

Differential Privacy in Constant Function Market Makers

Tarun Chitra¹, Guillermo Angeris², and Alex Evans³

¹ Gauntlet Networks, Inc., tarun@gauntlet.network

² Stanford University, angeris@stanford.edu

³ Bain Capital, aevans@baincapital.com

Abstract. Constant function market makers (CFMMs) are the most popular mechanism for facilitating decentralized trading. While these mechanisms have facilitated hundreds of billions of dollars of trades, they provide users with little to no privacy. Recent work illustrates that privacy cannot be achieved in CFMMs without forcing worse pricing and/or latency on end users. This paper quantifies the trade-off between pricing and privacy in CFMMs. We analyze a simple privacy-enhancing mechanism called *Uniform Random Execution* and prove that it provides (ϵ, δ) -differential privacy. The privacy parameter ϵ depends on the curvature of the CFMM trading function and the number of trades executed. This mechanism can be implemented in any blockchain system that allows smart contracts to access a verifiable random function. Our results provide an optimistic outlook on providing partial privacy in CFMMs.

1 Introduction

Constant function market makers (CFMMs) have become the most widely used decentralized product. In 2021, these market makers were facilitating over a billion dollars of daily (spot) volume, comparable to centralized exchanges such as Binance, Coinbase, or FTX. These market makers allow those looking for passive yield on a portfolio of assets to be automatically matched with traders looking to execute a swap against their assets. CFMMs work by ensuring that an invariant known as the trading function is kept constant before and after a trade is executed. The trading function, which is a function of the liquidity provided by those seeking passive yield, controls the price displayed by the CFMM that traders can execute a trade at. In order to ensure that liquidity providers (LPs) do not always lose money, as they are effectively buying the currency whose value is going down in exchange for one that is going up, a trading fee is applied to each transaction. Prior work [AC20, AAE⁺21, AEC20] has investigated necessary and sufficient conditions for the trading function and choice of fee to lead to profitable outcomes for LPs.

Privacy in CFMMs. One major problem with CFMMs is their lack of privacy. At a high-level, privacy in CFMMs boils down to preventing an adversary from

discerning trade sizes as a function of public prices and the knowledge of a feasible trade. Additionally, the dramatic increase in maximal extractable value (MEV) and front-running on Ethereum makes transaction-level privacy increasingly important [QZLG20, QZG21, ZQT⁺20, AEC21a]. Mechanisms that reduce the amount of information that an adversary has about user transactions can help reduce MEV and increase privacy. However, we do note that there are subtle distinctions between mechanisms that reduce MEV versus mechanisms that increase privacy (which we address in §3.1) In this paper, we study mechanisms purely in terms of privacy, but are motivated to study the problem in the hope that some form of privacy might address MEV.

Prior work [AEC21b] has shown that given any feasible trade and the (usually public) prices before and after executed trades, one can uniquely identify the size of the trade. This is a natural (although somewhat indirect) consequence of the concavity of the trading function [AC20, AAE⁺21]. This work implies that, even with modern cryptography such as zero-knowledge proofs (ZKPs), one will need to modify the CFMM mechanism to blind user’s trade sizes. In other words, simply hiding balances via ZKPs of reserves (which has been proposed and implemented in multiple protocols [CXZ20, Pow21]) is not sufficient for transaction-level privacy.

Proposed solutions. The two main options presented in [AEC21b] for recovering privacy involve either modifying prices (*e.g.*, adding noise to quoted prices) or batching transactions.⁴ Both of these changes often degrade the user experience: both options force traders to bear worse price impact while the latter option also means that users face higher latency for trade confirmation. Assuming that these are the only options available, a natural question to ask is: how well can we control the trade-off between worsened price and latency and improved transaction privacy? One might formulate this rather general question as the following:

- What is the minimum number of swaps, $n(\delta)$, that must be batched such that an adversary is unable to infer the true trade sizes, beyond a precision of δ ?
- How much worse is the worst price offered to any one user via such a mechanism?

Answers to the former question are analogous to sample complexity bounds from learning theory, whereas answers to the latter question measure the ‘cost of privacy’.

Differential privacy. One method for answering questions of this form is through the lens of differential privacy [DR⁺14]. Differentially private algorithms aim to hide individual user data (*e.g.*, trades) while simultaneously preserving aggregate statistics (*e.g.*, prices or averages). Many differentially private mechanisms work by adding targeted randomness to each individual users’ data. As a simple

⁴ There are two live batching CFMMs in production, CowSwap on Ethereum [Mar21] and Penumbra which relies on a specialized ZKP chain [dV21].

illustration, suppose for which we have a sequence of values x_1, \dots, x_n and we want to report the mean $\mu = \frac{1}{n} \sum_i x_i$. One can (in expectation) preserve the mean μ by adding i.i.d., mean zero noise to each x_i , before computing and reporting the new mean, $\tilde{\mu}$. Intuitively, as the variance of the added noise becomes large, it is harder to recover the original value of x_i , prior to the added noise, but the new reported mean $\tilde{\mu}$ is likely to be far away from the true mean μ .

In this sense, differentially private algorithms induce a natural trade-off between the privacy and accuracy of a query, much like the trade-off between price impact and privacy in CFMMs. We note that the methods from differential privacy have been used at scale and in production at the US Census [Dwo19], Google [ACG⁺16], and Apple [CJK⁺18]. Our threat model for the adversary (§3.1) involves an adversary trying to estimate an ordered vector of trades given prices. In this scenario, we can view the set of trades to be executed as “private” user data while the accuracy of the query is the deviation in price that users have to pay for privacy. Differential privacy is a natural way to study the expected worst case behavior of such an estimation process, similar to its usage within machine learning.

We note that achieving differential privacy, even with non-private noise, can help reduce expected MEV profits. Moreover, as differential privacy has often been used in machine learning to improve algorithmic fairness, we posit (without proof) that differentially private DeFi algorithms inherit fairness guarantees [DHP⁺12]. These fairness guarantees are distinct of those from cryptographic fair ordering [KZGJ20], as they provide explicit guarantees on the trade-off between (economic) utility, privacy, and fairness [CGKM19, XYW19].

Uniform random execution. To achieve differential privacy in CFMMs, we construct a black-box algorithm called *Uniform Random Execution* (URE). This algorithm can be viewed as the inverse of batching, as it breaks up and splits large trades before subsequently randomly permuting the trade ordering. Randomness is used for both splitting up large trades and for permuting the split up trades. Blockchains with smart contract capabilities that include CFMM ordering as part of consensus rules can execute the URE (*e.g.*, Celo [KOR19], Terra [MSS20], Penumbra [dV21], Osmosis [AO21]). In particular, any blockchain with a verifiable random function (VRF) [MRV99] that provides public randomness and consensus rules for executing trades in a particular order suffices for URE.

Summary. Our analysis of the differential privacy of URE utilizes a novel representation of a sequence of trades as a binary tree. The tree is constructed such that the height of the tree provides a lower bound on the worst case price impact. On the other hand, number of leaves of the tree controls how easy it is for one to invert the precise trades executed. Representing continuous objects (sequences of real-valued trades) as a random discrete data structure allows us to utilize traditional tools from differential privacy. We show that the trade tree controls the maximum price impact of a sequence of trades by utilizing curvature of a CFMM [AEC20, §2]. Curvature represents bounds on market impact

cost and liquidity and is crucial for relating the trade tree to worst-case price impact. Subsequently, we analyze the impact of splitting up and randomly permuting trades on the trade tree and then compute bounds on the price impact associated to these actions.

In order to achieve differential privacy, we first prove that splitting up trades can be executed in a differentially private manner (claim 3.4). To split a trade, we sample a random distribution π and the split up a single trade according to π . After splitting up the trades, we then show that randomly permuting the trades leads to an (ϵ, δ) -differentially private algorithm. We use composition laws [DR⁺14, KOV15] to combine these two results and show that the URE is differentially private. Note that ϵ and δ depend on the CFMM’s curvature and on the on the number of trades executed.

2 Preliminaries

We will cover preliminaries on CFMMs and differential privacy. For more details, please refer to review articles on CFMMs [AAE⁺21] and differential privacy [DR⁺14].

2.1 Constant function market makers

A *constant function market maker* is a contract that holds some amount of reserves $R, R' \geq 0$ of two assets and has a *trading function* $\psi : \mathbf{R}^2 \times \mathbf{R}^2 \rightarrow \mathbf{R}$. Traders can then submit a *trade* (Δ, Δ') denoting the amount they wish to tender (if negative) or receive (if positive) from the contract. The contract then accepts the trade if

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

and pays out (Δ, Δ') to the trader.

Curvature. We briefly summarize the main definitions and results of [AEC20] here. Suppose that the trading function ψ is differentiable (as most trading functions in practice are), then the marginal price for a trade of size Δ is

$$g(\Delta) = \frac{\partial_3 \psi(R, R', \Delta, \Delta')}{\partial_4 \psi(R, R', \Delta, \Delta')}.$$

Here ∂_i denotes the partial derivative with respect to the i th argument, while Δ' is specified by the implicit condition $\psi(R, R', \Delta, \Delta') = \psi(R, R', 0, 0)$; *i.e.*, the trade (Δ, Δ') is assumed to be valid. Additionally, the reserves R, R' are assumed to be fixed. The function g is known as the *price impact* function as it represents the final marginal price of a positive sized trade. When there are fees, one can show that $g^{fee}(\Delta) = \gamma g(\gamma \Delta)$ where $1 - \gamma$ denotes the percentage fee. We say that a CFMM is μ -stable if it satisfies

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

for all $\Delta \in [0, M]$ for some positive M . This is a linear upper bound on the maximum price impact that a bounded trade (bounded by M) can have. Similarly, we say that a CFMM is κ -liquid if it satisfies

$$g(0) - g(-\Delta) \geq \kappa\Delta$$

for all $\Delta \in [0, K]$ for some positive K . Simple methods for computing some μ and κ in common CFMMs are presented in [AEC20, §1.1].

Two-sided bounds. We can define similar upper and lower bounds for $g(\Delta) - g(0)$, with constants μ' and κ' , which hold when the trades Δ are in intervals $[0, M']$, $[0, K']$, respectively. For the remainder of this paper, we will refer to μ -stability as the upper bound for both $g(0) - g(-\Delta)$ and $g(\Delta) - g(0)$, and similarly for κ -liquidity. More specifically, given μ, μ' , we say that a CFMM is symmetrically μ'' -stable if

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

when $-M \leq \Delta \leq M'$, and symmetrically κ'' stable if

$$|g(\Delta) - g(0)| \geq \kappa|\Delta|.$$

when $-K \leq \Delta \leq K'$. From the above, it suffices to pick $\mu'' = \min\{\mu, \mu'\}$ and $\kappa'' = \min\{\kappa, \kappa'\}$.

For the remainder of this paper, we will focus on using CFMM curvature parameters to bound the impact cost realized, which in turn controls how easily an adversary can invert a trade size from prices.

2.2 Differential Privacy

Differential privacy is a framework for classifying how well a randomized algorithm \mathcal{A} anonymizes individual data points.

