dYdX Liquidity Staking Module Review

Justin Jiang
jjiang990@gmail.com

Thomas Cintra
tncintra@gmail.com

Max Holloway *
max@xenophonlabs.com

June 2022

Abstract

Since its inception, the Liquidity Staking Module (LSM) has suffered from a lack of utilization on behalf of community approved market makers (borrowers), despite providing them with interest-rate-free capital. In this paper, we analyze historical staking and borrowing through a series of Dune Analytics queries. We discuss factors that might be driving under-utilization, as well as the volatility in staked USDC supply over the past several months. We conclude that the Liquidity Staking Module is not adequately achieving its stated objectives and should be discontinued. Finally, we survey alternative uses for the remaining DYDX allocated for liquidity staking, and narrow them down to a set of recommendations.

*Disclosure: Reference to the DYDX price is necessary for this research. The authors do not own DYDX token, nor are they affiliated with dYdX Trading Inc. or any of its affiliates. This research was funded by the dYdX Grants Trust. Any opinions and results stated here are those of the authors, not of dYdX, its affiliates, nor the dYdX Grants Trust. This is not financial advice.

Contents

1 Introduction

1.1 Research Motivation

Cryptocurrency exchanges, whether centralized or decentralized, suffer from the limited number of high-quality market makers. The goal of the Liquidity Staking Module is to attract community-approved market makers to the dYdX platform by providing them with interest-rate-free, non-recourse USDC loans. These uncollateralized credit lines are used by market makers to pump up liquidity across all markets on the dYdX layer 2 perpetuals protocol. The foundation provides a more in-depth introduction of the module here.

The liquidity pool seems attractive to market makers and traders alike. Liquidity providers benefit from interest-rate-free capital to make markets on the exchange, while traders enjoy tighter spreads and more depth across all markets. While there are many different ways to measure liquidity, improvements in the most common metrics – increased volume, smaller spreads, deeper order levels – benefit *all* market participants by decreasing the friction to trade. Thus, it is imperative for an exchange to encourage market makers to provide liquidity, which dYdX has tried to do via the Liquidity Staking Module. As of May 9th, 2022 over 250M USDC is staked in the Liquidity Pool. However, over 67% of the available liquidity was not being utilized (borrowed) by any of the approved market makers, roughly 165M USDC.

The dYdX foundation has set aside 25M DYDX tokens to be disbursed as rewards for liquidity staking. The foundation is then paying rewards on a lot of USDC that is not being used to make markets on the exchange. This is bad for tokenholders as these tokens could be put to better use as rewards for other modules, or invested in other aspects through the community treasury. Furthermore, this is bad for traders as it detracts from the liquidity readily available on the exchange or available for personal use. Currently, over 3M DYDX tokens have been issued as liquidity staking rewards, a fraction of which was paid on liquidity that was actually borrowed.

Our research is organized into the following three sections:

1. Analyze the utilization of the Liquidity Staking Module and understand why the pool is currently under-utilized by borrowers;
2. Argue that tokens allocated for Liquidity Staking are better used elsewhere and propose alternative usages;
3. Consolidate our proposals into a set of recommendations.

1.1.1 An Aside of Dune Analytics

For this project, we relied heavily on data extracted from this Dune Analytics dashboard. Dune is essentially a PostgreSQL database, where transaction data from a variety of blockchains (including Ethereum) are indexed for quick and

easy access. The data and plots we refer to in this paper are all publicly available on the linked dashboard, along with their corresponding queries.

2 An Overview of the Liquidity Staking Module

In this section, we will investigate certain aspects of the Liquidity Staking Module; specifically, we will see how the price of DYDX token correlates with the amount of USDC staked in the Liquidity Pool, the historical utilization rates of Liquidity Pool by each market maker, and an investigation into the amount of liquidity rewards paid out on unused staked USDC.

As stated on the dYdX docs here, there are two goals of the Liquidity Module:

Goal 1: Incentivize the allocation of USDC for market-making purposes on the dYdX Layer 2 protocol.

Goal 2: Allocate capital to top-performing liquidity providers to increase spread, depth, and uptime on dYdX.

2.1 DYDX Token Price and Staked USDC

We claim that amount of staked USDC is significantly moderated by DYDX token price. Note that when we refer to USDC staked throughout this paper, we will not make a distinction between active and inactive staked USDC. We expect the amount of inactive staked USDC to increase throughout an epoch, but not so much to discourage borrowers from borrowing from the pool.

It is intuitive that the total amount of USDC staked in the Liquidity Pool has a positive correlation with the price of DYDX token, as an increase in price of DYDX token makes the interest paid out in DYDX more valuable and thus more profitable for a staker. Conversely, a drop in token price should decrease the size of the pool, since stakers can more likely find other, more profitable staking pools. In either case, volatility in token price can lead stakers and (as we will see later) borrowers to no longer be comfortable with the pool's risk profile.

In Figure 1, we visualize DYDX's dropping token price; it is easy to visually identify a correlation with the amount of USDC staked in the Liquidity Pool.

