convergence rates in forward–backward splitting*, SIAM Journal on Optimization **7** (1997), no. 2, 421–444.
- [17] Patrick Cheridito, Michael Kupper, and Ludovic Tangpi, *Duality formulas for robust pricing and hedging in discrete time*, SIAM Journal on Financial Mathematics **8** (2017), no. 1, 738–765.
- [18] Joseph Clark, *The replicating portfolio of a constant product market*, Available at SSRN 3550601 (2020).
- [19] Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels, *Flash boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability*, 2020 IEEE Symposium on Security and Privacy (SP), 2020, pp. 566–583.

- [20] Brady Dale, *What is yield farming?*, <https://www.coindesk.com/defi-yield-farming-comp-token-explained>, July 2020.
- [21] Michael Egorov, *StableSwap - efficient mechanism for Stablecoin liquidity*, 6 (en).
- [22] ———, *Stableswap — efficient mechanism for stablecoin liquidity*, unpublished, Nov 2019.
- [23] ———, *Curve*, Jan 2020.
- [24] ———, *Vulnerability disclosure: the discovery and the rescue*, Jan 2020.
- [25] Alex Evans, *Liquidity provider returns in geometric mean markets*, arXiv preprint arXiv:2006.08806 (2020).
- [26] Jim Gatheral, *No-dynamic-arbitrage and market impact*, Quantitative finance **10** (2010), no. 7, 749–759.
- [27] Lawrence R Glosten, *Is the electronic open limit order book inevitable?*, The Journal of Finance **49** (1994), no. 4, 1127–1161.
- [28] Lawrence R Glosten and Paul R Milgrom, *Bid, ask and transaction prices in a specialist market with heterogeneously informed traders*, Journal of financial economics **14** (1985), no. 1, 71–100.
- [29] Omkar Godbole, *Defi flipping comes to exchanges as uniswap topples coinbase in trading volume*, <https://www.coindesk.com/defi-flipping-uniswap-topples-coinbase-trading-volume>, September 2020.
- [30] Ron Goldman, *Curvature formulas for implicit curves and surfaces*, Computer Aided Geometric Design **22** (2005), no. 7, 632–658.
- [31] Wanyun Catherine Gu, Anika Raghuvanshi, and Dan Boneh, *Empirical measurements on pricing oracles and decentralized governance for stablecoins*, Available at SSRN 3611231 (2020).
- [32] Robin Hanson, *Combinatorial information market design*, Information Systems Frontiers **5** (2003), no. 1, 107–119.
- [33] jakub, *Is yield farming dead?*, <https://finematics.com/is-yield-farming-dead/>, October 2020.
- [34] John Jenq and Priscilla Jenq, *Order book data structures*, Proceedings of the International Conference on Foundations of Computer Science (FCS), The Steering Committee of The World Congress in Computer Science, Computer ..., 2018, pp. 49–55.
- [35] Liu Jie, *Mcdex white paper*, <https://mcdex.io/references/#/en-US/white-paper>.