Definition 1. *A randomized algorithm \mathcal{A} is (ϵ, δ) -differentially private if for all $S, S' \in \text{Dom } \mathcal{A}$ with $d(S, S') \leq 1$ we have for all measurable $B \subset \text{Range } \mathcal{A}$*

$$\mathbf{Prob}[\mathcal{A}(S) \in B] \leq e^\epsilon \mathbf{Prob}[\mathcal{A}(S') \in B] + \delta$$

In this definition, ϵ can be thought of as a uniform upper bound on the Kullback-Leibler divergence over the distribution induced by any pair of neighboring data sets. Traditionally, S, S' are thought of as discrete and the metric d corresponds to the Hamming metric. In this case, the intuition behind the definition is the following: changing one entry of the variable S' does not change the output distribution ‘too much,’ making it difficult to tell apart S from S' by looking only at the results of algorithm \mathcal{A} . In this paper, we will assume d is the L^1 norm [DR⁺14, NRS07]. We provide further details on differential privacy in appendix A.

3 Problem Construction

In the discussions of [AEC21b, §3], two ways of providing approximate privacy are presented:

1. *Randomizing price*: the protocol can randomly perturb the price quoted by the CFMM, in manner resistant to adversaries (while also not destroying liquidity provider returns).
2. *Batching orders*: picking a number of orders n to batch prior to execution.

Neither of these proposed solutions are perfect and [AEC21b] provides no adversarial model for assessing them. Here we first formulate a simple adversarial threat model for these solutions and then introduce URE. To construct URE, we first describe a simpler method called *Simple Uniform Random Execution* (SURE) which achieves differential privacy under restrictive conditions on trade sizes. We then prove that the URE achieves differential privacy by modifying SURE using extra randomness whose entropy is parametrized by the number of trades to execute and the curvature. For the remainder of the paper, we will assume that there are only two assets traded (in order to utilize curvature bounds) while n will refer to the number of trades executed.

3.1 Threat Model

Adversary definition and attack. We assume a simple model of an adversary that generalizes [AEC21b]. The adversary, who we will call Eve, attempts to discover the quantities traded by a set of agents referred to as Traders. Eve is unable to see the exact quantities the Traders use to trade with the CFMM, but knows when the Traders transactions $\Delta_1, \dots, \Delta_n$ are executed as a block. Eve does not know the order in which the trades are executed and her goal is to estimate the ordering and sizes of the trades. Her only ability is to interact with the CFMM in the state before the traders' transactions are executed and the state after their transactions are executed. Explicitly, Eve's goal is to produce a vector $(\tilde{\Delta}_1, \dots, \tilde{\Delta}_n)$ such that $\|(\tilde{\Delta}_1, \dots, \tilde{\Delta}_n) - (\Delta_1, \dots, \Delta_n)\|_1$ is small with high probability. Differential privacy provides a precise way of characterizing the probability of such a scenario occurring.

When a user submits a transaction to a blockchain, they send a transaction via a peer-to-peer network that reaches a miner or validator. In both proof-of-stake and layer 2 chains, the validator who chooses the final execution order of transactions is known as a *sequencer*. For the remainder of this paper, any reference to the sequencer will assume that the sequencer is honest (e.g. they execute a given ordering when received from an MEV auction). Unless the blockchain uses a fully homomorphic virtual machine (which does not currently exist), the sequencer necessarily sees a user's transaction in order to execute a valid state transition. Fair-ordering systems [KZGJ20, KDL⁺21] attempt to decentralize this sequencing operation, albeit with extra assumptions on validator behavior. Our threat model does not prevent the sequencer from discovering Traders' trades and front-running them as we assume that Eve is not the (honest) sequencer.

Action space. We assume that Eve has access to two queries:

- `marginalPrice()`: Computes the marginal price of the CFMM at its current reserves
- `isValid(Δ)`: Takes a trade $\Delta \in \mathbf{R}$, returns `True` if the trade is valid and `False` otherwise

We will denote the set of valid trades at reserves $R \in \mathbf{R}_+^n$ as $\mathcal{A}_\varphi(R)$ and note that it can effectively be thought of as the epigraph of the trading function φ [AC20].

3.2 Simple Uniform Random Execution

One of the simplest ways to introduce entropy into a CFMM is to randomly permute the set of trades to be executed. We will first describe the *simple uniform random execution* (SURE) mechanism that simply permutes observed trades. Formally, suppose that we are given a vector of valid trades

$$\Delta_1 \in \mathcal{A}_\varphi(R), \Delta_i \in \mathcal{A}_\varphi\left(R + \sum_{j=1}^{i-1} \Delta_j\right)$$

For brevity, we will refer to above condition as $\mathcal{A}_\varphi(\mathbf{\Delta})$ for a trade vector $\mathbf{\Delta}$. The SURE mechanism draws a random permutation $\pi \sim_{\text{Unif}} S_n$ and constructs a sequence of trades $\Delta_i^\pi = \Delta_{\pi(i)}$, which arise from permuting the order in which the trades are executed. Consider the marginal prices of the original trades p_1, \dots, p_n and the permuted prices p_1^π, \dots, p_n^π . Note that $p_n = p_{\pi(n)}$ if and only if the CFMM is path-independent (*e.g.*, feeless). Our goal is two-fold: first, we aim to bound the maximum deviation between the true price p and the permuted prices p^π . That is, we want to compute

$$\mathcal{E}_{SURE} = \mathbf{E}_{\pi \sim S_n} \left[\max_{i \in [n]} |p^\pi(i) - p(i)| \right]$$

This deviation effectively corresponds to a bound on the worst quoted price that a trader can receive (relative to their original order price). Secondly, we want to capture a notion of how difficult it is for an adversary to learn the values of π chosen given only the prices p_n^π .

Before analyzing the SURE mechanism for some classes of trades, let's look at some simple examples. If all of the trades Δ_i are unique — *e.g.*, $\nexists i, j \in [n]$ such that $\Delta_i = \Delta_j$ — then computing Δ_i given p^π is in some sense be difficult to invert to a precision higher than $\kappa \min_{i,j} |\Delta_i - \Delta_j|$. This is because if π is a single adjacent transposition ($i \ i + 1$), then $g(\sum_{j=1}^i \Delta_i) - g(\sum_{j=1}^{i-1} \Delta_j) \geq \kappa \min(\Delta_i, \Delta_{i-1}) \geq \kappa \min_{i,j} |\Delta_i - \Delta_j|$. Moreover, we should expect that SURE should work better when $\sum_{i=1}^n \text{sgn}(\Delta_i) \approx 0$. This is because the probability of having a long run of trades in the same direction is very low. For instance, if Δ_i is a Rademacher random variable (*e.g.*, uniformly drawn from $\{-1, 1\}$) then the expected maximum length of a run is $\Theta(\log n)$ [ER75, Theorem 1].

On the other hand, if there is a set $S \subset [n]$ with $|S| = \Omega(n)$ such that for all $i, j \in S$, $\Delta_i = \Delta_j$, then it will be much easier to invert the set of trades. There is a loss of entropy in the output trade sequences as there will be many permutations $\pi, \pi', \pi \neq \pi'$ such that $\Delta^\pi = \Delta^{\pi'}$. Let's consider an explicit numerical example. Let $\Delta_1 = 100$ and $\Delta_i = 1$ for all $i \in \{2, \dots, n\}$. Even though we are sampling from $n!$ permutations, there are only n output sequences that SURE outputs: $\Delta_j^\pi = \Delta_{\pi(1)} = 100$ in the j th position for $j \in [n]$. Suppose we consider a permutation π with $\pi(1) = j$. For any trade in position i with $\pi(i) < j$, the trade gets significantly better execution than they did initially. This is because their trade is executed before the trade of size 100 is executed, giving them significantly less impact. Therefore, SURE requires the trade distribution to have sufficient entropy and the distribution of trade sizes to not be too concentrated in order to work.

We will first analyze SURE on a subset of allowable input trades. This subset will be defined via simple constraints on $\min_{i,j} |\Delta_i - \Delta_j|$. We later relax these by splitting up large trades in a manner that ensure that the trade size distribution satisfies these constraints with high probability. To analyze SURE, we will start by obtaining upper and lower bounds on the worst case expected price discrepancy, $\mathbf{E}[\max_i |p^\pi(i) - p(i)|]$. This analysis will provide insight into what subset of admissible trades provide provable bounds on price discrepancy and identifiability.

Maximum of the Price Process and Random Binary Trees. Suppose the price impact function g is κ -liquid and μ -stable on an interval $[-M, M]$. By definition this implies that for all $i \in [n]$

$$\sum_{j=1}^i \kappa \Delta_{\pi(j)} - \mu \Delta_j \leq p^\pi(i) - p(i) \leq \sum_{j=1}^i \mu \Delta_{\pi(j)} - \kappa \Delta_j$$

This means that we have

$$\kappa \max_i \left| \sum_{j=1}^i \Delta_{\pi(j)} - \frac{\mu}{\kappa} \Delta_j \right| \leq \max_i |p^\pi(i) - p(i)| \leq \mu \max_i \left| \sum_{j=1}^i \Delta_{\pi(j)} - \frac{\kappa}{\mu} \Delta_j \right| \quad (1)$$

Therefore, bounds on partial sums of permuted trades will allow us to bound the worst case price impact of SURE. Define the partial sum

$$\rho_i(\Delta, \pi) = \sum_{j=1}^i \Delta_{\pi(j)} - \frac{\mu}{\kappa} \Delta_j \quad (2)$$

Now consider the binary search tree $T(\rho(\Delta, \pi))$ whose root is $\rho_1(\Delta, \pi)$. Each element $\rho_j(\Delta, \pi)$ is inserted sequentially to construct the tree (see Figure 1 for an example).

This representation of the partial sums as a tree provides a natural geometric description of the maximum price deviation. In particular, $\max_i \rho_i(\Delta, \pi)$ is




Fig. 1: Depiction of the tree $T(\rho(\Delta, \pi))$ where $\rho_i = \rho_i(\Delta, \pi)$ and $\rho_{11} < \rho_{10} < \rho_3 < \rho_{12} < \rho_4 < \rho_1 < \rho_8 < \rho_5 < \rho_9 < \rho_2 < \rho_7 < \rho_6$

necessarily a leaf node in this tree. This means that the maximum deviation $\max_{i,j} |\rho_i(\Delta, \pi) - \rho_j(\Delta, \pi)|$ is at most 2 times the height of the tree as the distance from ρ_1 to any element is maximized by the height. This provides the following bounds using (1)

$$\max_i |p^\pi(i) - p(i)| \leq \mu \left(|\rho_1(\Delta, \pi)| + \max_j \left| \Delta_{\pi(j)} - \frac{\kappa}{\mu} \Delta_j \right| \cdot 2 \cdot \text{height}(T(\rho(\Delta, \pi))) \right) \quad (3)$$

$$\max_i |p^\pi(i) - p(i)| \geq \kappa \left| \rho_1(\Delta, \pi) + \min_j \left(\Delta_{\pi(j)} - \frac{\mu}{\kappa} \Delta_j \right) \cdot 2 \cdot \text{height}(T(\rho(\Delta, \pi))) \right| + O(1) \quad (4)$$

Note that the second bound comes from bounded support of curvature:

$$\begin{aligned} \max_j |\rho_j(\Delta, \pi)| &\geq |\rho_j(\Delta, \pi)| = \left| \rho_1(\Delta, \pi) + \sum_{i=1}^j \left(\Delta_{\pi(i)} - \frac{\mu}{\kappa} \Delta_i \right) \right| \\ &\geq \left| \rho_1(\Delta, \pi) + \min_j \left(\Delta_{\pi(j)} - \frac{\mu}{\kappa} \Delta_j \right) \cdot 2 \cdot \text{height}(T(\rho(\Delta, \pi))) \right| + O(1) \end{aligned}$$

Moreover, the number of leaves in the tree represent the number of left-to-right local maxima of ρ_j . Note, furthermore, that by using curvature and the tree structure, we have reduced the maximum price deviation problem (a continuous problem) into a combinatorial one regarding a random tree. If the tree is roughly balanced (*e.g.*, height is $O(\log n)$) and there are $\Omega(n)$ leaf nodes then it is unlikely that a small change to the permutation π by a transposition will change the maximum value. We will formalize this by studying the behavior of the random variable $T(\Delta)$, which draws a permutation π randomly and sets $T(\Delta) = T(\rho(\Delta, \pi))$.