In section 6.1 in the appendix, we explore the relationship between DYDX token price and supply of staked USDC. We perform a Granger causality test using a VAR model with a maximum lag of 29 days. Unsurprisingly, **we find that changes in DYDX token price do Granger-cause the supply of staked USDC**. In other words, we expect that a change in DYDX token price provides statistically significant information about the future size of the Liquidity pool. This corroborates our intuition that DYDX token price is a meaningful predictor of total staked USDC supply, and we expect that drops in DYDX token price will drive staked USDC supply down.

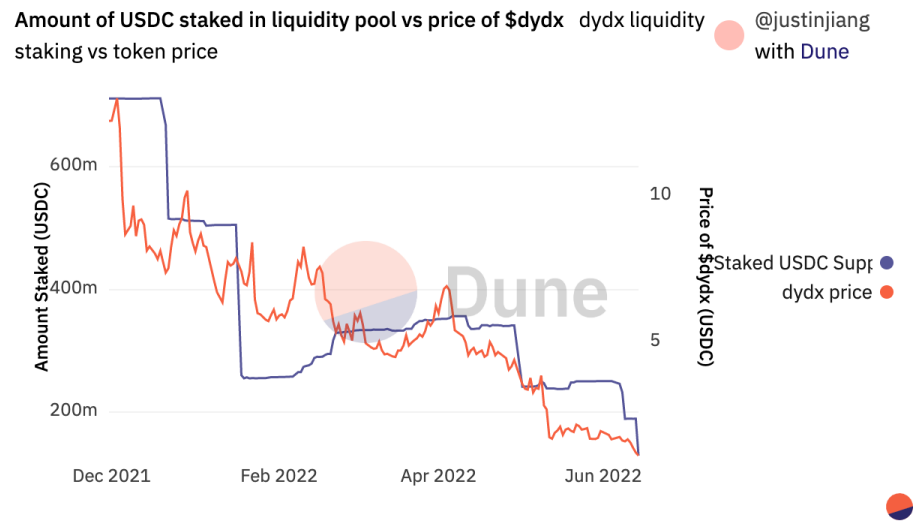

Figure 1: Graph of the size of the Liquidity Pool vs the price of DYDX token. Taken on June 15th 2022 at 5pm EST.

This variability in the size of the Liquidity Pool leads to variability in the amount of capital available to each market maker. The volatility of the size of the pool was the primary driver of changes in the percentage utilization of some of the market makers, as might be clear when we examine Figure 2. That is, even if market makers don't borrow or repay borrowings, their utilization drastically fluctuates due to changes in total supply. Let's take Wintermute as an example.

On October 27th Wintermute borrowed (notably for the last time) from the Liquidity pool. They borrowed 50M USDC which put their balance at 76M USDC. At the time this was only 37% of the total USDC they were allowed to borrow. In the next 3 months, DYDX price dropped considerably, from around $15 to $5. At the end of the January epoch, staked USDC supply had similarly dropped from around $800M to around $250M (note the seemingly linear relationship). This put Wintermute's utilization at 116% of their allotted share, and they consequently had to repay a significant portion of their borrowings. We suspect that this volatility in staked USDC supply has turned at least some market makers away from utilizing more of their allotted share of the pool. One would hesitate from borrowing even 50% of their allotted share, if in a matter of weeks this might represent over 100% of their share leading to their becoming indebted to the pool.

Notice that the primary incentive for staking USDC is the DYDX rewards awarded to stakers at the end of the epoch. We have shown through our Granger causality test that DYDX price is a powerful predictor of the amount of staked USDC. Hence, we argue that the sensitivity of staked USDC supply to DYDX price undermines Goal 1 of the liquidity module. The module's ability to provide liquidity in dYdX markets is inextricably tied to the price of DYDX token. **A downward-trending token with high volatility** suggests that the module cannot adequately accomplish Goal 1.

2.2 Utilization rate of Liquidity Pool

As of June 2022 there are 5 community-approved borrowers, each with a maximum percentage of the total pool that can be borrowed. For example, Amber Group can borrow up to 25% of the USDC in the Liquidity Pool, and as of writing of this paper, they are only borrowing roughly 6%. As you can see in figure 2, most market makers don't borrow anything close to their allotted share, although as of June 15th 2022, Wintermute has *exceeded* their maximum allotted share due to a sudden drop in staked USDC supply (observed in figure 3).

