Poolside - A Simple AMM for Value-Accruing and Rebasing Tokens

Manny Rincon-Cruz, Socks&Flops, Fiddlekins

August 15, 2023

1 Introduction

Buttonwood Poolside is a simple constant-product automated market maker (AMM) protocol optimized for use with value-accruing or rebasing tokens—VAR tokens for short. Value-accruing tokens serve as "receipts" that are then redeemed for a different underlying asset. Rebasing tokens are tokens whose supply changes dynamically in the wallets of all token holders. The most important VAR tokens to date include liquid staking derivative tokens (LSDs), real-world assets (RWA), bonds, and synthetic commodities like Ampleforth (AMPL).

Liquidity providers (LPs) of all VAR tokens, whether rebasing or non-rebasing, incur avoidable divergence losses.[1] For example, for native network assets and their corresponding LSDs, LPs will lose the yield from their LSDs. As of writing, LSD protocols have provided between $50M to $150M in yearly incentives to LPs in order to maintain liquidity for their tokens, which is crucial for broader integrations in decentralized finance (DeFi).[2] The two largest Ethereum LSD protocols are Lido and Rocket Pool. Both protocols used their governance tokens as rewards. For the year ending on May 2023, these two protocols had distributed 32,753,741 LDO tokens and 49,656 RPL tokens, which at today's prices are worth $60.1M and $1.4M.

Poolside mitigates these losses through the use of "reservoirs." These are pockets of inactive liquidity that modulate changes in the value of VAR tokens. Additionally, as reservoirs grow in size, LPs can opt to deposit more of the matching, scarcer token so that both tokens can flow back into the active liquidity pool. To prevent reservoir manipulation attempts, the reservoir-matching function has three guardrails: a flow limit, a volatility circuit-breaker, and an approximated time-weighted-average-price (TWAP) requirement.

The ideal VAR token is defined by two properties: 1) its change in value happens through on-chain processes and 2) its rate of change is sufficiently slow relative to a chain's execution speed. If

[1] Divergence loss refers to opportunity cost incurred by LPs when token prices diverge. For any one pool, they will own less of the appreciating token and more of the depreciating token. Divergence loss is also popularly known as "impermanent loss," however this older terminology is misleading as the loss is only "impermanent" and unrealized if token prices re-converge.

[2] https://research.lido.fi/t/an-update-on-rewards-v2/2797; https://xangle.io/en/research/detail/1129; https://thedefiant.io/rocketpool-reth-incentives; As of April 2023, Rocket Pool had allocated a budget of 67,500 RPL tokens, which at today's prices of $28 per token are $1.9M. Through 2022 Lido was spending over $10M per month, and in June 2023 phased out LDO rewards for stETH rewards of about $15M.

the change in a token's value happens through observable off-chain processes, then arbitrageurs will be able to front-run any oracle or transaction that could update a VAR token's contracts. Similarly, if a token's rate of change in value is too fast relative to a blockchain's execution speed, arbitrageurs can use memory-pool reordering to front-run transactions, a process otherwise known as "miner extractable value" (MEV). For that reason, Poolside is being launched with a fixed number of LSD-based pairs and a lock on additional pair creation. Protocol governance can, however, vote to remove the lock and allow permissionless pair creation.

The on-chain nature of LSD tokens and the relatively slow rate of staking rewards on Ethereum make Poolside an optimal solution for LSD liquidity, which is why we use LSDs as the canonical example throughout the white paper. However, we do expect similarly positive outcomes with collateral lending tokens and RWAs.

2 The problem of divergence loss and VAR tokens

Divergence loss is a well-known issue for LPs on almost all AMMs and is unavoidable with most token pairs. But this need not be the case with VAR tokens.