To study the behavior of $T(\Delta)$, we need to analyze the expected height of a random binary tree. It is known that the height of a random binary tree with

distinct elements (*e.g.*, such that every permutation is equiprobable) has height $\Theta(\log n)$ with high probability:

Theorem 1 (Theorem 1 [Ree03]). *Let Δ have unique elements. Then $\mathbf{E}[\text{height}(T(\Delta))] = \alpha \log n - \beta \log \log n$ and $\text{Var}[\text{height}(T(\Delta))] = O(1)$*

If we can guarantee that the elements of $T(\Delta)$ are distinct (*e.g.*, such that every permutation of Δ is equiprobable) then combining this result with (3) yields

$$\begin{aligned} \mathbf{E}[\max_i |p^\pi(i) - p(i)|] &\leq \mu \left(\mathbf{E}[\rho_1(\Delta, \pi)] + 2 \max_{i,j} \left| \Delta_i - \frac{\kappa}{\mu} \Delta_j \right| \cdot \mathbf{E}[\text{height}(T(\Delta))] \right) \\ &\leq \mu \left(\mathbf{E}[\rho_1(\Delta, \pi)] + 2 \max_{i,j} \left| \Delta_i - \frac{\kappa}{\mu} \Delta_j \right| (\alpha \log n - \beta \log \log n) \right) \\ &\leq 3\mu \left(\max_{i,j} \left| \Delta_i - \frac{\kappa}{\mu} \Delta_j \right| \right) (\alpha \log n - \beta \log \log n + 1) \quad (5) \end{aligned}$$

where we used the upper bounds $\max_j \left| \Delta_{\pi(j)} - \frac{\kappa}{\mu} \Delta_j \right| \leq \max_{i,j} \left| \Delta_i - \frac{\kappa}{\mu} \Delta_j \right|$ and

$$\mathbf{E}[\rho_1(\Delta, \pi)] = \frac{1}{n} \sum_{j=1}^n \left| \Delta_j - \frac{\kappa}{\mu} \Delta_1 \right| \leq \max_{i,j} \left| \Delta_i - \frac{\kappa}{\mu} \Delta_j \right|$$

Similarly, note that $\mathbf{E}[\rho_1(\Delta, \pi)] \geq \min_j \left(\Delta_{\pi(j)} - \frac{\mu}{\kappa} \Delta_j \right)$ so we have

$$\begin{aligned} \mathbf{E}[\max_i |p^\pi(i) - p(i)|] &\geq \kappa \left| \min_j \Delta_{\pi(j)} - \frac{\mu}{\kappa} \Delta_j \right| (2 \mathbf{E}[\text{height}(T(\Delta))] + O(1)) \\ &\geq 2\kappa \left| \min_{i,j} \Delta_i - \frac{\mu}{\kappa} \Delta_j \right| (\alpha \log n - \beta \log \log n + O(1)) \end{aligned}$$

Therefore, provided that the following two conditions hold

$$\Delta_{\min} = \left| \min_{i,j} \Delta_i - \frac{\mu}{\kappa} \Delta_j \right| = \Omega(1) \quad \Delta_{\max} = \left| \max_{i,j} \Delta_i - \frac{\kappa}{\mu} \Delta_j \right| = O(1) \quad (6)$$

we have $\mathcal{E}_{SURE} = \mathbf{E}[\max_i |p^\pi(i) - p(i)|] = \Theta(\log n)$. Such a bound is ideal as it ensures that there is always a minimum price discrepancy of $\Omega(\kappa \log n)$ so that an adversary cannot determine a trade size with precision greater than $\Omega(\kappa)$. On the other hand, the upper bound on price deviation means that the mechanism will not cause too great of a price impact for users.

Note that the usage of Theorem 1 is prefaced on every permutation of the elements of $\rho_j(\Delta, \pi)$ being equiprobable. One simple example of when this isn't true is from the threshold trades, $\Delta = (T, 1, \dots, 1) \in \mathbf{R}^T$ when $\mu \geq 100\kappa$. When this is true, neither of the conditions (6) hold and moreover, the conditions of Theorem 1 do not hold. This means that SURE only works when (a) all permutations of partial sums are unique and (b) when $\mu \leq (\max_i \Delta_i)\kappa$. In the next section, we will achieve (a) by adding noise dependent on Δ, μ, κ to the trades and (b) by splitting trades.

3.3 Uniform Random Execution

We have seen the SURE mechanism works well at providing privacy while minimizing price discrepancy when (6) holds, when elements of Δ are unique, and when $\frac{\mu}{\kappa}$ is not too large. However, we're not guaranteed that both of these conditions hold in general as illustrated by the example at the end of the last section. This section will focus on using randomization to ensure that a) (6) holds with high probability and b) the elements of Δ are unique. We will do this by performing two actions: splitting large trades to ensure the maximum condition holds and adding noise to trades to ensure that trades are not too close in size. Applying these two actions to Δ and subsequently executing SURE is termed the *Uniform Random Execution* mechanism. There are three parameters that control the URE mechanism:

- c_{\min} : Lower bound on Δ_{\min}
- $s \in \mathbf{R}_+$: Split threshold that controls the average chunk size for a big trade
- $k \in \mathbf{N}$: Multiple of $(1 + s)\Delta_{\min}$ that requires splitting

Lower bounding the minimum by adding Laplace noise. Our goal is to construct random variables ξ_1, \dots, ξ_n drawn i.i.d. from a distribution that can depend on a particular Δ but guarantees that $\tilde{\Delta} = \Delta + \xi$ satisfies the left hand side of (6) with high probability. In particular, we would like to control $\mathbf{Prob} \left[\left| \min_{i,j} \tilde{\Delta}_i - \frac{\kappa}{\mu} \tilde{\Delta}_j \right| > c_{\min} \right]$ for a constant $c_{\min} > 0$. We desire the following condition to hold bounded above by $\delta \in (0, 1)$:

$$\begin{aligned}
 \mathbf{Prob} \left[\left| \min_{i,j} \tilde{\Delta}_i - \frac{\kappa}{\mu} \tilde{\Delta}_j \right| \leq c_{\min} \right] &= \mathbf{Prob} \left[\left| \min_{i,j} \Delta_i + \xi_i - \frac{\kappa}{\mu} (\Delta_j + \xi_j) \right| \leq c_{\min} \right] \\
 &\leq \mathbf{Prob} \left[- \left| \min_{i,j} \Delta_i - \frac{\kappa}{\mu} \Delta_j \right| + \left| \min_{i,j} \xi_i - \frac{\kappa}{\mu} \xi_j \right| \leq c_{\min} \right] \\
 &= \mathbf{Prob} \left[\left| \min_{i,j} \xi_i - \frac{\kappa}{\mu} \xi_j \right| \leq c_{\min} + \left| \min_{i,j} \Delta_i - \frac{\kappa}{\mu} \Delta_j \right| \right] \leq \delta
 \end{aligned} \tag{7}$$

In Appendix C, we prove the following claim:

Claim. There exists $a \in \mathbf{R}$ dependent on Δ, μ, κ and $\xi_i \sim \text{Lap}(a, |a|)$ such that (7) holds

This mechanism can be naturally modified to inherit the ϵ -privacy guarantees of the Laplace mechanism [DR⁺14, §3.2]. Note that the dependence of the noise parameter a on Δ is similar to smoothed sensitivity in differential privacy [NRS07]. We note that this added noise ensures both that the lower bound of (6) holds and ensures that the elements of $\Delta + \xi$ are unique so that Theorem 1 holds.

Upper bounding the maximum by splitting trades. One way to reduce the upper bound on error in (5) is to split up a trade Δ_i . This reduces Δ_{\max} and as explained in Appendix G, also increases the privacy of SURE. More precisely, we split Δ_i into Δ'_i, Δ''_i with $\Delta_i = \Delta'_i + \Delta''_i$ and then consider the pricing error associated to $p(\mathbf{\Delta}')$ where $\mathbf{\Delta}' = (\Delta_1, \dots, \Delta_{i-1}, \Delta'_i, \Delta''_i, \Delta_{i+1}, \dots, \Delta_n)$. This process can be iterated until all trades meet a particular criteria. Instead of splitting trades in two, we instead split trades into $m(\Delta_i)$ pieces, where $m(\Delta_i)$ is defined as

$$m(\Delta_i) = \max\left(1, \left\lceil \frac{|\Delta_i|}{(1+s)\Delta_{\min}} \right\rceil\right)$$

That is, the mechanism splits the trade into $m(\Delta_i)$ pieces who sizes are roughly $(1+s)\Delta_{\min}$. Let $\mathbf{1}^m = (1, \dots, 1)$. For any trade Δ_i with $m(\Delta_i) > 1$, we draw $\pi \sim \text{Dir}(\mathbf{1}^{m(\Delta_i)})$ and split Δ_i into trades $\Delta_{i,j} = \Delta_i \pi_j$. Since $\sum_{j=1}^n \pi_j = 1$, this provides a natural mechanism for splitting trades in a single step. As the Dirichlet distribution is sub-Gaussian when using uniform weights [MA17] and as the expected order statistics of a Dirichlet process decay exponentially [BJP12], $\text{Prob}[\Delta_{i,j} - (1+s)\Delta_{\min} > k\Delta_{\min}]$ also decays exponentially in k . This ensures that we have very few chunks that are significantly greater than $(1+s)\Delta_{\min}$, which ensures that with high probability $\max_i \Delta_i < (1+s+k)\Delta_{\min}$. As described in Appendices D and G, this condition ensures that SURE is effective with high probability. We note that the precise price impact of splitting trades (as a function of curvature) is analyzed in [AEC20].