- [36] Albert S Kyle, *Continuous auctions and insider trading*, Econometrica: Journal of the Econometric Society (1985), 1315–1335.
- [37] Ismael Lemhadri, *Market impact in a latent order book*, arXiv preprint arXiv:1802.06101 (2018).
- [38] Tianyi Lin, Chi Jin, and Michael I Jordan, *On gradient descent ascent for nonconvex-concave minimax problems*, arXiv preprint arXiv:1906.00331 (2019).
- [39] Alan Lu, *Gnosis conditional token automated market maker*, <https://github.com/gnosis/conditional-tokens-market-makers>, March 2020.
- [40] Richard K Lyons and Ganesh Viswanath-Natraj, *What keeps stablecoins stable?*, Tech. report, National Bureau of Economic Research, 2020.
- [41] Dilip B Madan, *Jointly modeling american depository receipts, the local stock and the local price of the us dollar*, Available at SSRN 2234785 (2012).
- [42] Dominik Maltritz and Stefan Eichler, *Currency crisis prediction using adr market data: An options-based approach*, International Journal of Forecasting **26** (2010), no. 4, 858–884.
- [43] Fernando Martinelli and Nikolai Mushegian, *Balancer: A non-custodial portfolio manager, liquidity provider, and price sensor*, (2019).
- [44] Saniya More, *Sushiswap initiates migration of over \$800m in liquidity funds from uniswap*, <https://www.theblockcrypto.com/post/77324/sushiswap-initiates-migration-of-over-800m-in-liquidity-funds-from-uniswap>, August 2020.
- [45] Justin J Moses, *Contract 0xdc6a51ea3ca5d3fd898fd6564757c7aaec3ca92 (curve rewards)*, <https://etherscan.io/address/0xdc6a51ea3ca5d3fd898fd6564757c7aaec3ca92#code>, October 2020.
- [46] Allan Niemerg, Dan Robinson, and Lev Livnev, *Yieldspace: An automated liquidity provider for fixed yield tokens*, (2020).
- [47] Abraham Othman, David M Pennock, Daniel M Reeves, and Tuomas Sandholm, *A practical liquidity-sensitive automated market maker*, ACM Transactions on Economics and Computation (TEAC) **1** (2013), no. 3, 1–25.
- [48] Abraham Othman and Tuomas Sandholm, *Automated market-making in the large: the Gates Hillman prediction market*, Proceedings of the 11th ACM conference on Electronic commerce, 2010, pp. 367–376.
- [49] Manfred Schäl, *Martingale measures and hedging for discrete-time financial markets*, Mathematics of operations research **24** (1999), no. 2, 509–528.

- [50] Theodore Shifrin, *Differential geometry: a first course in curves and surfaces*, <http://www.dma.unifi.it/~verdiani/penn2016/ShifrinDiffGeo.pdf>, 2015.
- [51] Martin Tassy and David White, *Growth rate of a liquidity provider's wealth in $xy = c$ automated market makers*, https://math.dartmouth.edu/~mtassy/articles/AMM_returns.pdf, Nov 2020.
- [52] Teddy Woodward, *Notional finance amm*, <https://docs.notional.finance/traders/advanced-topics/notional-amm>.
- [53] Peter Zeitz, *Curve vulnerability report*, https://medium.com/@peter_4205/curve-vulnerability-report-a1d7630140ec, Oct 2020.
- [54] Victor Zhang, *Cofix, the future of on-chain market making*, <https://medium.com/cofix/cofix-8b904eea507e>, Oct 2020.

A Form equivalence

In [4], the trading function is defined as a function $\varphi : \mathbf{R}_+^n \times \mathbf{R}_+^n \times \mathbf{R}_+^n \rightarrow \mathbf{R}$, which maps the reserves, *input trades*, and *output trades* to a real value. In this paper, we do not make an explicit distinction between the input and output trades; these are, instead, specified by the sign of the trade amounts Δ and Δ' . In this case, we can make the following equivalences between the trade (Δ, Δ') (the notation as used in this paper) and the input trade $\Delta^0 \in \mathbf{R}_+^2$ and output trade $\Lambda^0 \in \mathbf{R}_+^2$ as used in [4], for $n = 2$:

$$\Delta^0 = ((-\Delta)_+, (\Delta')_+), \quad \Lambda^0 = ((\Delta)_+, (-\Delta')_+),$$

while the reserves are simply $R^0 = (R, R')$. We then have

$$\varphi(R^0, \Delta^0, \Lambda^0) = \psi(R, R', \Delta, \Delta'),$$

as expected. The update equations remain identical, since $R^0 \leftarrow R^0 + \Delta - \Lambda$ is equivalent to $R \leftarrow R - \Delta$ and $R' \leftarrow R' + \Delta'$, which means that all results from [4] hold as stated.

B Lower bounds for portfolio value with fees

The analysis of CFMMs with fees is often much harder than the analysis of fee-less CFMMs. This construction gives a simple lower bound which shows that it often suffices to consider a fee-less CFMM with fees taken separately; *i.e.*, it often suffices to consider a CFMM where the fee is not reinvested into the reserves, but is instead given to LPs directly.

Statement. For simplicity, we will use the notation from [4], which results in a simple proof for any number of coins n . (See appendix A.)