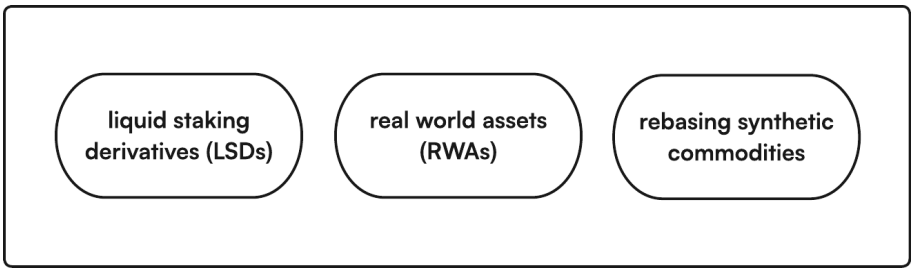

The three main varieties of VAR token are LSDs, RWAs, and synthetic commodities. LSD tokens such as those from Lido, Rocket Pool, or BENQI, represent network native tokens—ETH in the case of Lido and Rocketpool and AVAX in the case of Benqi—that have been deposited with the protocol. The protocols allow an underlying set of validators to stake ETH or AVAX to earn more ETH or AVAX. Some of this yield is given to the protocol, some of it to validators, and the rest flows back to users, who can then redeem their LSD tokens for a greater amount of ETH or AVAX.

RWA and coupon bonds have a similar mechanic. Instead of depositing a native network token for the RWA bond, users deposit or exchange on-chain fiat or stablecoins for bond tokens. These bond tokens are "receipts" for the original purchase amount plus interest. Both LSD tokens and RWA tokens can be rebasing or non-rebasing.[3]

The last category of VAR tokens are synthetic commodities, of which AMPL is the canonical example. AMPL is rebasing but cannot be redeemed for anything else. Instead of tracking predetermined staking rewards rate or bond interest rate, AMPL rebases according to the average AMPL

[3]For example, Lido and Aave tokens are rebasing, while RocketPool and Compound v2 tokens are non-rebasing.

price over the previous 24 hours. Thus, these changes in supply and value can be both negative or positive.[4]

ETH LSDs best satisfy both properties of well-behaved VAR tokens: 1) the minting and redemption rates between ETH and LSDs is entirely on-chain, which makes arbitrage unattractive, and 2) staked ETH accumulates rewards at about 0.016% per day, which makes MEV front-running difficult to execute in a profitable manner. For these reasons we think an AMM designed for LSDs is desirable and feasible.

Below we illustrate divergence loss for two imaginary LSD tokens, one being the rebasing rbETH and the other being the non-rebasing nrETH.

In the case of rebasing tokens like rbETH, divergence loss is caused by the depreciation of rbETH's marginal price against ETH. As the rbETH protocol accumulates rewards, it increases the balance of rbETH tokens owned by LPs. An AMM, however, will interpret the increase in rbETH tokens as a decrease in the value of rbETH relative to ETH, because on an AMM the marginal price of any token is equal to the ratio of tokens within the active liquidity pool. In this case, the pool has 12 rbETH and 10 ETH, so the marginal price of 1 rbETH is 10/12.5 ETH, or 0.8 ETH. Any trader can then swap 0.8 ETH for 1 rbETH, even though the underlying redemption value of 1 rbETH is 1 ETH.

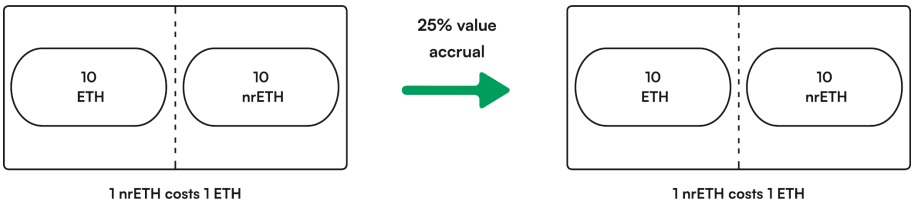

For non-rebasing tokens like nrETH, divergence loss is caused by a failure of the AMM to appreciate the marginal price of nrETH against ETH. As the nrETH protocol accumulates rewards, the value of each nrETH increases. Assuming an accumulation of rewards equal to 25% of their original ETH, 1 nrETH can be redeemed for 1.25 ETH. However, on an AMM the marginal price of nrETH and ETH remains unchanged at 1 nrETH for 1 ETH. An enterprising trader could then exchange 1 ETH for 1 nrETH, which he can then redeem for 1.25 worth of ETH from the LSD protocol.

[4]Evan Kuo, Brandon Iles, and Manny Rincon Cruz. Ampleforth: A new synthetic commodity. Ampleforth White Paper, 2019.

3 The solution

The core of Poolside is an $xy = k$ pricing curve. Each pair on Poolside consists of a pool of active liquidity and two reservoirs of inactive liquidity. When a VAR token grows in value, that extra value accumulates in the reservoir, not the main pool. If a VAR token loses value, as is possible with rebasing synthetic commodities, then it is the other token in the pair which flows into the reservoir—contracting the active liquidity whilst retaining the marginal price.