3.4 Differential Privacy

We are now in a position to prove that the URE mechanism satisfies (ϵ, δ) -differential privacy, where $\epsilon = O(\mu \log n + \max_i \Delta_i)$. Our proof proceeds in two steps. First, we prove the following claim in the Appendix E.

Claim (Splitting is differentially private). Suppose that we have a sequence of admissible trades $\mathbf{\Delta} \in \mathbf{R}^n$ and after adding noise we have $\tilde{\mathbf{\Delta}}$ with $\tilde{\Delta}_{\min} > 0$. For each $k \in \mathbf{N}$ define $S_k = \{j : \tilde{\Delta}_j > k\tilde{\Delta}_{\min}\}$. If $\eta^* = \max_j \frac{\Delta_j}{\Delta_{\min}} = O(n)$ and there exists $k > 0$ such that $|S_k| = O(1)$, then there exists an (ϵ, δ) -differentially private algorithm $\text{Split}(\tilde{\mathbf{\Delta}})$ for splitting trades in $\tilde{\mathbf{\Delta}}$ such that $\text{height}(T(\text{Split}(\tilde{\mathbf{\Delta}}))) = O(\log n)$ where $\epsilon = O(\eta^*)$

This claim ensures that under mild conditions on the maximum trade size, we can generate a partial sum trade tree of height $O(\log n)$. Note that we can get the claim's conditions to be satisfied by varying s , the scale parameter, which leads to a privacy-utility trade-off. Second, we show that when a partial sum trade tree has height $O(\log n)$, permuting the trades provides $(O(\mu \log n), \delta)$ -differential privacy for the maximum price impact (Claim 3.4). We combine these two differentially private algorithms using standard composition theorems (see Appendix), resulting in a differentially private CFMM.

Claim (SURE is differentially private). Suppose that we have a sequence of admissible trades $\mathbf{\Delta} \in \mathbf{R}^n$ such that $\text{height}(T(\mathbf{\Delta})) = O(\log n)$ and all trade

sizes are unique. Then randomly permuting the trades Δ^π can be made into a $(\mu \log n, \delta)$ -differential private algorithm for the minimum and maximum price impact

While the full proof of the theorem is in the appendix, we sketch the steps of the proof below. First, we show that if a set of trades satisfies (6), then we can achieve differential privacy. We do this by first bounding the local sensitivity [DR⁺14] of the price impact vector $p_j(\pi, \Delta)$ as a function of Δ . This is done by reducing the problem to analyzing two different price trees (Appendix B). We make an analogue of smooth sensitivity [NRS07] that rounds a vector of trades to an integer lattice whose length is Δ_{\min} . These steps ensure that the maximum difference in price impact between neighboring sets of trades will be $O(\mu \log n)$. This immediately leads to achieving (ϵ, δ) -differential privacy, where $\epsilon = O(\mu \log n)$.

Using the composition property of differential privacy, we are able to compose these two mechanisms to achieve $(\mu \log n + \max_i \Delta_i, \delta)$ -differential privacy where $\delta = F^{-1}(O(\frac{1}{\epsilon}))$ and F^{-1} is the inverse Laplace CDF. While the constants can likely be improved, this suggests that permuting and splitting up trades is a simple and viable mechanism for adding differential privacy to CFMMs. Finally, note that in Appendix F we provide a convex program that can split up trades more efficiently than the Dirichlet mechanism of Theorem 3.4. This is likely useful to practitioners where randomness is a constrained resource (*e.g.*, on a blockchain).

4 Worst-Case Bounds and Path Deficiency

In this section, we'll explore if we can do better than the URE mechanism by analyzing the curvature of the mechanism and generalizing the previous work using Generic Chaining. Our goal will be to consider classes of mechanisms, \mathcal{F} , that can provide (ϵ, δ) -differential privacy for CFMMs and attempt to compute worst-case bounds. We first provide some necessary conditions that elements of such a class have to satisfy. We will also show that extending the results of §3.4 to the path-deficient (positive fee) case involves proving bounds over a class of functions \mathcal{F} . Finally, we'll investigate connections to private PAC learning which suggest that one cannot do significantly better than the URE unless curvature is dynamically adjusted.

4.1 Mechanism Curvature

Instead of directly working with a mechanism, can we say something about the set of all mechanisms that ensure that $|p^m(i) - p^t(i)| > \delta$ where $p^m(i)$ is the i th price of the mechanism and $p^t(i)$ is the non-private or true price? Using a curvature definition analogous to those of [AEC20], we can provide a simple bounds related to this question.

Note that bounds of the form $|p_i^m - p_i^t| > \delta$ involve bounding changes between two different price processes. Suppose that we define “curvatures” of the form

$$\begin{aligned}\kappa_t |\Delta_i| &< |p^t(i) - p^t(i-1)| < \mu_t |\Delta_i| \\ \kappa_m |\Delta_i| &< |p^m(i) - p^m(i-1)| < \mu_m |\Delta_i| \\ \kappa_{mt} |\Delta_i| &< |p^m(i) - p^t(i)| < \mu_{mt} |\Delta_i|\end{aligned}$$

First, let’s look at the difference between the mechanism price at time i and the true price at time $i - 1$:

$$\begin{aligned}|p^m(i) - p^t(i)| &= |(p^m(i) - p^t(i)) - (p^t(i-1) - p^t(i))| \\ &\geq |p^m(i) - p^t(i)| - |p^t(i) - p^t(i-1)| \\ &\geq (\kappa_m - \mu_{mt}) |\Delta_i|\end{aligned}$$

This says that we can ensure that the predictive value of previous price information on a trade cannot be resolved more than a multiplicative amount of $\kappa_m - \mu_{mt}$ times the trade size. In particular, $\kappa_m > \alpha + \mu_{mt}$ ensures that an adversary never has more than a precision α of information about the trade size. This provides a necessary condition in terms of mechanism curvature for a class \mathcal{F} of mechanisms to provide differential privacy bounds.

4.2 Path Deficiency

Any CFMM that has non-zero fees (*e.g.*, $\gamma = 1 - f < 1$) is path-deficient and has strictly negative expected value for round trip trades [AEC20]. Such CFMMs have price path $p^t(i)$ that are explicitly dependent on the trade ordering. Note that almost all CFMMs that are used in practice have non-zero fees to attract liquidity, so this is an important scenario to study. Previous work on path-deficient CFMMs has focused on analyzing how a particular price process (such as a geometric brownian motion) interacts with the expected returns from fees [EAC21]. Moreover, [AEC20, §2] illustrated that when fees are present $g^f(\Delta) = \gamma g(\gamma \Delta)$, where g^f is the price impact function with fees and g is the feeless price impact function. This suggests that we can analyze the path-dependent case by uniformly bounding the geometric parameters of §3.2 (*e.g.*, height and number of leaves) as a function of the fee.

Suppose that given a trade vector Δ , we have a bound of the form

$$\mathbf{E}_{\pi \in S_n} \left[\max_{i \in [n]} |p_f^\pi(\Delta) - p^\pi(\Delta)| \right] = O(\gamma^k) \quad (8)$$

In Appendix I, we compute a lower bound that allows one to prove such a bound for Uniswap (the most commonly used CFMM). Then we can bound the deviation in height between the set of trade and price trees (see Appendix B) as a function of γ and transfer path-independent returns to the path-deficient case with extra polylogarithmic terms in γ . Two ways of proving bounds of the form (8) are using generic chaining [Tal21, Ch. 3] and smoothed analysis [HRS20]. We discuss how this analysis can be applied to CFMMs in Appendix H.

4.3 Private PAC Learning and Adversarial Bounds

A number of recent results have shown that differentially private PAC learning and online learning are closely related. In particular, the finiteness of an integer-valued complexity measure known as the Littlestone dimension controls whether a particular algorithm can be learned in both an online and differentially private manner [ALMM19, BLM20]. The Littlestone dimension of a class of functions \mathcal{F} from $X \rightarrow Y$, $\text{LDim}(\mathcal{F})$, is defined as the maximum depth $d \in \mathbf{N}$ of a tree made up of sequences $x_1, \dots, x_d \in X$ such that there exists $f \in \mathcal{F}$ with $f(x_i) = y_i$ for every possible $y_i \in Y$. Consider the set $\mathcal{F}^\pi(\Delta)$ which is the set of all trees constructable from any permutation $\pi \in S_n$ for a fixed $\Delta \in \mathbf{R}^n$. The results of 3.2 show that $\text{LDim}(\mathcal{F}^\pi(\Delta)) = \Omega(\mu \log n)$. State-of-the-art results for blackbox constructions of online learners [GL21] show that the regret of a differentially private online learning algorithm is $O(2^{2^{\text{LDim}(\mathcal{F})}})$. This implies that the best online learners can do again the URE, in a blackbox manner, is $O(2^{2^\mu})$. This means that any algorithm that has non-zero curvature is unlikely to do asymptotically better than the URE. If it were possible to construct a polynomial time algorithm to privately PAC learn trades, then there would be significantly degraded privacy guarantees for users. However, this would require a mechanism for which $\text{LDim}(\mathcal{F}^\pi) = O(\log \log n)$, which appears unlikely except for constant-sum market makers that have $\mu = 0$. One other piece of evidence that Littlestone dimension is the correct complexity measure for CFMM privacy comes from the fact that the worst case instances for Littlestone dimension and CFMMs are thresholds (*cf.*, §3.2 and [ALMM19]).

5 Differentially (Non)-Private MEV Reduction

In previous sections, we assumed an honest sequencer who implements a differentially private mechanism for CFMM trades. We had explored this in the hopes that privacy might hinder MEV. Interestingly, it may be possible to prevent MEV by instantiating the “sequencer” and our mechanism on a public blockchain with access to a verifiable random function [MRV99], which exists on chains such as Polkadot [BCC⁺20] and Cosmos [Buc16].⁵ While this would not necessarily be differentially private — the noise is from public but unpredictable random coins — it could still prevent MEV. And the cost to users of doing so is modeled by the price impact analysis of §3.2. We see a similarity between this and results in machine learning relating differential privacy and fairness [DHP⁺12].

In practice, the majority of front-running and sandwich attacks are executed via maximal extractable value (MEV) auctions [BDKJ21]. These auctions separate the roles of sequencing (choosing an execution ordering) from *searching*,

⁵ We note, however, that the precise design in this paper is not immediately implementable — there are a number of practical and technical hurdles to overcome. These include, but are not limited to, determining how to allow applications to use randomness generated by consensus and figuring out how transaction submission and the pending transaction queue are affected by random orderings.

which is the process of finding the optimal front-run or sandwich transactions to maximize profit. Searchers bid for priority of transaction placement — they place a trade of size X before another user’s trade of size Y in order to front-run them. Sequencers collect these bids and construct a final transaction ordering based on which bids generate the maximum profit for them. Within this context, we can view the searchers as Eve (§3.1) — they do not know the final ordering and they can only affect it by placing a bid with the sequencer. Note, however, that when consensus-provided randomness is used to dictate the transaction order and sizing (e.g. via a verifiable random function), searchers match the description of Eve as they have a negligible edge over a coin flip in determining the order of trades. Even if searchers colluded with the sequencer to try to force a particular ordering, they would need to successfully execute a grinding attack against the VRF. In this paper, we implicitly have assumed that a VRF for which grinding attacks are hard to execute is used by the base protocol.