Historically, market makers rarely borrow all that is allocated to them. To measure this, we define *utilization* as

$$utilization := \frac{amountBorrowed}{amountAvailable} \tag{1}$$

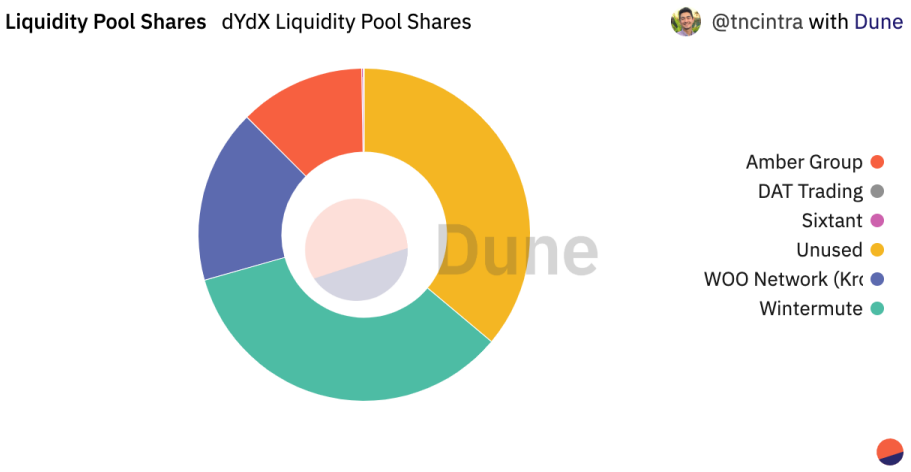

Figure 2: Pie chart of pool utilization by borrower, including unused USDC. Notice that Sixtant and DAT Trading don't appear in the chart since their utilization is negligible. Also note that Wintermute's utilization is currently above its allotted share at around 35%. Taken on June 15th 2022 at 5pm EST.

In figure 3, we can see the historical utilization of each market maker.

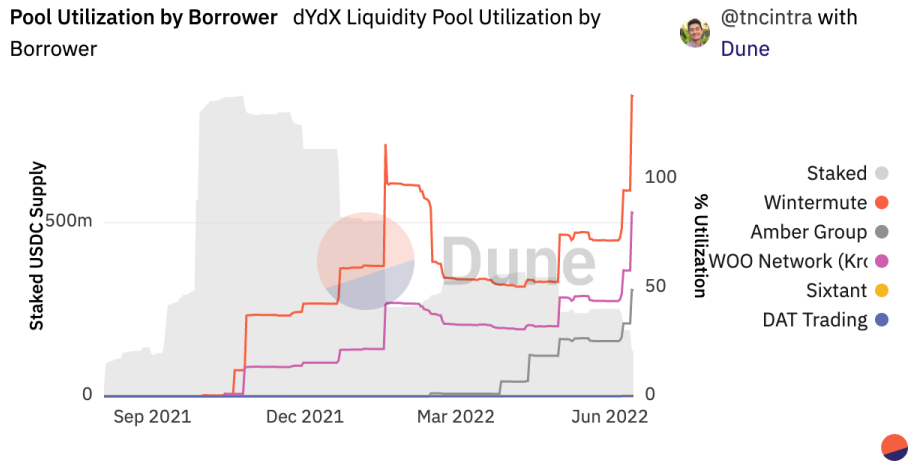

Figure 3: We can see that most market makers have not used more than 50% of the liquidity available to them to borrow. In fact, Sixtant and DAT Trading barely use any of their allotted share. Taken on June 15th 2022 at 5pm EST.

It is clear that a large portion of the Liquidity Pool is not lent out; in fact, four of the five market makers have never borrowed more than 50% of the share of the pool allocated to them; Sixtant and DAT Trading have borrowed $100K and $10K respectively, a mostly negligible amount.

Now, we investigate what the under-utilization of the Liquidity Pool by each market maker means for the total utilization of the pool. Changes in utilization are a function of both the borrower balance and the amount of staked USDC, as is clear from Equation 1. Below we visualize the utilization of the Liquidity Pool and compare it to the amount of USDC staked and borrowed over time.

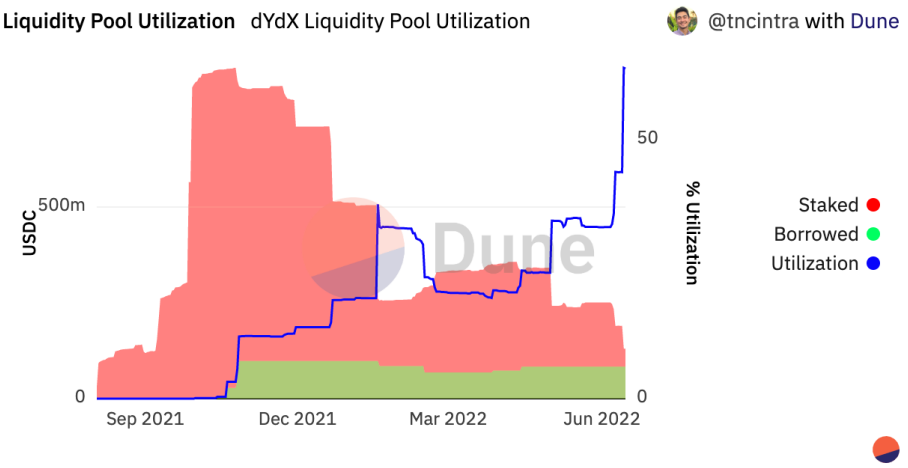

Figure 4: Graph of staked USDC vs the amount of USDC lent. Note that utilization of staked USDC has been constantly low throughout the lifespan of the module. Note: utilization spiked as we were reviewing the paper to around 60% as staked USDC supply plummeted mid-June 2022. Taken on June 15th 2022 at 5pm EST.