Let $\varphi : \mathbf{R}_+^n \times \mathbf{R}_+^n \times \mathbf{R}_+^n \rightarrow \mathbf{R}$ be a trading function for a CFMM that can be written as (with some slight abuse of notation):

$$\varphi(R^0, \Delta^0, \Lambda^0) = \varphi(R^0 + \gamma\Delta^0 - \Lambda^0),$$

where $R^0 \in \mathbf{R}_+^n$ are the reserves, $\Delta^0 \in \mathbf{R}_+^n$ is the input trade, and $\Lambda^0 \in \mathbf{R}_+^n$ is the output trade. We will assume that φ is increasing in its arguments, and let $0 < \gamma \leq 1$ such that $(1 - \gamma)$ is the fee taken for the CFMM. (See, *e.g.*, [4, §2.3.1].) In this case, we will consider the resulting portfolio value of an LP at some cost vector $c \in \mathbf{R}_+^n$. We will then show that the portfolio value of the LP after any feasible trade (Δ^0, Λ^0) is at least as large as the equivalent portfolio value at the previous reserves with an extra factor of $(1 - \gamma)c^T \Delta$.

Proof. The proof is nearly immediate. Let $R^1 = R^0 + \Delta^0 - \Lambda^0$ be the post-trade reserves and R^0 be the pre-trade reserves, then

$$c^T R^1 = c^T (R^0 + \Delta^0 - \Lambda^0) = c^T (R^0 + \gamma\Delta^0 - \Lambda^0) + (1 - \gamma)c^T \Delta^0 \geq c^T R^* + (1 - \gamma)c^T \Delta^0,$$

where R^* is the solution to the fee-less portfolio-value problem [4, §2.5],

$$p_{R^0}(c) = \inf_{\psi(R) \geq \psi(R^0)} c^T R, \quad (18)$$

with variable $R \in \mathbf{R}_+^n$. Note that the second inequality follows since, by definition, (Δ^0, Λ^0) is a feasible trade only when

$$\psi(R^0 + \gamma\Delta^0 - \Lambda^0) \geq \psi(R^0),$$

and so $R = R^0 + \gamma\Delta^0 - \Lambda^0$ is a feasible point for the portfolio-value problem (18). Repeatedly applying this statement to any number of feasible trades (Δ^k, Λ^k) yields the following lower bound for the portfolio value at time k :

$$c^T R^k \geq p_{R^0}(c) + (1 - \gamma) \sum_{k=1}^n c^T \Delta^k.$$

In many cases, because we are interested in finding a lower bound to the portfolio value of liquidity providers, it will often suffice to use this statement in order to achieve a reasonable lower bound. This allows us to side-step the potentially very complicated analysis of CFMMs with fees and the fees' interactions with the CFMM's reserves.

C Relationship between g and curvature of ψ

From equation (4), we can write,

$$g(\Delta) = \frac{\frac{\partial \psi}{\partial \Delta}}{\frac{\partial \psi}{\partial \Delta'}} = \frac{\frac{\partial \psi}{\partial \Delta'} \frac{\partial \Delta'}{\partial \Delta}}{\frac{\partial \psi}{\partial \Delta'}} = \frac{d\Delta'}{d\Delta}$$

Given a CFMM invariant function, given initial reserves, $\psi(\Delta, \Delta')$, we can write:

$$\frac{\partial \psi}{\partial \Delta} = \frac{\partial \psi}{\partial \Delta'} \frac{\partial \Delta'}{\partial \Delta} = \frac{\partial \psi}{\partial \Delta'} g(\Delta)$$

Therefore, $g(\Delta) = \left(\frac{\partial \psi}{\partial \Delta'}\right)^{-1} \frac{\partial \psi}{\partial \Delta}$. Thus μ -stability condition, for sufficiently smooth g , relies on the first derivative of g :

$$\begin{aligned} \frac{dg}{d\Delta} &= - \left(\frac{\partial \psi}{\partial \Delta'} \right)^{-2} \left(\frac{\partial \psi}{\partial \Delta \partial \Delta'} \right) \left(\frac{\partial \psi}{\partial \Delta} \right) + \left(\frac{\partial \psi}{\partial \Delta'} \right)^{-1} \frac{\partial^2 \psi}{\partial \Delta^2} \\ &= \left(\frac{\partial \psi}{\partial \Delta'} \right)^{-1} \left(g \frac{\partial \psi}{\partial \Delta \partial \Delta'} + \frac{\partial^2 \psi}{\partial \Delta^2} \right) \end{aligned}$$