This observation demonstrates that our threat model is one in which searchers (not sequencers) are thwarted by the mechanisms of the subsequent sections. Given that $>50\%$ of CFMM extractable value from front-running is executed via the largest MEV auction, Flashbot [DOS], our model more closely models the real agents who are front-running users.

6 Conclusion

In this paper, we demonstrated that there exists a novel, practical mechanism for providing differential privacy to users of constant function market makers. This mechanism, unlike previous methods such as batching, has provable guarantees on the worst case price impact and strong privacy guarantees. As a number of new blockchain protocols implement CFMMs directly in their consensus mechanism, the randomness needed to execute this algorithm will become more plentiful and easier to source. Our analysis used novel techniques combining results from stochastic processes, concentration inequalities, and differential privacy. The results in this paper can likely be improved by providing tighter bounds on the minimal amount of noise needed to achieve (ϵ, δ) -differential privacy. Moreover, numerical studies of the utility loss (e.g., worsened price impact) would justify practical usage of URE on networks such as Osmosis [AO21] and Penumbra [dV21]. Finally, we note that differential privacy has been explored in path-independent prediction markets [FW17], where similar bounds to the ones found in this paper exist. These bounds utilize different proof techniques as prediction market makers do not directly translate to CFMMs (cf., [AC20, §3.2]). We note that a consequence of using this mechanism is that it likely provides better fairness for end users. Unlike fair ordering solutions [KZGJ20, KDL⁺21], our results provide economic guarantees on fairness for a particular application. Future work involves demonstrating that fairness is inherently present when a DeFi protocol can guarantee differential privacy.

7 Acknowledgements

The authors want to thank Ian Miers, Yi Sun, GaussianProcess, Tim Roughgarden, Kobi Gurkan, Dev Ojha, Henry de Valence, and the anonymous reviewers for helpful comments and feedback.

References

- [AAE⁺21] Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. Constant function market makers: Multi-asset trades via convex optimization. 2021.
- [ABC20] Louigi Addario-Berry and Benoît Corsini. The height of mallows trees. *arXiv preprint arXiv:2007.13728*, 2020.
- [ABF13] Louigi Addario-Berry and Kevin Ford. Poisson–dirichlet branching random walks. *The Annals of Applied Probability*, 23(1):283–307, 2013.
- [AC20] Guillermo Angeris and Tarun Chitra. Improved Price Oracles: Constant Function Market Makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91, New York NY USA, October 2020. ACM.
- [ACG⁺16] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 308–318, 2016.
- [AEC20] Guillermo Angeris, Alex Evans, and Tarun Chitra. When does the tail wag the dog? curvature and market making. *arXiv preprint arXiv:2012.08040*, 2020.
- [AEC21a] Guillermo Angeris, Alex Evans, and Tarun Chitra. A note on bundle profit maximization. 2021.
- [AEC21b] Guillermo Angeris, Alex Evans, and Tarun Chitra. A note on privacy in constant function market makers. *arXiv preprint arXiv:2103.01193*, 2021.
- [Alb19] Mélisande Albert. Concentration inequalities for randomly permuted sums. In *High Dimensional Probability VIII*, pages 341–383. Springer, 2019.
- [ALMM19] Noga Alon, Roi Livni, Maryanthe Malliaris, and Shay Moran. Private pac learning implies finite littlestone dimension. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 852–860, 2019.
- [AO21] Sunny Agrawal and Dev Ojha. Vision for osmosis, May 2021.
- [BCC⁺20] Jeff Burdges, Alfonso Cevallos, Peter Czaban, Rob Habermeier, Syed Hosseini, Fabio Lama, Handan Kiliç Alper, Ximin Luo, Fatemeh Shirazi, Alstair Stewart, et al. Overview of polkadot and its design considerations. *arXiv preprint arXiv:2005.13456*, 2020.
- [BDKJ21] Kushal Babel, Philip Daian, Mahimna Kelkar, and Ari Juels. Clockwork finance: Automated analysis of economic security in smart contracts. *arXiv preprint arXiv:2109.04347*, 2021.
- [BJP12] Tamara Broderick, Michael I Jordan, and Jim Pitman. Beta processes, stick-breaking and power laws. *Bayesian analysis*, 7(2):439–476, 2012.
- [BLM20] Mark Bun, Roi Livni, and Shay Moran. An equivalence between private classification and online prediction. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 389–402. IEEE, 2020.

- [BNSV15] Mark Bun, Kobbi Nissim, Uri Stemmer, and Salil Vadhan. Differentially private release and learning of threshold functions. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 634–649. IEEE, 2015.
- [Buc16] Ethan Buchman. *Tendermint: Byzantine fault tolerance in the age of blockchains*. PhD thesis, 2016.
- [CGKM19] Rachel Cummings, Varun Gupta, Dhamma Kimpara, and Jamie Morgenstern. On the compatibility of privacy and fairness. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pages 309–315, 2019.
- [Cha07] Sourav Chatterjee. Stein’s method for concentration inequalities. *Probability Theory and Related Fields*, 138(1):305–321, 2007.
- [CJK⁺18] Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data*, pages 1655–1658, 2018.
- [CXZ20] Shumo Chu, Qiudong Xia, and Zhenfei Zhang. Manta: Privacy preserving decentralized exchange. *IACR Cryptol. ePrint Arch.*, 2020:1607, 2020.
- [DHP⁺12] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- [Dia88] Persi Diaconis. Metrics on groups, and their statistical uses. In *Group representations in probability and statistics*, pages 102–130. Institute of Mathematical Statistics, 1988.
- [DOS] Phil Daian, Alex Obadia, and Luke Setters. Mev explore.
- [DR⁺14] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407, 2014.
- [dV21] Henry de Valence. Sealed-bid batch auctions, 2021.
- [Dwo19] Cynthia Dwork. Differential privacy and the us census. In *Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 1–1, 2019.
- [EAC21] Alex Evans, Guillermo Angeris, and Tarun Chitra. Optimal fees for geometric mean market makers. *arXiv preprint arXiv:2104.00446*, 2021.
- [ER75] Paul Erdos and Pál Révész. On the length of the longest head-run. *Topics in information theory*, 16:219–228, 1975.
- [FW17] Rafael Frongillo and Bo Waggoner. Bounded-loss private prediction markets. *arXiv preprint arXiv:1703.00899*, 2017.
- [GL21] Noah Golowich and Roi Livni. Littlestone classes are privately online learnable. *arXiv preprint arXiv:2106.13513*, 2021.
- [GWH⁺21] Parham Gohari, Bo Wu, Calvin Hawkins, Matthew Hale, and Ufuk Topcu. Differential privacy on the unit simplex via the dirichlet mechanism. *IEEE Transactions on Information Forensics and Security*, 16:2326–2340, 2021.
- [HRS20] Nika Haghtalab, Tim Roughgarden, and Abhishek Shetty. Smoothed analysis of online and differentially private learning. *arXiv preprint arXiv:2006.10129*, 2020.
- [KDL⁺21] Mahimna Kelkar, Soubhik Deb, Sishan Long, Ari Juels, and Sreeram Kannan. Themis: Fast, strong order-fairness in byzantine consensus. *Cryptology ePrint Archive*, 2021.
- [KOR19] Sep Kamvar, Marek Olszewski, and Rene Reinsberg. Celso: A multi-asset cryptographic protocol for decentralized social payments. *DRAFT version*

- 0.24 https://storage.googleapis.com/celo-whitepapers/Celo_A_Multi_Asset_Cryptographic_Protocol_for_Decentralized_Social_Payments.pdf, 2019.
- [KOV15] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. The composition theorem for differential privacy. In *International conference on machine learning*, pages 1376–1385. PMLR, 2015.
- [KZGJ20] Mahimna Kelkar, Fan Zhang, Steven Goldfeder, and Ari Juels. Order-fairness for byzantine consensus. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020*, pages 451–480, Cham, 2020. Springer International Publishing.
- [MA17] Olivier Marchal and Julyan Arbel. On the sub-gaussianity of the beta and dirichlet distributions. *Electronic Communications in Probability*, 22:1–14, 2017.
- [Mar21] Fernando Martinelli. The crypto cinematic universe crossover event of the summer: Balancer-gnosis-protocol (bgp), Apr 2021.
- [MRV99] Silvio Micali, Michael Rabin, and Salil Vadhan. Verifiable random functions. In *40th annual symposium on foundations of computer science (cat. No. 99CB37039)*, pages 120–130. IEEE, 1999.
- [MSS20] Amani Moin, Kevin Sekniqi, and Emin Gun Sirer. Sok: A classification framework for stablecoin designs. In *International Conference on Financial Cryptography and Data Security*, pages 174–197. Springer, 2020.
- [Nad07] Saralees Nadarajah. The linear combination, product and ratio of laplace random variables. *Statistics*, 41(6):535–545, 2007.
- [NRS07] Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. Smooth sensitivity and sampling in private data analysis. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 75–84, 2007.
- [Pow21] Benjamin Powers. Secretswap is the secret network’s answer to defi privacy, Feb 2021.
- [QZG21] Kaihua Qin, Liyi Zhou, and Arthur Gervais. Quantifying blockchain extractable value: How dark is the forest? *arXiv preprint arXiv:2101.05511*, 2021.
- [QZLG20] Kaihua Qin, Liyi Zhou, Benjamin Livshits, and Arthur Gervais. Attacking the defi ecosystem with flash loans for fun and profit. *arXiv preprint arXiv:2003.03810*, 2020.
- [Ree03] Bruce Reed. The height of a random binary search tree. *Journal of the ACM (JACM)*, 50(3):306–332, 2003.
- [Tal14] Michel Talagrand. *Upper and lower bounds for stochastic processes: modern methods and classical problems*, volume 60. Springer Science & Business Media, 2014.
- [Tal21] Michel Talagrand. Upper and lower bounds for stochastic processes: modern methods and classical problems (2nd edition, preprint). 2021.
- [XYW19] Depeng Xu, Shuhan Yuan, and Xintao Wu. Achieving differential privacy and fairness in logistic regression. In *Companion Proceedings of The 2019 World Wide Web Conference*, pages 594–599, 2019.
- [ZQT⁺20] Liyi Zhou, Kaihua Qin, Christof Ferreira Torres, Duc V Le, and Arthur Gervais. High-frequency trading on decentralized on-chain exchanges. *arXiv preprint arXiv:2009.14021*, 2020.