We can see that the total utilization of the Liquidity Pool has historically remained low; the utilization rate of the pool has never exceeded 40%. Furthermore, the increase in utilization from October 2021 to June 2022 has been largely a function of decrease in staked USDC supply, which we have argued is explained by a drop in DYDX price.

This raises the question of the efficacy of the Liquidity Staking Module: if the goal of the module is to provide accessible liquidity to market makers, how efficient has it been in accomplishing its goal?

2.3 Rewards Awarded for Unused Staked USDC

In Section 1.2, we saw how a majority of the USDC staked in the Liquidity Pool is not lent to market makers and thus does not fulfill its purpose. However, the Liquidity Pool Module still pays out DYDX token to stakers even if the USDC they staked is unused. In this section, we will find the total amount of Liquidity Pool rewards that was distributed for staked USDC that was not lent out. To begin, we define a few terms that we will be using.

Let's first define β_i as the utilization of the liquidity pool on day i where utilization is as defined on Equation 1. Next, let $\bar{\beta}_i$ be the proportion of unused staked USDC, such that $\bar{\beta}_i = 1 - \beta$, where i denotes days. Define r_i to be the total DYDX rewards distributed on day i such that the rewards awarded for unused staked USDC is $X_i = \bar{\beta}_i \times r_i$. It follows that the total amount of DYDX rewards awarded for unused DYDX over the interval $[a, b]$ is:

$$\chi = \sum_{i=a}^{b} X_i = \sum_{i=a}^{b} (\bar{\beta}_i \times r_i) \tag{2}$$

We perform a Dune query to determine χ on the interval [09.06.2021−05.29.2022]. Note that 09.06.2021 is the first day that rewards started being distributed for the Liquidity Staking Module.

For our defined interval, we found that $\chi = 3,034,602.93$ DYDX tokens. With a total of $3,748,233.17$ DYDX tokens distributed in the history of the Liquidity Staking Module, this means that **roughly** 81% **of the rewarded token has been given for USDC that was not lent to any market maker**.

Tokens Awarded	
Rewarded for un-borrowed USDC	Total Token Rewarded
3034602.93	3748233.17

For all figures and tables present herein, please refer to this Dune dashboard for the source queries and data.

2.4 Summary of Historical Analysis

In this section, we dug deep into the data behind the Liquidity Staking Module to investigate the efficacy of the module in accomplishing its goals. Below, we summarize our findings in regards to each goal.

2.4.1 Goal 1: Incentivizing USDC Allocation in the Liquidity Pool

In 2.1 we found a Granger-causal relationship between the amount of USDC staked in the Liquidity Pool and the price of DYDX token. Since the interest paid in DYDX token is the primary incentive for users to stake their USDC, the Module is missing a way to incentivize stakers when the price of DYDX token is low. To illustrate this, we note that there was roughly a 70% decrease in staked

USDC supply from it's all time high in late September to its current all time low.

Furthermore, the sensitivity of staked USDC supply to DYDX token price leads to a highly volatile liquidity pool. A volatile liquidity pool puts market makers in a difficult position: if they borrow an amount within their allocation on day i, this might exceed their allocation on day $i + n$ and put them in debt. This happened to Wintermute in January of 2022. We suspect that the volatility of the size of the pool due to its sensitivity to DYDX price leads to distrust in the module on the part of the market makers.

The supply of liquidity from the LSM seems to be largely a function of DYDX token price, an undesirable feature for a liquidity pool. We conclude that the Liquidity Staking Module does not adequately incentivize allocation of USDC for market-making purposes.

2.4.2 Goal 2: Allocating Capital to Top Market Makers To Build Better Markets on dYdX

In 2.2, we saw that market makers rarely borrow from the Liquidity Pool — four of the five market makers have never borrowed more than 50% of their share. In addition, the utilization rate of the Pool is around 34% on June of 2022, and has historically been much lower than that.

The process for defining "top market makers" and allocating them a percentage of the pool is also significantly cumbersome. **Restricting access to this capital to a small selection of market makers, a majority of which barely use it, leads to an inefficient use of stakers' capital, as well as of DYDX token rewards.** In the coming section, we will explore an alternative mechanism for providing liquidity by rewarding low-capital liquidity providers, not just "top market makers".

We then draw two conclusions:

1. The LSM does not accomplish its second goal of allocating capital to top market makers.
2. A better strategy might reward low-capital market makers in smaller-cap markets, in addition to (or instead of) rewarding top market makers. An example will be discussed in the next section.

2.4.3 Conclusion

As we saw above, the Liquidity Module does not accomplish either of its stated goals, and its lack of utilization by market makers has lead to roughly 81% of the rewarded tokens to be awarded for USDC that was not borrowed, as we saw in 1.3.

Furthermore, the sensitivity of the pool size to DYDX token price leads to an unstable source of liquidity, and burdens market makers with the possibility of

going over their allotted share, as Wintermute did in October of 2021.