From [30, Prop. 3.1], we see that for an implicit function $F : \mathbf{R}^2 \rightarrow \mathbf{R}$, the Gaussian curvature κ_F is defined as,

$$\kappa_F = \frac{(\partial_2 F)^2 (\partial_1^2 F) - 2(\partial_1 F)(\partial_2 F)(\partial_1 \partial_2 F) + (\partial_1 F)^2 (\partial_2^2 F)}{((\partial_1 F)^2 + (\partial_2 F)^2)^{3/2}} \quad (19)$$

where $\partial_i F$ is the i th partial derivative of F . Using implicit substitution (*e.g.*, writing $\Delta'(\Delta)$) and substituting it into eq. (19)), we can see that the two formulas are equivalent.

D Curve's price impact function is convex

In this section, we claim that if the following assumption on the reserve sizes holds, then Curve's impact function is convex:

$$R' - \Delta' > \frac{1}{R} \quad (20)$$

Recall that for Curve, the trading function is defined for $\alpha, \beta > 0$ as,

$$\psi(\Delta, \Delta') = \alpha(R - \Delta + R' + \Delta') + \frac{\beta}{(R - \Delta)(R' + \Delta')} - k$$

From [50, Prop. 3.8], a curve defined via a sufficiently smooth implicit function $F(x, y) = 0$ is convex if and only if its Gaussian curvature (*i.e.*, equation (19)) is non-negative. Therefore, the claim of convexity is equivalent to showing that $\kappa_\psi \geq 0$. Since the denominator of equation (19) is non-negative, we only need the numerator to be positive, *e.g.*,

$$(\partial_2 \psi)^2 (\partial_1^2 \psi^2) + (\partial_1 \psi)^2 (\partial_2^2 \psi)^2 \geq 2(\partial_1 \psi)(\partial_2 \psi)(\partial_1 \partial_2 \psi) \quad (21)$$

Using the definition of ψ , we have

$$\begin{aligned} \partial_1 \psi &= \alpha - \frac{\beta}{(R - \Delta)^2(R' + \Delta')} & \partial_2 \psi &= -\alpha + \frac{\beta}{(R - \Delta)(R' + \Delta')^2} \\ \partial_1^2 \psi &= \frac{2\beta}{(R - \Delta)^3(R' + \Delta')} & \partial_2^2 \psi &= \frac{2\beta}{(R - \Delta)(R' + \Delta')^3} \\ \partial_1 \partial_2 \psi &= \partial_2 \partial_1 \psi = \frac{-\beta}{(R - \Delta)^2(R' + \Delta')^2} \end{aligned}$$

Note that

$$\begin{aligned} (R - \Delta)^5(R' + \Delta')^5(\partial_1 \psi)^2(\partial_2^2 \psi) &= 2b(\alpha(R - \Delta)^2(R' + \Delta') - \beta)^2 \\ (R - \Delta)^5(R' + \Delta')^5(\partial_2 \psi)^2(\partial_1^2 \psi) &= 2b(\alpha(R - \Delta)(R' + \Delta')^2 - \beta)^2 \\ (R - \Delta)^5(R' + \Delta')^5(\partial_1 \psi)(\partial_2 \psi)(\partial_1 \partial_2 \psi) &= \beta(\alpha(R - \Delta)^2(R' + \Delta') - \beta)(\alpha(R - \Delta)(R' + \Delta')^2 - \beta) \end{aligned}$$

As $(R - \Delta)^5(R' + \Delta')^5 \geq 0$, combining results with (21) yields the positivity condition

$$\begin{aligned} 2((\alpha(R - \Delta)^2(R' + \Delta') - \beta)^2 + (\alpha(R - \Delta)(R' + \Delta')^2 - \beta)^2) \\ \geq \\ (\alpha(R - \Delta)^2(R' + \Delta') - \beta)(\alpha(R - \Delta)(R' + \Delta')^2 - \beta) \end{aligned}$$