We illustrate these mechanics for LSDs by using rbETH and nrETH, as in the examples above.

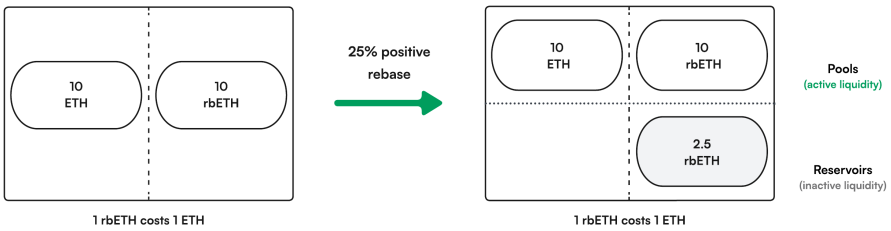

In the case of a rebasing token like rbETH, a 25% increase in the token's underlying value would produce an increase of 2.5 rbETH tokens, which instead of staying in the main pool would flow into one of the reservoirs. The reservoir thus preserves the marginal price of 1 ETH for 1 rbETH.

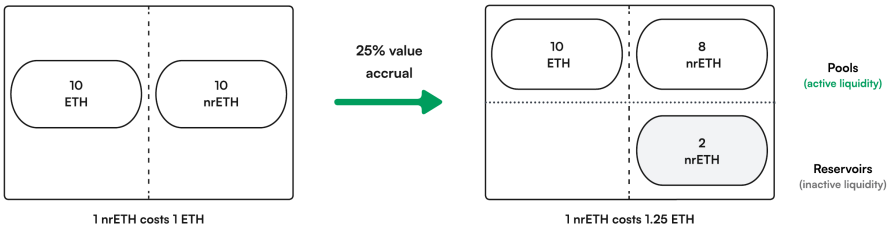

In the case of non-rebasing tokens like nrETH, a 25% increase in the token's underlying value leaves the token supply unchanged, but as we noted above, the marginal price of nrETH needs to increase to 1.25 ETH in order to stave off divergence losses for LPs. For this reason, 2 nrETH tokens flow into the reservoir, which leads to an adjustment of nrETH's price up from 1 ETH to 1.25 ETH.

Users and developers should note that non-rebasing tokens on Poolside require the use of a special token wrapper before they are deposited into a pair, much like ETH is often wrapped into wETH on most DeFi protocols. This special wrapper allows non-rebasing tokens to behave as if they were rebasing tokens, although much like wETH, this process is handled silently by a router and is completely abstracted away from the eyes of final users. Nonetheless, taking advantage of Reservoir's unique features will require the instantiation of a wrapper for any new, non-rebasing VAR tokens.

We chose this rebasing-first design framework because it appears likely that rebasing tokens will continue to grow in their total value locked (TVL) and to gain traction in the broader DeFi

ecosystem. Already, the largest LSD and lending protocols use rebasing tokens. Lido's rebasing stETH represents $14.7B in ETH, while RocketPool's non-rebasing rETH represents only $1.8B in ETH. Similarly, Aave's rebasing aTokens hold $5.8B in TVL, while Compound's non-rebasing cTokens hold $1.3B in TVL. In addition, Compound v3 has recently launched its own growing suite of rebasing tokens totaling $932M in TVL.[5] It seems likely that RWA protocols will end up adopting similar designs.

Consider an RWA asset, say a tokenized Treasury bill, which represents its interest accrual through a rebasing token. Very much like an LSD, the accrued interest would flow into the Treasury-USDT pair's reservoir.

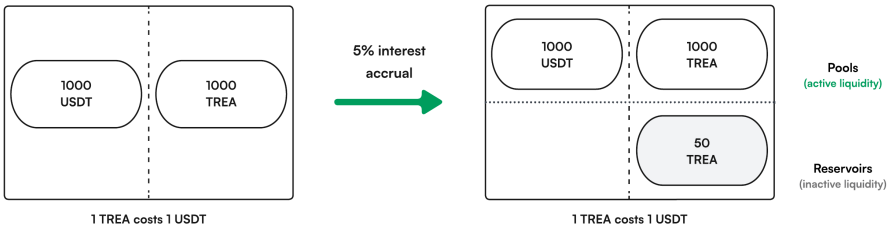

Lastly consider the pair of AMPL with the SPOT stablecoin[6], since both of them track the value of 2019 US dollar. A positive rebase in the AMPL supply would cause excess AMPL to flow into a reservoir.