By failing to accomplish either of its stated goals, and since it requires a significant allocation of DYDX token, we conclude that the Liquidity Staking Module should be discontinued. A proposal to discontinue liquidity staking would likely require governance to stop paying USDC stakers. This would lead stakers to withdraw their full balances, and market makers would then have to repay their borrowings. Stakers would then withdraw the debt repaid by LPs.

In the next section, we will propose a few alternative uses for the token share allocated to the Liquidity Staking Module that aim to accomplish the goals it defined.

3 Survey of Alternative Uses for Token Allocated to LSM

In this section, we will propose alternative uses for the remaining 20 million DYDX token originally allocated to be paid out with the Liquidity Staking Module. We will measure each proposal against the objectives set by Liquidity Staking Module or how they tackle other relevant problems within the dYdX ecosystem. Note that if liquidity staking is discontinued, remaining DYDX is transferred to the community treasury until a proposal is passed on how it should be deployed.

3.1 Other Mechanisms to Reward Market Makers

Currently, there is a mechanism in-place that rewards anyone who provides liquidity on both sides of the spread: called the Liquidity Provider rewards program, it aims to reward anyone who provides liquidity on dYdX "based on a formula rewarding a combination of uptime, two-sided depth, bid-ask spreads, and the number of markets supported". The design of dYdX's program is similar to those on other decentralized exchanges, which all follow Nasdaq's Designated Liquidity Provider Program [1]. On dYdX, to qualify for a percentage of the 1,150,685 DYDX token distributed per 28-day epoch, a trader must account for at least 0.25% of the maker volume in the previous epoch. With maker volume at $2.2 billion for epoch 10, a trader needs to provide almost $66 million in volume to qualify for rewards in epoch 11 [2]. Clearly, a user needs a large amount of capital in order to just qualify for the Liquidity Provider rewards.

3.1.1 Partnering with Hummingbot

Another route to source maker liquidity lies outside of dYdX: a partnership with the liquidity-mining platform Hummingbot. Hummingbot is an open-source

[1] Nasdaq DLP Program for ETFs
[2] Data source

market making platform: users can download Hummingbot's ready-to-use trading bot and use their own capital to provide liquidity on a number of exchanges. Certain exchanges have partnerships with dYdX, which entail reward campaigns where users can earn a portion of a reward for providing liquidity. During their partnership with Binance, Hummingbot users provided $1.7 billion in filled orders on 19 trading pairs to earn $200,000 USDC in rewards [3]

Functionally, Hummingbot's reward system is similar to dYdX's Liquidity Provider reward system, and dYdX has partnered with Hummingbot in the past: in July 2021, dYdX gave $15,000 to users who created the best strategies to run Hummingbot's trading bot [4]

Another advantage of using Hummingbot is it has no requirements to qualify to earn a share of the rewards, unlike dYdX's Liquidity Provider mechanism. Many smaller traders do not have the capital to meet the 0.25% of maker volume requirement, and so a partnership with Hummingbot would increase the number of low-liquidity market makers. Furthermore, the decentralized nature of Hummingbot's platform makes it especially attractive to a decentralized exchange like dYdX: if a large market maker were to pull its liquidity for whatever reason, traders on dYdX would immediately feel the impact. However, a single Hummingbot user does not have nearly as much effect on the market — it is the combination of the thousands of Hummingbot users that provide such a large amount of liquidity. A partnership with Hummingbot would decrease reliance on large market makers and decrease pressure on those points of failure, providing another security blanket for regular trader on dYdX.

A possible rewards partnership would resemble what other exchanges do with Hummingbot: select illiquid pairs on dYdX and pay out rewards in DYDX token. A case study published by Hummingbot investigates the partnership between them and Binance on a liquidity rewards program on the XEM/USDC pair. The report found that in a month's time, "142 distinct individuals have participated and earned rewards ... They accounted for as much as 11% of total XEM trading volume on Binance", which is about $23.5 million of volume provided by the Hummingbot users. This equates to about $4 million per week of liquidity for just 13,800 XEM (about $1,700 USDC) [5].

A glance at trading pairs on dYdX can show that there are multiple pairs that have 24H volume in the tens or hundreds of thousands of dollars and 24H trades in the low hundreds — each pair also has spreads in the tenths of a percentage point. A reasonable explanation to this low liquidity and high spread is the design of the primary mechanism on dYdX that encourages liquidity, the Liquidity Provider rewards. The design of the Liquidity Provider mechanism means that the percentage of the reward allocated to a certain pair is based on the fraction of the market cap of that pair to the rest of the pairs on dYdX. That is to say, smaller cap coins get less of the total reward allocation.

[3] Hummingbot-Binance partnership link
[4] Hummingbot-dYdX partnership Link
[5] Hummingbot NEM Case Study

Reasonably, one could see how certain pairs have too little volume for large market makers to participate. To solve this lack of liquidity, Hummingbot can bring in smaller market makers to provide liquidity to smaller market-cap pairs. We believe that funding Hummingbot reward campaigns would make dYdX a more attractive trading platform by making it easier to trade smaller-cap coins.