Let $A = \alpha(R - \Delta)^2(R' + \Delta')$ and $B = \alpha(R - \Delta)(R' + \Delta')^2$. Then this condition can be rewritten as

$$2((A - b)^2 + (B - a)^2) \geq (A - b)(B - a)$$

Dividing through by $(A - \beta)(B - \alpha)$ gives the final condition

$$2 \left(\frac{A - \beta}{B - \alpha} + \frac{B - \alpha}{A - \beta} \right) \geq 1$$

Provided that $A > \beta$ and $B > \alpha$, then this condition is always true. The first condition is equivalent to $(R - \Delta)^2(R' + \Delta') > \frac{\alpha}{\beta}$ and the second condition is equivalent to $(R - \Delta)(R' + \Delta')^2 > \frac{\beta}{\alpha}$. Combining these gives the condition $(R - \Delta)(R' + \Delta') > 1$. Minorizing $R - \Delta$ by R is equivalent to the assumption (20), proving the claim.

E Portfolio Greeks for n coin CFMM

In the n coin scenario, the portfolio value [4] under no arbitrage is:

$$P_V = \sum_{i=1}^n \frac{\partial_i \Psi(R)}{\partial_n \Psi(R)} R_i = \sum_{i=1}^n m_i R_i$$

where $R \in \mathbf{R}_+^n$ is the reserve for each asset and we assume asset n is the numéraire, such that $m_n = 1$. Implicit differentiation of Ψ gives,

$$\frac{d\Psi(R)}{dm_i} = \sum_{i=1}^n \partial_i \Psi(R) \frac{dR_i}{dm_i} = 0$$

Dividing through by $\partial_n \Psi(R)$ gives

$$\frac{1}{\Psi_n(R)} \frac{d\Psi(R)}{dm_i} = \sum_{i=1}^n m_i \frac{dR_i}{dm_i} = 0$$

This yields

$$(P_\Delta)_i = \frac{dP_V}{dm_i} = R_i + m_i \frac{dR_i}{dm_i} + \sum_{j \neq i} \frac{dR_j}{dm_i} = R_i + \sum_{j \neq i} (m_j - 1) \frac{dR_j}{dm_i}$$

and

$$(P_\Gamma)_{ii} = \frac{d^2 P_V}{dm_i^2} = \frac{dR_i}{dm_i} + \sum_{j \neq i} (m_j - 1) \frac{d^2 R_j}{dm_i^2}$$

Note that when $m_j \approx 1$, which is what happens in stablecoin–stablecoin trading, P_Δ and P_Γ resemble the two-asset trading hedges.

F Conjecture: Delta hedging impermanent Loss

The linear bounds of (17) suggest that one can hedge P_Δ if $\frac{d\Delta}{dm}$ and $\frac{d\Delta'}{dm}$ can be replicated using options on impact. We note that the results presented here do not hold as stated since the assumption that the market price is decreasing, used in the definitions of μ and κ , is broken here. We suspect that more general results of this form do hold in practice, but leave this as an open question.

First, note that $\frac{d\Delta}{dm}$ is the inverse of the price impact function $\frac{dm}{d\Delta}$. Suppose that an LP wants to hedge their impermanent loss when the realized price impact is greater than a fixed quantity ξ . For instance, an LP in a stablecoin pool might believe that the price of the assets never deviate by more than 10% and fixed $\xi = 0.1$. Mathematically, ξ corresponds to an lower bound on the impact, e.g. $\frac{dm}{d\Delta} \geq \xi$. Recall that the Carr-Madan expansion [13, App. 1] says that any payoff $f \in L^2(\mathbf{R}_+)$ can be expressed, for $\kappa > 0$ as

$$f(F) = f(\kappa) + f'(\kappa) [(F - \kappa)_+ - (\kappa - F)_+] + \int_0^\kappa f''(K)(K - F)_+ dK + \int_\kappa^\infty f''(K)(F - K)_+ dK$$

Consider the function $f(F) = \frac{1}{F} \mathbf{1}_{F > \xi + \epsilon} \in L^2(\mathbf{R}_+)$ for some $\epsilon > 0$. Using the Carr-Madan expansion at $F = \frac{dm}{d\Delta}$ and $\kappa = \xi$, we have:

$$\left(\frac{dm}{d\Delta}\right)^{-1} = \int_{\xi+\epsilon}^\infty \frac{2}{K^3} \left(\frac{dm}{d\Delta} - K\right)_+ dK$$