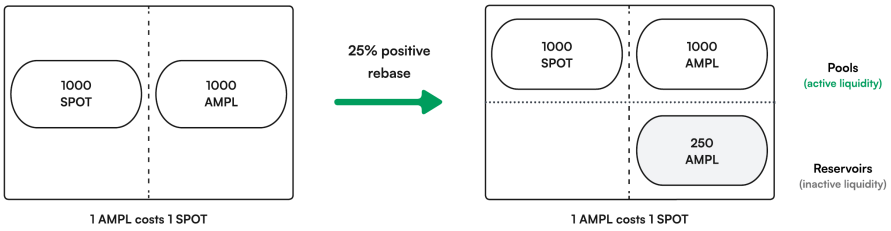

Conversely, a negative rebase in the AMPL supply would cause SPOT to flow into the reservoir.

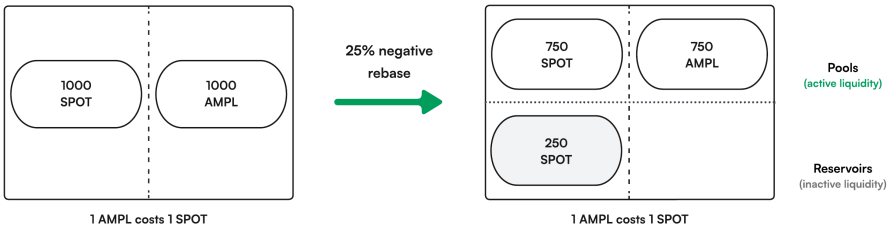

[5] https://defillama.com/protocol/compound-v3

[6] Evan Kuo, Brandon Iles, Nithin Krishna, and Manny Rincon Cruz. SPOT—An Inflation Resistant Store of Value. SPOT White Paper, 2022.

4 Single-sided liquidity and reservoir guardrails

Although Poolside mitigates LPs' divergence losses, as more tokens accumulate in a reservoir a pair's capital efficiency decreases. For LSD and RWA tokens, this is a problem that worsens over a span of years, while for synthetic commodities it might occur in months. At current Ethereum staking reward rates, it would take over 8 years for an LSD-pair's reservoir assets to equal 25% of the ETH value of the main pool, while it would take 10 years for a tokenized ten-year Treasury bond and USDT pair.

For this reason, Poolside includes a set of methods for rebalancing a pair's pool and reservoir through single-sided liquidity provisioning. Below we revisit our fictional example tokens rbETH and nrETH, but the same dynamics apply to RWAs and synthetic commodities.

Take the following case for a rebasing asset like rbETH where the pair has a marginal price of 1 ETH to 1 rbETH, and 2.5 rbETH has accumulated in the reservoir. A user can deposit 2.5 ETH, whereupon the corresponding 2.5 rbETH tokens flow into the active liquidity pool. Overall this increases the liquidity depth from 10 ETH/rbETH to 12.5 ETH/rbETH.

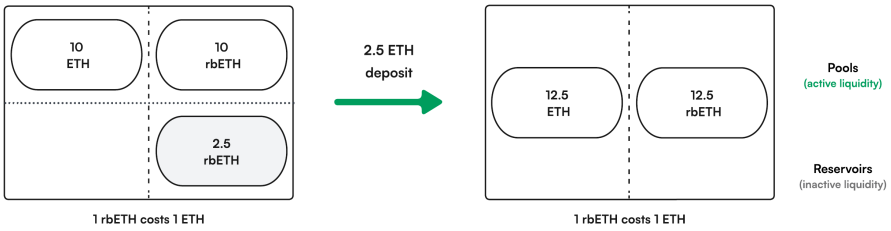

In the case of a non-rebasing asset like nrETH:

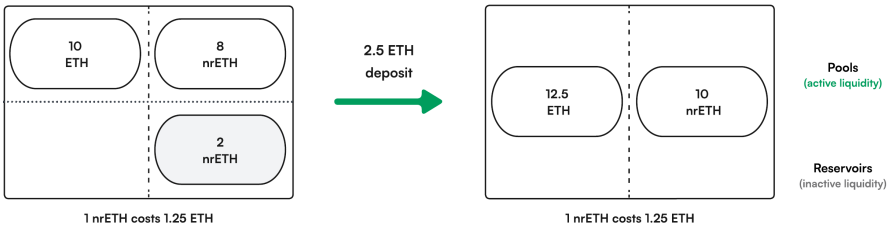

4.1 Reservoir guardrails

The amount of tokens needed for single-sided liquidity cannot be determined solely by the current ratio of tokens in the pair's liquidity pool, since this price could be subject to manipulation.

Instead, Poolside uses a time-weighted average price (TWAP) alongside a volatility circuit breaker and reservoir throttle. This serves to limit single-sided liquidity provisioning to times of relatively stable prices and low volatility. All three of these mechanisms have default parameter values. These are 24 hours for the TWAP window, 7% as the threshold for the volatility circuit

breaker, and 5% for the maximum reservoir value that can flow back into the liquidity pool. All of these can be calibrated by governance for specific pairs.

While guardrails make single-sided liquidity intermittently available, ETH-based LSDs will not need that many transactions of this type per year to maintain high capital efficiency. In return, the guardrails can guarantee that LPs preserve as much value as possible from the liquidity they provide.

4.1.1 TWAP implementation

In order to calculate a "stable" price, our approach is to compute a modified TWAP. TWAPs typically use a weighted moving average approach where each observation is weighted by the amount of time since the last observation. But this approach requires maintaining a sliding window of previous trades, which is gas-intensive.

Our approach is to approximate a time weighted moving average, not by using a list of the previous trades, but by iteratively updating the moving average at every swap. The default window is 24 hours long, which we use in the examples below, but can be configured to be longer or shorter.

We weight each price by scalars W_j, the amount of time since the last observation (truncated to fit inside the sliding window). If the observation happens outside of the moving average window, then $W_j = 0$. In essence, we average over the entire time window, whilst expiring observations that are outside of it. Therefore for any number of W_j, their total sum will always equal the size of the window, 24 hours. This allows us to write our modified TWAP in an approximate recursive fashion:

$$S_j = \frac{\sum_{i \leq j} P_i \cdot W_i}{\sum_{i \leq j} W_i}$$

$$S_j \approx \frac{P_j \cdot W_j + S_{j-1} \cdot \left((\sum_{i \leq j} W_i) - W_j \right)}{\sum_{i \leq j} W_i}$$

$$S_j \approx \frac{P_j \cdot W_j + S_{j-1} \cdot ((24hrs) - W_j)}{24hrs}$$

$$S_j \approx \alpha_j \cdot P_j + (1 - \alpha) S_{j-1}$$

where:

- S_j is the moving average at time of measurement j
- $\alpha_j = \frac{W_j}{24hrs}$ is the relative weight of the current observation over the last 24 hours at time of measurement j

As a result, our contracts need only store three values, with the rest of the necessary information accessible from the block.

4.1.2 Volatility circuit breaker

The circuit breaker disables single-sided liquidity operations in response to price volatility. This safeguard has three configurable parameters: minimum and maximum length of time during which operations are disabled, which we call a timelock, and the volatility threshold that triggers the maximum time lock. The default minimum time lock is 24 seconds, the default maximum time lock is 24 hours, and the default volatility threshold is 7%.

If volatility is defined as $|\texttt{TWAP} - \texttt{lastPrice}| \,/\, \texttt{TWAP}$, and the `calculatedTimelock` is

$$\frac{\texttt{volatility}}{\texttt{volatilityThreshold}} * (\texttt{maximumTimelock} - \texttt{minimumTimelock}) + \texttt{minimumTimelock}$$

then each trade triggers a timelock of the following length:

$$\texttt{effectiveTimelock} = \min(\texttt{calculatedTimelock}, \texttt{maximumTimelock})$$

In other words, when a trade with zero volatility happens, then $\frac{\texttt{volatility}}{\texttt{volatilityThreshold}}$ equals zero. Thus the `effectiveTimelock` is merely the `minimumTimelock`.

If `volatility` is equal to or greater than the `volatilityThreshold`, the `calculatedTimelock` will be greater than or equal to the `maximumTimelock`, so the `effectiveTimelock` will be the `maximumTimelock` of 24 hours.