Now, we will recommend how a Hummingbot rewards campaign would be structured. A Hummingbot rewards campaign for a single market has three components:

1. The number of days D that this campaign will run. We recommend that $D = 28$, the length of an epoch on dYdX.

2. The amount of reward R allocated to this pair. Rewards can be given in DYDX, USDT, or the underlying token itself. Also, Hummingbot recently launched a program that rewards users in their own token HBOT. This makes for a total of four different tokens that can be allocated to a market.

   We suggest that the reward be paid in DYDX, given the logistical difficulty of selling DYDX for USDT.

3. The maximum spread S for an order to be eligible for reward. For example, if $S = 2.0\%$ and the mid-price of the coin is \$100.00, then a bid must be $\geq$ \$98.00 and an ask must be $\leq$ \$102.00.

Now, we will recommend a value for S. The Liquidity Provider mechanism sets a value of 40 basis points from the mid-price for any market that is not BTC or ETH[6]. For a coin with mid-price of \$100, this equates to an acceptable range of [\$99.60, \$100.40] or what is equivalent to $S = 0.4\%$. An examination of max spreads for the current campaigns on Hummingbot shows that the smallest max spread is 1.0% [7]. Thus, we will recommend a range of values for S such that $S \in [0.5, 1.5]$. There are currently 37 perpetual contracts traded on dYdX; we recommend that a contract with a 30-day average volume below the 33^{rd} percentile of all contracts have a $S = 1.5$. Similarly, a contract with a 30-day average volume in the top third of all contracts should have a $S = 0.5$ (Table 1). Intuitively, contracts with more volume (and therefore have tighter spreads than contracts with less volume) should have stricter spread requirements for market makers to encourage tighter spreads on their orders. Simply put, a smaller value of S will encourage more aggressive passive quoting (think limit orders closer to the best bid/ask).

Finally, we recommend a percentile-based system to decide values of R for a specific coin, similar to how we recommended values for S. For a coin with a 30-day average volume in the lower third of all coins should have highest value of R, while a coin with a 30-day average volume in the higher third of all coins should have lowest value of R. This is because that a larger R value for a coin makes it more profitable for a market maker to participate, and thus, the

[6] Liquidity Provider spreads
[7] Current Hummingbot campaigns

amount of liquidity provided will be greater for contracts with a larger R. Since our goal is to increase liquidity on contracts with less volume, then we ought to give out more rewards to those campaigns. For now, we suggest three values of R in USDT for the corresponding third-percentile splits: $R = \{500, 1,000, 1,500\}$, with further details in Table 1. Finally, we note that our proposed parameter structure follows closely to Nasdaq's Designated Liquidity Provider program, where ETFs with less volume have looser spread requirements and a larger rebate[8].

Percentile Volume	S	R
$< 33\%$	1.5	1,500
$\geq 33\%$	1.0	1,000
$\geq 66\%$	0.5	500

Table 1: Specific parameter values of S and R for a coin based on the percentile of that coin's volume

Another potential avenue that dYdX can explore is a partnership with the governance of the coin dYdX is creating a campaign on. For example, Avalanche, a smart-contract platform, partnered with Hummingbot to provide rewards on their token AVAX on multiple trading pairs on Binance[9]. We suggest that dYdX attempt to partner with each smaller-cap coin to secure funding from their governance. From the perspective of a project's governance, increased liquidity on dYdX means it is easier for traders to buy and sell their token, an obvious benefit.

We acknowledge that the values we suggested are by no means final, because market forces are dynamic. An ideal reward campaign is one that has a short time horizon to enable adaptability to market changes and also one that is informed by data. We believe that our recommendations are a reasonable starting point; however, these campaigns need to adapt to market forces with fine-tuning of these parameters. Therefore, we recommend that governance create a committee to oversee the execution and development of this partnership with Hummingbot.

3.1.2 Increasing Liquidity Provider Rewards

An obvious choice to allocate the extra DYDX token from Liquidity Staking is to increase the rewards for the Liquidity Provider program. With $2.2 billion of maker volume in the last epoch, this program accounts for a significant percentage of the total volume on dYdX. While we haven't done research into it, we suspect that, like Trader Rewards, reward given has a positive correlation to volume.

We do not recommend allocating more token to increasing Liquidity Provider

[8] DLP Factsheet
[9] AVAX Rewards Campaign

Rewards, mainly because the design of the rewards system prioritizes large-cap coins: 40% of rewards are allocated to market makers on BTC and ETH, and the rest is allocated to each market based on volume. We also suspect that with already so much of the token supply allocated to this mechanism, allocating more supply might cause diminishing returns. Finally, increasing Liquidity Provider Rewards will not incentivize more volume on the smaller coins on dYdX, and so we believe that with Hummingbot is a better solution to increasing liquidity on dYdX.

3.1.3 Properties

Benefit: The Hummingbot proposal rewards top market makers on smaller-cap coins (goal 2).