This equation states that we can replicate the exposure $\frac{d\Delta}{dm} \mathbf{1}_{\frac{d\Delta}{dm} \geq \xi}$ by holding a portfolio of $\frac{2}{K^3}$ call options at strike K for all $K \in [\xi, \infty]$. Thus to hedge, we short this portfolio. Using put-call parity, this is equivalent to holding a quantity of the asset α and selling put options at strike prices K greater than our cutoff, weighted by $\frac{2}{K^3}$. This effectively provides a way for an LP to insure themselves against impermanent losses up to price impacts of size ξ by selling a portfolio of put options.

G Multiperiod informed trading

Suppose that we have a discrete time series of probabilities $\alpha_t \in [1/2, 1)$. At time t , an informed trader has probability α_t of predicting whether the true price $m(t+1)$ at time $t+1$ is equal to the trader's predicted price p_{t+1} . Associated to each α_t is a pair of quantities $\Delta(t), \Delta'(t) \in \mathbf{R}$ that represent the trade quantities needed to take the true price $m(t)$ to the predicted price p_{t+1} . We can compute these quantities implicitly via the following differential equation implied by (4):

$$\partial_1 \Psi(R(t) + \Delta, R'(t) - \Delta') = p_{t+1} \partial_2 \Psi(R(t) + \Delta, R'(t) - \Delta')$$

To execute this in practice, an informed trader would need compute the optimal quantities Δ^*, Δ'^* that satisfy these questions. However, given that Ψ can be quite complicated, such a trader would likely have to rely on approximate methods. In particular, it is likely that greedy methods like gradient descent would be employed given that there are latency constraints in realistic settings.

How would we compute these optimal quantities to trade? To utilize a local method like gradient descent, we first need to compute an objective function $h(\Delta, \Delta')$ that we minimize using $\partial_\Delta h, \partial_{\Delta'} h$. In light of the above equation, is natural to define the following objective function

$$h(\Delta, -\Delta', p) = \partial_1 \Psi(R(t) + \Delta, R'(t) - \Delta') - p \partial_2 \Psi(R(t) + \Delta, R'(t) - \Delta')$$

Let us assume that the informed trader is computationally constrained and attempts to approximate Δ, Δ' via gradient descent:

$$\Delta(t)_n = \Delta(t)_{n-1} - \eta_\alpha \partial_1 h(\Delta, \Delta', p) \quad (22)$$

$$\Delta'(t)_n = \Delta'(t)_{n-1} + \eta_\beta \partial_2 h(\Delta, \Delta', p) \quad (23)$$

This recursion is the *gradient descent-ascent algorithm* (GDA) [38] or forward-backward iteration [16]. Note that due to the natural constraint of increasing one quantity while reducing another, we take gradients in the opposite directions for each quantity. Moreover, this is how a number of CFMMs compute trade quantities in practice (*e.g.*, Curve, where it has been the source of a number of vulnerabilities [24, 53]).

One natural question to ask is: how many time steps n does an informed trader need to run this algorithm to compute the optimal trade quantities up to some prescribed error? It turns out that this is directly related to the curvature of the CFMM. Suppose that this algorithm is run for a maximum of $T > 0$ time steps. Then we can write a recurrence for the expected reserve quantities as a function of time:

$$E[R(t+1)|R(t)] = R(t) + \alpha_t \Delta(t)_T + (1 - \alpha_t) \hat{\Delta}(t) \quad (24)$$

$$E[R'(t+1)|R'(t)] = R'(t) + \alpha_t \Delta'(t)_T + (1 - \alpha_t) \hat{\Delta}'(t) \quad (25)$$

where the $\hat{\Delta}$ are the quantities that would need to trade if an oracle provided the optimal quantities to move from $m(t)$ to $m(t+1)$. In particular, the quantities $\hat{\Delta}, \hat{\Delta}'$ are the no-arbitrage quantities that are traded when a roundtrip trade is made by an arbitrageur and the LP books a profit (akin to eq. (7)).