In all other situations, price `volatility` will be somewhere between 0 and 100% of the `volatilityThreshold`, and thus the `effectiveTimelock` will be between the `minimumTimelock` and the `maximumTimelock`. For example, volatility of 3.5% is 50% of the 7% `volatilityThreshold`, thus the `effectiveTimelock` is halfway between 24 seconds and 24 hours.

The circuit breaker ensures that single sided liquidity operations are only live during times of relative price stability.

4.1.3 Reservoir throttling

Reservoir throttling limits how much of the reservoir can be used in a single-sided operation over a window of time. This safeguard has two configurable parameters: the fraction of the corresponding pool balance which can be used in a given timeframe, and the length of that timeframe. By default these values are set to 5% of the value of the main liquidity pool over 24 hours.

In an extreme scenario we can imagine the reservoir being larger than the active liquidity. In this case it doesn't make sense to allow the entire reservoir to be valued using a TWAP from a smaller liquidity pool. For reservoirs smaller than 5% of the value of their liquidity pool this has no impact on the user experience.

4.2 Liquidity tokens

4.2.1 Minting

As with other protocols, liquidity tokens represent a proportional share of the liquidity pool and the reservoirs. Liquidity providers can mint liquidity tokens in two ways:

1. Double-sided deposit
2. Single-sided deposit

For double-sided deposits, liquidity providers provide both assets in the same proportion as the current total balances. Depositors receive liquidity tokens relative to the proportion of tokens they have deposited into the pair.

$$\frac{L_{user}}{L_{total}} = \frac{A_{user}}{A_{total}} = \frac{B_{user}}{B_{total}}$$

On the other hand, single-sided deposits allow users to mint LP shares by only depositing the asset opposite of the reservoir.[7] This benefits existing liquidity providers by converting the reservoir back into active liquidity and improving their capital efficiency.

4.2.2 Burning

Liquidity tokens can be redeemed for their proportional share of token balances contained in the pools and reservoir. They can be burned in two ways:

1. Double-sided withdrawal
2. Single-sided withdrawal

As with deposits, for double-sided withdrawals, the liquidity provider receives a proportional share of token balances from the pools and the reservoir.

$$\frac{L_{user}}{L_{total}} = \frac{A_{user}}{A_{pool} + A_{reservoir}} = \frac{B_{user}}{B_{pool} + B_{reservoir}}$$

[7] Single-sided deposits must be converted into double-sided deposits in order to fairly distribute the correct amount of liquidity tokens. This is done by using the TWAP to trade some of the depositing asset for a portion of the reservoir until the depositor has the same ratio of asset balances as the pair. Both assets are then deposited into the pools in exchange for a proportional amount of liquidity tokens. This is executed in a single atomic transaction.

Similarly, single-sided withdrawals enable redeeming exclusively from the reservoir, within the protections of the reservoir guardrails.[8]

5 Conclusion

We believe Poolside provides a simple implementation of an idea that limits divergence losses for liquidity providers of liquid staking derivatives, real world assets, and synthetic commodities, tokens which we collectively term value-accruing and rebasing (VAR) assets. Today, Poolside's biggest potential use case is liquid staking derivative tokens, given that their value accrual is driven by an on-chain process and their rate of value accrual is slow enough to discourage maximal extractable value (MEV) arbitrage.

The benefits of using Poolside diminish when pairs include tokens whose value accrual is driven by an off-chain process, which allows front-running, or when the rate of value accrual is so fast that MEV attacks are profitable. In those circumstances, LP performance converges with that of a constant-product AMM like Uniswap v2. As a precautionary measure, Poolside will be deployed as a "guarded" launch. This means that new pair creation is disabled until governance re-enables permissionless pair creation.

Lastly, we believe Poolside has room for expansion. Future features that can be funded by the Poolside DAO include concentrated liquidity curves, cross-chain deployments, and a universal wrapper for non-rebasing, value-accruing tokens.

[8]Single-sided withdrawals must also be converted into double-sided withdrawals in order to calculate the correct amount of liquidity tokens to burn. First, the liquidity tokens are burned for a proportional amount of both assets in the pair. The TWAP is then used to trade the asset opposite the reservoir, until the withdrawer only has the reservoir asset. This is executed in a single atomic transaction.