Benefit: A partnership with Hummingbot encourages more market makers to participate, which injects USDC into markets — especially if illiquid coins are chosen (goal 1). Increasing liquidity on these small-cap and illiquid coins makes dYdX a more attractive trading platform.

Drawback: A poorly designed rewards campaign could cause backlash from the dYdX or Hummingbot community.

3.2 Token Burn

In this section we explore the possibility of burning the remaining DYDX token from the Liquidity Staking Module. This would represent burning approximately 2% of total DYDX supply. We will define what token burning is, explain why many projects burn their native tokens, discuss current community support for burning DYDX, and explore how to implement a token burn through governance.

3.2.1 Overview

A token burn can be performed in a variety of ways. In its simplest terms, a token burn occurs when the team in charge of a project sends some portion of their native token to a burner address. A burner address is simply an address for which it is extremely unlikely for someone to have the private key, and can therefore never send the token anywhere else. Any token sent to a burner address is effectively permanently removed from circulation. By making the token more scarce, the team is theoretically also inflating each unit token's price (assuming the underlying token market cap stays constant). Token burns are not native to DeFi; in traditional finance, a similar effect is observed when a company issues a stock buyback or a reverse stock split.

Token burns are often hardwired into the tokenomics of a project. Exchanges such as Binance, for example, perform quarterly token burns where 50% of that quarter's revenue is used to buyback some portion of circulating BNB token and burn it. Binance's token burning has contributed significantly to the rising

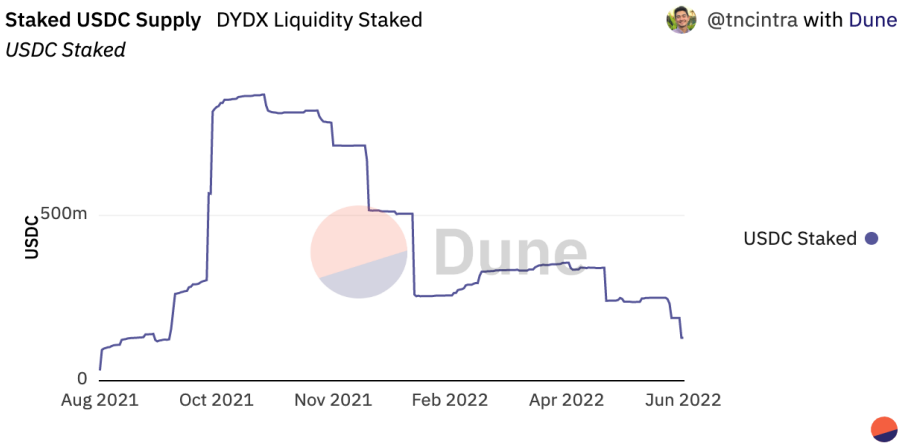

Figure 5: DYDX token Price. Taken June 15th 2022 at 6pm EST.

price of BNB token, and we observe similar token burning procedures on the protocol level in many exchanges.

Binance's strategy is better defined as "Buy and Burn" since it requires the exchange to buy BNB token from token holders before they burn it. This differs from the burn strategy we are proposing in this paper, since no DYDX token will need to be bought back. This is a key observation: as has been discussed at length in a number of community forum threads, DYDX governance does not have the ability to deploy exchange revenue to buyback any DYDX token. Since the remaining token from the LSM is not currently in circulation, this token burn does not violate the rules set forth for governance.

3.2.2 Motivation

DeFi projects burn token for a variety of reasons. We have touched on why token burns are an efficient way to increase scarcity and thereby increase the price of the token. This is arguably the primary reason many projects may create contracts to burn some amount of their native token, and is sometimes explained as a move against inflation as more and more token are minted and distributed.

As of June 2022, DYDX price remains at a low $1\$ - \$2$ per token. The significant drop in token price from its all time high in September of last year has arguably driven a lot of changes in how traders have used the platform. Xenophon Labs briefly touched on how the drop in price had affected open-interest and fees-paid in our Trading Rewards Review a few months ago. Intuitively, a less valuable token undermines the incentives in place to encourage trading and market making on dYdX, as the rewards for trading and liquidity provision

are denoted in DYDX. This is the case for any program that relies on DYDX rewards as an incentive structure.

Notably, burning DYDX has been a common and controversial topic on many community forum threads over the past several months [10][11][12]. As was repeatedly explained by members of the foundation, governance does not have permission to buyback the token, but it can burn it. Burning the remaining LSM token could illustrate the foundations commitment to the long-term value of DYDX token, as well as its commitment to engaging with ideas popular within the community.

Although DYDX is inherently a *governance token*, increasing the value of DYDX could lead to widespread improvements for the platform at large. Greater engagement with Trading Rewards and LP rewards would lead to more liquidity and therefore better markets on the exchange, as well as more revenue for dYdX Trading Inc. A pricier token would lead to greater visibility for dYdX to the broader crypto community, and more dedication to the long-term success of the project on behalf of users.