If the informed trader has a very high amount of edge or information (*e.g.* $\liminf_{t \rightarrow \infty} \alpha_t = 1$), then (24) is completely dominated by the GDA time steps of the informed trader. From [38, Theorem 1], we find the result that when the GDA steps dominate the evolution of (24), then if $T = \Omega\left(\frac{\mu^2}{\kappa}\right)$, the expected reserves will be very close to the optimal reserves Δ^*, Δ'^* , *e.g.*

$$\forall \epsilon > 0, \exists t'(\epsilon) \in O(T) \text{ such that } |\Delta(t)_{T+t'(\epsilon)} - \Delta^*| < \epsilon$$

In particular, [38] uses effectively the same definition of curvature for solving minimax problems using GDA. We note that their results depend weakly on dimension, so that this recursion can be extended to informed traders interacting with n -asset CFMM markets.

This illustrates that the curvature of a CFMM also controls the computational complexity that an informed trader needs to act on perfect information (*e.g.*, $\alpha_t \approx 1$) and trade with a CFMM. This formulation was inspired by constructions in robust machine learning that resemble adversarial information aggregation in a CFMM. We conjecture that this provides a way to extend some of our results to higher dimensions.

H Optimal Balancer yield farming

Suppose that we have two Balancer pools consisting of two assets that have the same spot price p at time t' . Pool i for $i \in \{1, 2\}$ has reserves $R^i, R'^i \in \mathbf{R}_+$ and exponent $\tau_i \in (0, 1)$ so

that the trading function for the i th pool is

$$\Psi_i(R^i, R'^i) = (R^i)^{\tau_i} (R'^i)^{1-\tau_i}$$

By definition, the spot price equivalence means that:

$$p = \frac{R^1}{R'^1} \left(\frac{1 - \tau_1}{\tau_1} \right) = \frac{R^2}{R'^2} \left(\frac{1 - \tau_2}{\tau_2} \right) \quad (26)$$

Let $p_i(t)$ be the price of the i th pool at time t so that $p_1(t') = p_2(t') = p$. Assume that a liquidity provider owns the same fraction $b \in [0, 1]$ of each pool and that at time $t' + 1$ a trade of size Δ from x to y occurs. In the feeless case, this gives a quantity change $\tilde{\Delta}_i$ to each pool, where $\tilde{\Delta}_i$ is implicitly specified via the CFMM constraints

$$k_i = (R^i)^{\tau_i} (R'^i)^{1-\tau_i} = (x_i - \Delta)^{\tau_i} (y_i + \tilde{\Delta}_i)^{1-\tau_i}$$

Let $P_V(R^i, R'^i, t)$ be the portfolio value of pool i at time t in units of x . Then we have,

$$P_V(R^i, R'^i, t) = b (R^i + p_i(t') R'^i) = b R^i \left(1 + \frac{1 - \tau_i}{\tau_i} \right) = \frac{b}{\tau_i} R^i \quad (27)$$

$$P_V(R^i, R'^i, t' + 1) = \frac{b}{\tau_i} (R^i - \Delta) \quad (28)$$

This implies that if $R^1 = R^2$, then $\tau_1 > \tau_2$ implies $P_V(R^1, R'^1, t' + 1) < P_V(R^2, R'^2, t' + 1)$. Less formally, this says that losses to portfolio value, given an equal reserve of x are higher for less sharply curved Balancer pool (*e.g.*, τ_1 loses more than τ_2). If we wanted to incentivize liquidity in pool τ_1 by issuing new x to liquidity providers, how much would we have to pay? This value is exactly equal to $P_V(R^2, R'^2, t' + 1) - P_V(R^1, R'^1, t' + 1)$, which represents the ‘excess loss’ covered by printing token α (we’re assuming that’s the portfolio numéraire). Using this equation we have:

$$\begin{aligned} \delta(t' + 1) &= \frac{P_V(R^2, R'^2, t' + 1) - P_V(R^1, R'^1, t' + 1)}{b} = \frac{1}{\tau_2} (R^2 - \Delta) - \frac{1}{\tau_1} (R^1 - \Delta) \\ &= \frac{R^2}{\tau_2} - \frac{R^1}{\tau_1} + \Delta \left(\frac{1}{\tau_1} - \frac{1}{\tau_2} \right) \\ &= \left(\frac{R'^2}{\tau_2} \right) \left(1 - \left(\frac{R'^1}{R'^2} \right) \frac{1 - \tau_2}{1 - \tau_1} \right) \\ &\quad + \Delta \left(\frac{1}{\tau_1} - \frac{1}{\tau_2} \right) \end{aligned}$$