A Differential Privacy Results

We implicitly use a number of differential privacy results on composition and provide them here for convenience. First we note the serial composition theorem:

Theorem 2 (Composition theorem 3.16 [DR⁺14]). *Let $\mathcal{A}_1, \dots, \mathcal{A}_n$ be a sequence of (ϵ_i, δ_i) algorithms such that $\text{Range } \mathcal{A}_i \subseteq \text{Dom } \mathcal{A}_{i+1}$. Then the composition $\mathcal{A}_n \circ \dots \circ \mathcal{A}_1$ is $(\sum_{i=1}^n \epsilon_i, \sum_{i=1}^n \delta_i)$ -differentially private*

Secondly, we note the parallel composition theorem

Theorem 3. *Let $\mathcal{A}_1, \dots, \mathcal{A}_n$ be algorithms whose domains (databases) are independent and each algorithm is (ϵ_i, δ_i) -differentially private. Then $(\mathcal{A}_1, \dots, \mathcal{A}_n)$ is $\max_i \epsilon_i$ differentially private*

Finally, we note that the serial composition rule can be improved from $(\sum_i \epsilon_i, \sum_i \delta_i)$ to $(n\epsilon^2 + \epsilon\sqrt{n \log(1/\tilde{\delta})}, n\delta + \tilde{\delta})$ where $\tilde{\delta} = O(n\delta)$ if $\epsilon_i = \epsilon, \delta_i = \delta$ for all i [KOV15]. We will not need to use this result, only the generic composition rules. However, it is possible that one can improve our constants using results such as this.

B Price Tree Height is close to Trade Tree Height

Suppose that we have an admissible trade vector $\Delta = (\Delta_1, \dots, \Delta_n) \in \mathcal{A}_\varphi$. Given $\pi \in S_n$, we can write a sequence of prices in terms of the price impact function:

$$p_j(\pi) = g\left(\sum_{i=1}^j \Delta_{\pi(i)}\right)$$

We generate a random binary tree from the price vector by uniformly sampling $j \sim [n]$ and making $p_j(\pi)$ the root before inserting the remaining prices sequentially as per π . Under this framework, we have

$$\begin{aligned} \mathbf{E}_{\pi \sim S_n} \left[\max_j p_j(\pi) \right] &\leq \mathbf{E}_{j \sim [n]} [p_j(\pi)] + \max_i |p_i(\pi) - p_{i-1}(\pi)| \mathbf{E}_{\pi \sim S_n} [\text{height}(T(p_j(\pi)))] \\ &\leq \mathbf{E}_{j \sim [n]} [p_j(\pi)] + \mu(\max_i \Delta_i) \mathbf{E}_{\pi \sim S_n} [\text{height}(T(p_j(\pi)))] \end{aligned}$$

We can later remove this constraint by adding a small amount of noise to each entry, which will make the entries unique a.s.. Note that the height of the tree generated by P_j represents the number of trades in the longest sequential deviation from the mean price. Let's consider when the trade tree and price tree differ in branching. On average, this occurs when the j th price $p_{\pi(j)}$ is a left branch whereas the $j+1$ st price $p_{\pi(j+1)}$ is a right branch, but both trades $\Delta_{\pi(j)}, \Delta_{\pi(j+1)}$ are left branches. When this happens, the price tree has an average height that is 1 less than the trade tree.

We will first illustrate this when the first two elements of the permutation after the pivot (which is random) differ from the expected pivot value. Explicitly, suppose that we have

$$p_{\pi(2)} - \frac{1}{n} \sum_{i=1}^n p_{\pi(i)} < 0 \qquad p_{\pi(3)} - \frac{1}{n} \sum_{i=1}^n p_{\pi(i)} > 0$$

Using curvature bounds, the first equation gives

$$0 \geq p_{\pi(2)} - \frac{1}{n} \sum_{i=1}^n p_{\pi(i)} \geq \kappa \Delta_{\pi(2)} - \frac{\mu}{n} \sum_{i=1}^n \Delta_i$$

Similarly, the second equation gives

$$0 \leq p_{\pi(3)} - \frac{1}{n} \sum_{i=1}^n p_{\pi(i)} \leq \mu \Delta_{\pi(3)} - \frac{\kappa}{n} \sum_{i=1}^n \Delta_i$$

which when combined gives

$$\Delta_{\pi(2)} \leq \frac{\mu}{\kappa} \left(\frac{1}{n} \sum_{i=1}^n \Delta_i \right) = \eta_+ \quad (9)$$

$$\Delta_{\pi(3)} \geq \frac{\kappa}{\mu} \left(\frac{1}{n} \sum_{i=1}^n \Delta_i \right) = \eta_- \quad (10)$$

Let ρ_i be as in (2) and let $\bar{\Delta} = \frac{1}{n} \sum_{i=1}^n \Delta_i$. On the other hand, suppose that $\rho_2(\pi) - \bar{\Delta}(\pi), \rho_3(\pi) - \bar{\Delta}(\pi)$ are both greater than zero (*e.g.*, they are both left nodes of their parent). This implies that $\Delta_{\pi(2)} + \Delta_{\pi(3)} \geq \frac{1}{n} \sum_{i=1}^n \Delta_i$. This means that we can only end up in a state where $\text{height}(T(\rho_j(\pi))) > \text{height}(T(p_j(\pi)))$ if the trades are within the interval $[\eta_-, \eta_+]$. For instance, when the drift $\frac{1}{n} \sum_{i=1}^n \Delta_i = 0$, then interval has size zero (its a mean-reverting set of trades) and we never enter this error condition. This matches intuition: if there's a lot of drift in the trades, then we shouldn't expect our price and trade vectors to 'sort' the same way. In particular, the higher the curvature of the CFMM, the less drift we can tolerate because large trades cause more noticeable price impact. The length of the interval $[\eta_-, \eta_+]$ is

$$\left(\frac{\mu}{\kappa} - \frac{\kappa}{\mu} \right) \left(\frac{1}{n} \sum_{i=1}^n \Delta_i \right)$$

Note that we can recurse the above argument as we go down the tree and get a set of intervals $I_1 = [\eta_-(1), \eta_+(1)], I_2 = [\eta_-(2), \eta_+(2)], \dots, I_n = [\eta_-(n), \eta_+(n)]$. Performing the same calculation as above yields

$$\eta_-(i) = \frac{\kappa}{\mu} \left(\frac{1}{n-i} \sum_{i=i}^n \Delta_{\pi(i)} \right) \quad \eta_+(i) = \frac{\mu}{\kappa} \left(\frac{1}{n-i} \sum_{i=i}^n \Delta_{\pi(i)} \right)$$

Given that the maximum interval size is μM is the max trade size for which curvature is valid), we can use this to bound the probability p_j that vertex j has a height difference between the trade and price trees. This probability is upper bounded by ratio of the length of I_j and the interval length μM , *e.g.*, $p_j \leq \frac{|I_j|}{\mu M}$. We can upper bound the interval length by the maximum mean-drift

subsequence:

$$|I_j| \leq \left(\frac{\mu}{\kappa} - \frac{\kappa}{\mu} \right) \left(\max_{J \subset [n]} \frac{1}{|J|} \sum_{j \in J} \Delta_j \right)$$

Define $R^*(\Delta) = \max_{J \subset [n]} \frac{1}{|J|} \sum_{j \in J} \Delta_j$. Finally, performing a union bound gives an upper bound on the probability p_{diff} of the heights of the trade tree and price tree different

$$p_{\text{diff}} \leq \sum_{j=1}^n p_j = n \left(\frac{1}{M\kappa} - \frac{\kappa}{\mu^2 M} \right) R^*(\Delta) \quad (11)$$

If this quantity is sufficiently small (*e.g.*, we have tight curvature bounds), then bounds on the trade tree transfer to the price tree with high probability. For the rest of the paper, we will assume that (11) is sufficiently small. We note that fee adjustments and curvature adjustments are intricately related [AEC20, §3] and in practice, this can be enforced by dynamic updates to a CFMM curve.

C Proof of Claim 3.3

Suppose that $\xi_i \sim_{iid} \text{Lap}(a, b)$. We need to analyze the distribution of $\xi_i - \frac{\mu}{\kappa} \xi_j$. Recall that if $X \sim \text{Lap}(a, b)$ then $kX \sim \text{Lap}(ka, |k|b)$. Therefore we are trying to bound the distribution of $Z(a, b) = X + Y$ where $X \sim \text{Lap}(a, b)$, $Y \sim \text{Lap}(-\frac{\mu}{\kappa}a, \frac{\mu}{\kappa}b)$. In particular, given $\delta < 0$ we want to choose a, b such that

$$F_Z(k) \leq \mathbf{Prob}[X + Y \leq k] \leq \delta$$

where $k = c_{\min} + \left| \min_{i,j} \Delta_i - \frac{\kappa}{\mu} \Delta_j \right|$. Nadarajah [Nad07, Theorem 1] explicitly computes the CDF $F_{Z(a,b)}(k)$ and shows that it is monotone, continuous, and differentiable in a, b except at one value of k for all a, b . Moreover, it is supported on the entire real line. Therefore, $\exists a^*$ such that $F_{Z(a^*, |a^*|)}(k) = \delta$.

D Proof of Claim 3.4

Our proof works by differentially privately sampling a probability distribution $\pi \sim \text{Dir}(\vec{1})$ multiple times using the mechanism of [GWH⁺21]. The Dirichlet mechanism on k nodes $\mathcal{M}_D^{(k)}(\pi)$ samples a Dirichlet distribution centered at π , where $\pi \in P_k = \{x \in \mathbf{R}^k : \sum_i x_i = 1, x_i \geq 0\}$. One can think of it as sampling a increment $d\pi$, adding it to π and renormalizing. First, we reproduce a theorem on differentially private Dirichlet sampling

Theorem 4 ([GWH⁺21], **Theorem 1, Corollary 1**). *The Dirichlet mechanism $\mathcal{M}_D^{(k)}(\pi)$ achieves (ϵ, δ) -differential privacy where $\epsilon = O(k(1 + \log(o(k))))$ and $\delta = 1 - \min_{\pi} \mathbf{Prob}[\mathcal{M}_D^k(\pi) - \pi > \Omega(\epsilon)]$*

Define the vector $\eta(\Delta)$ as follows:

$$\eta(\Delta) = \left(\left\lceil \frac{\Delta_1}{\Delta_{\min}} \right\rceil, \dots, \left\lceil \frac{\Delta_n}{\Delta_{\min}} \right\rceil \right)$$

Each coordinate represents rounding each trade to an integer lattice with width Δ_{\min} . Define $S_k = \{i : \eta(\Delta)_i > k\}$ and $S_k^c = [n] - S_k$. For each $j \in S_k$, privately sample $\pi \sim \text{Dir}(\vec{1})$ where $\vec{1} = (1, \dots, 1) \in \mathbf{R}^{\eta(\Delta)_j}$. Let $\hat{\Delta}_{j,k} = \Delta_j \pi_k$ with $\sum_k \hat{\Delta}_{j,k} = \Delta_j$. We can view each Dirichlet sample π as providing a mechanism for splitting the trade Δ_j . Our goal is to find $k \in \mathbf{N}$ such that the following two conditions hold

1. $\text{height}(T(\Delta_{S_k^c})) = \Theta(\log n)$
2. $\text{height}(T(\hat{\Delta}_{j,k})) = \Theta(\log \eta_j)$ with high probability

We can show that the latter condition holds with high probability when the distribution sampled is Dirichlet centered at the centroid $(\frac{1}{n}, \dots, \frac{1}{n})$. Constructing a partial sum tree from a Dirichlet sample is the same as drawing a sample from a Poisson-Dirichlet branching random walk [ABF13]. These walks satisfy $\mathbf{Prob}[\text{height}(T(\hat{\Delta}_{j,k})) - c \log \eta(\Delta)_j \geq k] = O(e^{-k})$ for a universal constant c [ABF13, Corollary 1.3]. Therefore, the probability that all of the Dirichlet constructed trees $T(\hat{\Delta}_{j,k})$ have height greater than $c \log \eta(\Delta)_j$ is

$$\mathbf{Prob} \left[\exists j \in S_k \mid \text{height}(T(\hat{\Delta}_{j,k})) - c \log \eta_j \geq c' \log \eta_j \right] \leq \left(\frac{|S_k|}{\eta_j^{c'}} \right)$$

which directly follows from the independent sampling from the private Dirichlet distribution and inclusion-exclusion. Therefore, with probability $p^* = 1 - \frac{|S_k|}{\delta_j^{c'}}$, we have the maximum height of a tree constructed from all $|S_k|$ vectors $\hat{\Delta}_{j,k}$ is

$$\sum_{j \in S_k} \log \eta_j \leq |S_k| \max_j \log \eta_j$$

which under our assumptions is $O(\log n)$. Our claim about differential privacy then follows immediately from Theorem 4.