Thus, if $\delta(t' + 1)$ is paid out pro-rata to holders via a Synthetix-like yield farming mechanism, we can encourage users to stay in the worse pool. This is an improved formula over those used by pro-active market makers, such as Dodo [9], which don’t try to perform this accounting on a trade-by-trade basis and instead rely on an oracle. You can effectively force the user to lock liquidity into pool 1 from height t' to $t' + k$ and then pay out $\sum_{i=0}^k \delta(t' + i)$ upon redemption.

I Conjecture: yield farming is superhedging

Before making the comparison to superhedging in discrete time [14, §3.4], we will demonstrate that equations (13) and (14) can be recast as an optimization problem. Recall that for a function f , g is a subgradient of f at x , *e.g.*, $g \in \partial f(x)$ if $\forall y \in \mathbf{dom}(f)$, $f(y) - f(x) \geq g^T(y - x)$. Let $P_\varphi(R, c)$ be the portfolio value for an LP at price p , reserves R for trading function φ . Equation 13 can be restated in terms of portfolio value as,

$$P_\varphi(R, m_0^s) - P_\varphi(R, m_0^e) \geq -\frac{\mu}{\kappa}(m_0^s - m_0^e)$$

Let $f(p) = P_\varphi(R, p)$. Then this condition is equivalent to $-\frac{\mu}{\kappa} \in \partial f(m_0^e)$. Similarly, equation (14) corresponds to showing any admissible subsidy R_ℓ must satisfy $0 \in R_\ell + \partial f(m_0^e)$. This connects the subsidy to an optimization problem, as optima for a function f are found by showing that $0 \in \partial f$ [14, §1.3]. These subgradient conditions illustrate that there is a connection between the yield farming subsidy and optimizing changes in portfolio value. Since the change in portfolio value is equivalent to solving a dual optimization problem [4], this shows that yield farming subsidies are equivalent to bounding the payoff an LP engenders under an arbitrage trade by a simpler curvature dependent payoff. In mathematical finance, bounding complex payoffs with simpler payoffs is known as *superhedging* [14, §3.4].

Suppose, instead, that we start with the opposite problem of trying to find a CFMM trading function Ψ that is equivalent to given CFMM trading function φ and a subsidy R_ℓ . Another way to arrive at equation (13) is to consider the set of CFMMs that can have a valid portfolio value at $(R + (R_\ell, 0), p)$ and find an upper bound for this portfolio value. More formally, suppose that we define the following set of trading functions:

$$\mathcal{P}(R, R_\ell, R', m', \mu, \kappa) = \{\varphi : (R + R_\ell, R', m') \in \mathbf{dom}(P_\varphi), \varphi \text{ is } \mu\text{-stable and } \kappa\text{-liquid}\}$$

We can consider $\mathcal{P}(R, R_\ell, R', m', \mu, \kappa)$ to be a subset of the set of closed, proper, l.s.c. convex functions [4]. If we can find a quasiconvex function ψ such that

$$\psi = \sup \mathcal{P}(R, R_\ell, R', m', \mu, \kappa) \tag{29}$$

then ψ must satisfy (13) as per the discussion in §3.3. With some slight modifications, trading functions can be put in bijective correspondence with portfolio values, which are contingent claims. Finding a contingent claim whose payoff is the supremum over a set of admissible contingent claims is known as a superhedge [14, §3.4], [17]. Both [14, 17] map superhedging to a convex dual problem analogous to finding portfolio value that is dual to (29). These formulations are very similar to the relationship between trading function and portfolio value from [4]. While traditional superhedging involves taking a supremum over a set of equivalent martingale measures, we are instead taking a supremum over a set of contingent claims whose curvatures are constrained. There is some literature on superhedging over sets of contingent claim payoffs that are not equivalent martingale measures (*e.g.*, [49]), however, none of the dual frameworks known to the authors map cleanly to our definitions of curvature. We conjecture a curvature claim result analogous to [17, Prop. 2.3] exists.