E Proof of Claim 3.4

We will prove differential privacy by using the smooth sensitivity framework of [NRS07]. First, we will recall definitions and introduce preliminaries on this framework before specializing it to SURE. Smooth sensitivity places an upper bound on the local sensitivity of a function f , which is defined as

$$LS_f(x) = \max_{d(x,y) \leq 1} |f(x) - f(y)|$$

Note that unlike the global sensitivity, which is used in the generic Laplace mechanism [DR⁺14], the local sensitivity depends on the particular input x .

Often times, it is too difficult to get uniform bounds on local sensitivity and instead it is easier to use a smooth proxy. A β -smooth upper bound $S : \mathbf{Dom} f \rightarrow \mathbf{R}$ for $LS_f(x)$ satisfies $S(x) \geq LS_f(x)$ for all $x \in \mathbf{Dom} f$ and $S(x) \leq e^\beta S(y)$ for all $x, y \in \mathbf{Dom} f$ with $d(x, y) = 1$. We are now in a position to recall two results of Nissim, et. al:

Theorem 5 ([NRS07], **Lem. 2.6**). *Let h be an (α, β) -admissible noise probability density function and let $Z \sim h$. For a function $f : D^n \rightarrow \mathbf{R}^d$, let S be a β -smooth upper bound in the local sensitivity of f , then $\mathcal{A}(x) = f(x) + \frac{S(x)}{\alpha}Z$ is (ϵ, δ) -differentially private.*

Theorem 6 ([NRS07], **Lem. 2.9**). *For $\epsilon, \delta \in (0, 1)$, the d -dimensional Laplace distribution, $h(z) = 2^{-d}e^{-\|z\|_1}$ is (α, β) -admissible with $\alpha = \frac{\epsilon}{2}$, $\beta = \frac{\epsilon}{2\rho_\delta(Y)(\|Z\|_1)}$ where $\rho_\delta(Y)$ is the $1 - \delta$ quantile of Y*

Combined, these results illustrate that if we can construct a β -smooth upper bound, we can immediately construct a Laplace mechanism that achieves (ϵ, δ) -differential privacy. Section 3 of [NRS07] provides a mechanism for computing a β -smooth upper bound by first defining the sensitivity at distance k ,

$$LS_f^k(x) = \max_{\substack{y \in \mathbf{Dom} f \\ d(x, y) \leq k}} LS_f(x)$$

A β -smooth upper bound on local sensitivity is defined as,

$$S_{f, \beta}(x) = \max_{k \in \{0, 1, \dots, n\}} e^{-k\beta} LS_f^k(x)$$

Therefore, we need to construct a function f that represents price impact and compute an analogue of local sensitivity.

For a differentially private CFMM, we want to minimize the worst case price impact in a neighborhood of a trade Δ . We define $f(\Delta)$ as

$$f(\Delta) = \max_{j \in [n]} p_j(\Delta)$$

Now we need to modify the definition of local sensitivity to account for trade admissibility and discretization. Normally, local sensitivity is defined for discrete spaces where the distance d is taken to be the Hamming metric. We can discretize our trade space in terms of Δ_{\min} . Recall that we ensure that $\Delta_{\min} > 0$ by adding Laplace noise to all trades (whose parameter will be tuned in accordance with the above theorem). Note that moving to such a discretization simply changes our choice of β . Using this definition, we can define the local trade sensitivity as

$$TS_f^k(\Delta) = \sup_{\substack{\Delta' \in \mathbf{Dom} f \cap \mathcal{A}(R) \\ d(\Delta, \Delta') \leq k\Delta_{\min}}} |f(\Delta) - f(\Delta')|$$

where $\mathcal{A}(R)$ is the set of admissible trades. From the results of §3.2, we know that $TS_f^k(\Delta) = O(k\mu(\max_i \Delta) \log n)$ since the depth of the tree quantifies the largest

price impact. In particular, each element Δ'_i such that $|\Delta_i - \Delta'_i| > \Delta_{\min}$ can cause price impact of at most $\mu(\max_i \Delta) \log n$ and we can add these independently over the at most k coordinates that have prices changed by more than Δ_{\min} . We can define an analogous smooth sensitivity bound,

$$\tilde{S}_{f,\beta}(x) = \max_{\ell} e^{-\ell\beta} T S_f^{\ell}(\Delta) = \max_{\ell} e^{-\ell\beta} \ell \mu(\max_i \Delta) \log n$$

This is minimized when $\ell = \frac{1}{\beta}$, giving

$$\tilde{S}_{f,\beta}(x) = \frac{\mu}{e\beta} (\max_i \Delta) \log n$$

Therefore, provided that a) the partial sum tree has height $O(\log n)$ b) the noise added ensures that $\Delta_{\min} > 0$, and c) the noise is rescaled by $\frac{2\tilde{S}_{f,\beta}(x)}{\epsilon}$, we achieve differential privacy.

Note that in particular, our bound depends on $\max_i \Delta$ and the curvature upper bound. By splitting trades using Claim 3.4, we reduce $\max_i \Delta$ and can ensure that the noise added is reasonable. Moreover, as we saw, without splitting trades, we run into issues with trades of the form $(T, 1, \dots, 1)$. Note that algorithms that try to learn where the trade T occurs (after applying a permutation π) is equivalent to privately learning threshold functions [BNSV15, ALMM19].

F Convex Trade Splitting

When we are considering CFMM arbitrage, it can be shown that a necessary condition for stability is path-deficiency. Path deficiency ensures that no rational trader (*e.g.*, profit optimizing) is incentivized to split a desired trade size Δ into two trades $\Delta_1 + \Delta_2 = \Delta$. However, if a trader also desires privacy, splitting up trades can become necessary. To see why, consider a trader who makes a trade of size T and a sequence of trades $\Delta = (T, 1, \dots, 1) \in \mathbf{R}^{T+1}$. Using curvature, we know that the price impact is at least κT after a trade of size T and of size κ after each trade of size 1. This means that an adversary can easily discern where my trade is, even if Δ is randomly permuted due to the T times larger price impact. Therefore, splitting up the trade of size T into trades close to size 1 will make it hard for an adversary to reconstruct the total trade size.

Our goal is to split up trades such that the probability of an adversary detecting the position of a single trade is small relative to the curvature. Suppose that a trade Δ_1 is split into trades $\Delta'_1, \dots, \Delta'_j$ and let $\tilde{\Delta} = (\Delta'_1, \dots, \Delta'_j, \Delta_2, \dots, \Delta_n)$. A splitting adversary is a binary classifier $\ell(\Delta, \tilde{\Delta})$ that returns 1 if $\Delta \in \{\Delta'_1, \dots, \Delta'_j\}$ and 0 otherwise. We say that a splitting mechanism is (δ, ϵ) indistinguishable if

$$\mathbf{Prob} \left[\left| \frac{1}{n} \sum_{i=1}^n \ell(\tilde{\Delta}_i, \tilde{\Delta}^{\pi}) - \frac{j}{n} \right| < \epsilon \right] < \delta$$

over some suitable set of splitting classifiers. The inequalities in Appendix G can directly be used to prove that this holds for the L^2 norm.

However, path-deficiency implies that splitting trades will cost a user an extra fee. This trade-off between best execution price and privacy can be explored via a simple, convex objective function that trades off price impact vs. improved privacy via splitting. Recall that the L^2 norm strictly decreases under splitting, *e.g.*,

$$\begin{aligned} \|(\Delta_1, \dots, \Delta_n)\|_2^2 &= \sum_{i=1}^n \Delta_i^2 = \Delta_1^2 + \sum_{i=2}^n \Delta_i^2 \\ &= a\Delta_1^2 + (1-a)\Delta_1^2 \sum_{i=2}^n \Delta_i^2 > a^2\Delta_1^2 + (1-a)^2\Delta_1^2 + \sum_{i=2}^n \Delta_i^2 \\ &= \|(a\Delta_1, (1-a)\Delta_1, \dots, \Delta_n)\|_2^2 \end{aligned}$$

where $a \in (0, 1)$ represents the splitting fraction.

This property allows us to quantify the privacy benefit to splitting trades, as the more minimal the L^2 norm, the less noise that is needed to ensure that the random binary tree has height $\Theta(\log n)$ and $\Omega(n)$ leaves. In particular, the Cauchy and Gaussian mechanisms for differential privacy utilize distributions whose variances are proportional to the L^2 norm.

Given that we want to minimize price impact while maximizing the amount of trade splitting necessary for indistinguishable, we construct a convex optimization problem. Define the function f as:

$$f(\Delta_1, \dots, \Delta_n) = \sum_{i=1}^n \gamma g\left(\gamma \sum_{j=1}^i \Delta_j\right) + \eta \sum_{i=1}^n \Delta_i^2$$

The first term in f represents an upper bound on the price impact and the second term represents the L^2 splitting term. Our goal is to minimize f over sequences of trades $(\Delta_1, \dots, \Delta_k) \in \sqcup_{i=1}^{\infty} \mathbf{R}^i$ such that $\sum_{i=1}^k \Delta_i = \Delta^*$, *e.g.*,

$$\begin{aligned} &\text{minimize} && f(\Delta_1, \dots, \Delta_n) \\ &\text{subject to} && \Delta_1 + \dots + \Delta_n = \Delta^* \end{aligned} \tag{12}$$

Using curvature bounds, we can construct a simple descent algorithm to solve this. Firstly, note that the definition of curvature yields

$$\kappa\gamma^2 \sum_{i=1}^n \sum_{j=1}^i \Delta_i \leq f(\Delta_1, \dots, \Delta_n) - \eta \sum_{i=1}^n \Delta_i^2 \leq \mu\gamma^2 \sum_{i=1}^n \sum_{j=1}^i \Delta_i$$

Furthermore, note that we can rewrite the double sum as

$$\sum_{i=1}^n \sum_{j=1}^i \Delta_i = \sum_{i=1}^n (n-i+1)\Delta_i$$

Next, note that we can upper bound the split function, $f(a\Delta_1, (1-a)\Delta_1, \dots, \Delta_n)$ as

$$\begin{aligned} f(a\Delta_1, (1-a)\Delta_1, \dots, \Delta_n) &\leq \mu\gamma^2 \left((n+1)a\Delta_1 + n(1-a)\Delta_1 + \sum_{i=2}^n (n-i+2)\Delta_i \right) \\ &\quad + \eta \left(a^2\Delta_1^2 + (1-a)^2\Delta_1^2 + \sum_{i=2}^n \Delta_i^2 \right) \\ &= \mu\gamma^2 \left((n+a)\Delta_1 + \sum_{i=2}^n (n-i+1)\Delta_i + \Delta^* \right) \\ &\quad + \eta \left(a^2\Delta_1^2 + (1-a)^2\Delta_1^2 + \sum_{i=2}^n \Delta_i^2 \right) \end{aligned}$$

Combining these gives the following

$$\begin{aligned} f(\Delta_1, \dots, \Delta_n) - f(a\Delta_1, (1-a)\Delta_1, \dots, \Delta_n) &\geq \gamma^2(\kappa - \mu) \sum_{i=2}^n (n-i+1)\Delta_i - \Delta^* \\ &\quad - \mu\gamma^2(n+a)\Delta_1 + \eta\Delta_1^2(1-a^2 - (1-a)^2) \end{aligned} \tag{13}$$

Maximize the right-hand side in a provide a mechanism for deciding whether to split trade Δ_1 . Optimizing over a yields

$$a^* = \max \left(\frac{1}{2} - \frac{\mu\gamma^2}{4\eta\Delta_1}, 0 \right)$$

If we substitute a^* into (1) and the right-hand side is position, we split the trade Δ_1 into two trades of size $a^*\Delta_1$ and $(1-a^*)\Delta_1$.

G Splitting Trades: Concentration

Chatterjee proved a concentration bound using Stein's method that provides intuition as to why splitting trades improves the effectiveness of SURE. Theorem 7 shows that the variance of concentration around the mean for a randomly permuted sum is linear in the expected value.

Theorem 7 ([Cha07], Prop. 1.1). *Let $\{a_{i,j}\}_{1 \leq i,j \leq n}$ be a collection of numbers from $[0, 1]$. Let $X = \sum_{i=1}^n a_{i,\pi(i)}$ where $\pi \sim S_n$ uniformly. Then*

$$\mathbf{Prob}[|X - \mathbf{E}[X]| \geq t] \leq 2 \exp \left(-\frac{t^2}{4\mathbf{E}[X] + 2t} \right) \tag{14}$$

Note that unlike Bernstein-like inequalities there is no direct dependence on n . Moreover, unlike Talagrand-like inequalities [Tal21], we do not have terms dependent on ϵ -nets. If we let $t = k\mathbf{E}[X]$, we have

$$\exp \left(-\frac{t^2}{4\mathbf{E}[X] + 2t} \right) = \exp \left(-\frac{k^2\mathbf{E}[X]}{2k+4} \right) \leq \exp(-k\mathbf{E}[X])$$

For positive trade sizes, this implies that if we can split big trades into smaller trades (which reduces in turn reduces $\mathbf{E}[X]$) we can achieve the sufficient condition. More specifically, suppose that $a_{i,j} = \Delta_j - \frac{\kappa}{\mu} \Delta_i$. Then $X = \sum_{i=1}^n a_{i,\pi(i)}$ is the upper bound from (6) and the theorem claims that reducing the maximum will reduce the variance of SURE's utility.

We also note that better asymptotic results exist for non-negative sums:

Theorem 8 ([Alb19], Corollary 2.2). *Let a_{ij} be a connection of any real numbers and $\pi \sim S_n$ as uniform random permutation. Let $Z_n = \sum_{i=1}^n a_{i,\pi(i)}$. Then for all $x > 0$*

$$\mathbf{Prob}(|Z_n - \mathbf{E}[Z_n]| \geq t) \leq 16e^{1/16} \exp\left(\frac{-t^2}{256(\mathbf{Var}[Z_n] + \max_{i,j} |a_{ij}|t)}\right)$$

This bound explicitly includes a maximum term, directly justifying the improvement to SURE provided by splitting trades.

H Path Dependency and Generic Chaining

Suppose that we want to try to find the worst case price deviation given that we have fees, $\gamma < 1$. If we define $X_j = p^\pi(i) - p(i)$, then we want to study the extremal behavior of this process, albeit without being able to directly bound price impact using methods from §3.2. We will be most interested in the behavior of the random variable $X^* = \max_j X_j$, which quantifies the worse execution price received by a user under this mechanism. To do this, we will utilize the theory of empirical processes. Roughly speaking, one can show that for a metric space (T, d) , $\mathbf{E} \sup_{t \in T} X_t = \Theta(\text{Diam}(T) \sqrt{\log \text{card} T})$ by looking at simple bounds for empirical processes [Tal14, Tal21]. Our goal is to define a metric space T_γ that depends on fees and such that S_n acts faithfully on T_γ . We want the action to be faithful because that will be equivalent to the condition of unique elements of the form $|\Delta_i - \frac{\mu}{\kappa} \Delta_j|$. We can then attempt to bound, using chaining arguments, the worst case price deviation.

Chaining bounds rely on tail bounds on increments, *e.g.*, showing that for some metric d on our space T_γ , we have the following two conditions:

$$\forall u > 0, \mathbf{Prob}[|X_s - X_t| \geq u] \leq 2 \exp\left(-\frac{u^2}{2d(s,t)^2}\right) \quad (15)$$

$$\exists u > 0, \sum_{s \in T} \mathbf{Prob}[X_s \geq u] \geq 1 \quad (16)$$

In our case, we need to construct a metric space that takes advantage of our trading function curvature and the randomness induced by the choice of permutation.

Our goal is to construct a metric on S_n that depends on both φ . We need to construct metric $d_{\varphi, \rho_0, \Delta} : S_n \times S_n \rightarrow \mathbf{R}_+$ that we can use to find a formula like

eq. (15). A natural metric to construct is the raw price differences:

$$d_{\varphi, \rho_0, \Delta}(\pi_1, \pi_2) = \sum_{i=1}^n |p_{\pi_1(i)}^t - p_{\pi_2(i)}^t|$$

Note that if we took an infimum over one of the two permutations, we arrive at the Wasserstein distance. Suppose we have $d_{\varphi, \rho_0, \Delta}(\pi_1, \pi_2) \leq f(\varphi, \rho_0, \Delta)d(\pi_1, \pi_2)$ for some natural metric on the symmetric group (e.g., Mallows metric [Dia88]). Moreover, suppose there exists $\kappa > 0$ such that $\mathbf{Prob}[X_s \geq \sqrt{\log n}(\kappa + \sum_i \Delta_i)] \geq 1$. Then we have the lower bound [Tal21, Eq. 2.15]

$$C \left(\kappa + \sum_i \Delta_i \right) \sqrt{\log n} \leq \mathbf{E} \sup_{t \in T} X_t \leq C' \left(\kappa + \sum_i \Delta_i \right) \text{Diam}_d(T) \sqrt{\log n}$$

One simple idea for a metric upper bound is:

$$d^{ub}(\pi_1, \pi_2) = \mu \sum_{i=1}^n |\Delta_{\pi_1(i)} - \Delta_{\pi_2(i)}|$$

Under this metric, we need to show that

$$\mathbf{Prob}[|X_\pi - X_{\pi'}| \geq u] \leq 2 \exp\left(-\frac{u^2}{2d(\pi, \pi')^2}\right)$$

This is effectively direct from Azuma's inequality since Δ_i is in a bounded ball (in order for us to use curvature). Next, we need to show $\mathbf{Prob}[X_s \geq \sqrt{\log n}(\kappa + \sum_i \Delta_i)] \geq 1$. For each permutation $\pi \in S_n$, we can construct a binary tree T_π from the partial sums $S_i \sum_i \Delta_{\pi(i)}$, where $S_i < S_j$ implies S_i is in the left subtree of S_j (and vice versa). Assume, first, that each S_i is unique. Then, it can be shown that the expected height and the tail bounds for the height of this subtree satisfies [ABC20, Ree03]

$$\mathbf{Prob}[h(T_\pi) \geq \sqrt{\log n}] \geq \frac{c}{n}$$

Our conjecture is that $\kappa h(T_\pi) \leq X_s \leq \mu h(T_\pi)$ which would immediately imply $\sum_{\pi \in S_n} \mathbf{Prob}[X_\pi \geq u] \geq 1$. Unfortunately to find bounds of this form with fees, one needs to find universal bounds on $g(\Delta) - \gamma g(\gamma \Delta)$. We illustrate such bounds for Uniswap in Appendix I.

I Path Dependency in Uniswap

Getting bounds such as (15) relies on bounding how far away the path-dependent case strays from the path independent case. For a fixed Δ , p_n^{pi} only depends on $\sum_i \Delta_i$ for path-independent, whereas $p_n^{pd}(\pi)$ does depend on the path $\Delta_{\pi(1)}, \dots, \Delta_{\pi(n)}$. However, if we can uniformly bound $\max_{\pi \in S_n} |p_n^{pd}(\pi) - p_n^{pi}|$ as a function of fees and curvature.

For Uniswap, we have $g_{uni}(\Delta) = \frac{k}{(R-\Delta)^2}$. This gives a difference between the impact of a single path independent trade and a single path dependent trade as (see [AEC20] for the formulae):

$$\begin{aligned}
g(\Delta) - \gamma g(\gamma\Delta) &= k \left(\frac{1}{(R-\Delta)^2} - \frac{\gamma}{(R-\gamma\Delta)^2} \right) = \frac{k}{(R-\Delta)^2} \left(1 - \frac{\gamma(R-\Delta)^2}{(R-\gamma\Delta)^2} \right) \\
&= g(\Delta) \left(1 - \frac{\gamma(R-\Delta)^2}{R^2} \frac{1}{\left(1 - \frac{\gamma\Delta}{R}\right)^2} \right) \\
&\leq g(\Delta) \left(1 - \frac{\gamma(R-\Delta)^2}{R^2} \left(1 - \frac{c\gamma\Delta}{R} \right) \right) \\
&= g(\Delta) \left(1 - \frac{\gamma(R-\Delta)^2}{R^2} - \frac{c\gamma\Delta(R-\Delta)}{R^3} \right) \\
&= g(\Delta) \left(1 - \gamma \left(\frac{R-\Delta}{R} \right)^2 \left(R - \left(1 + \frac{c}{R} \right) \Delta \right) \right)
\end{aligned}$$

where we assume that $\frac{\gamma\Delta}{R} < 1$ and use the geometric series (so $c < 1$). When $R \gg 1$ and $R - \Delta \leq kR$ for some $k < 1$, this gives us the bound

$$\frac{g(\Delta) - \gamma g(\gamma\Delta)}{g(\Delta)} \leq 1 - \gamma \frac{(R-\Delta)^3}{R^2} \leq 1 - \gamma k^3 R$$