However, there are a variety of programs that rely on DYDX token for funding. Burning such a large portion of token reserves could pose an asymmetrically high opportunity cost to the community treasury. Furthermore, it is difficult to provide empirical evidence that a large token burn such as the one proposed here is guaranteed, or even likely, to raise the value of DYDX. Supply-and-demand economics leads us to posit it will, but approximating by how much is beyond the scope of this paper.

3.2.3 A Token-Burning Smart Contract

Here we propose how DYDX token should be burned. We consider two options: (1) burning all $21M$ DYDX token when the liquidity staking module is discontinued, or (2) gradually burning DYDX token at the same rate as liquidity staking module would have been disbursed. Since none of this token is currently in circulation, and the total supply of DYDX would be the same, there seems to be minimal material differences between either option.

3.2.4 Properties

Benefit: Token burning would theoretically increase the value of DYDX by enforcing greater scarcity.

Benefit: Increasing the price of DYDX would likely lead to greater engagement in Trading Rewards, LP Rewards, and any other programs that rely on DYDX rewards as an incentive structure. Furthermore, higher token price leads to

[10] Thread #1
[11] Thread #2
[12] Thread #3

greater visibility in the broader crypto community, thereby attracting more users to the exchange.

Benefit: Token burning is a popular idea amongst the dYdX community, and passing a token burning proposal through governance would illustrate the project's commitment to engaging with community members.

Drawback: There is little evidence to derive a causal relationship between token burns and token value. Burning token in hopes of driving DYDX price up could pose a large opportunity cost for the community treasury.

Drawback: There is no evidence that a token burn will necessarily drive up DYDX price, and there are a lot of more actionable ways this token can be deployed.

3.3 Increasing Trader Rewards

The current iteration of the Trader Rewards program, as reviewed by Xenophon Labs here, distributes DYDX token to individual traders based on how much they pay in trading fees and the amount of open interest on all of dYdX's trading pairs. At the end of every 28-day epoch, each trader receives r DYDX tokens, where r is given by:

$$r = R \times \frac{w}{\sum_n w_n}, n = 1, 2, \ldots, k \tag{3}$$

Note that w is the score of an individual trader and $R = 3,835,616$ is the total reward to be split between all n traders. Essentially, the term $\frac{w}{\sum_n w_n}$is the ratio of an individual trader's score to the sum of all the traders' scores, and so this term represents the fraction of R that each trader deserves. For further discussion on how w is calculated, visit this page.

3.3.1 Overview

An increase in the R distributed every epoch will create more volume, as R has a positive correlation with exchange fees and open interest. In this proposal, we analyze how increasing R would affect the total amount paid in fees at each epoch. We perform this analysis using the same repository we developed for our Trader Rewards Review (linked above). For a more detailed analysis of how modifying R might affect Trading Rewards, please refer to section 6.2.

As discussed in the appendix, there exists a unique Nash equilibrium for the amount traders will pay in fees at the end of the epoch. Using a variation of Newton's method, we can find a gradient for how the amount paid in fees (dYdX revenue) increases as R changes. By devoting all the remaining DYDX to Trading Rewards, we expect an increase of $500,000$ USDC per epoch, although this quantity is *very* sensitive to changes in DYDX price and outstanding total open interest.

While an increase in fees paid is most likely to follow an increase in R for Trading Rewards, this does not equate to improved liquidity across dYdX markets. As detailed in our review, trading for the purpose of gaming Trading Rewards is largely excluded to large, market-neutral trades to accrue open interest, as well as a sudden increase in trades placed on the last couple days of each epoch. In neither case will there be a significant gain in liquidity on dYdX by increasing R, we simply expect a marginal increase in inorganic order-flow to game the amount paid in fees at the end of every epoch. For these reasons, we find that increasing Trading Rewards would not benefit the exchange, or the dYdX community in a meaningful way.

3.3.2 Properties

Benefit: Since there is a positive correlation between R and exchange fees, then increasing R should increase total volume on the exchange and lead to more fees paid (more revenue for the foundation).

Drawback: Trading for the purpose of trading rewards leads to largely inorganic, market-neutral trades as we have explained in our previous paper. Increasing rewards is not expected to significantly increase the liquidity in any given market.

Drawback: Most additional trades resulting from this modification will be limited to the very end of each epoch.

Drawback: USDC staked in the LSM will no longer be available for top market makers, and instead will be employed to game the trader rewards mechanism. We find it hard to argue that this proposal will accomplish either of the two current goals of the LSM.

3.4 Increasing Affiliate Program Rewards

3.4.1 Overview

dYdX recently announced a new Affiliate Program, similar to those at Binance and Coinbase [13] [14].

All these programs share a similar trait: refer a new user to the exchange and get a percentage of the new user's trading fees. On dYdX, this percentage is determined by the amount of *stkDYDX* the referrer has in their wallet: a balance of 0 stkDYDX means the referrer gets 20% of the new user's trading fees, and at a maximum, a balance of 500,000 stkDYDX earns the referrer 40% of trading fees. In addition, each referred user gets 5% off their own trading fees.

Since this is a relatively new program to when this paper was written, we will not try to analyze the efficacy of the program. Instead, we will consider three

[13] Binance AP
[14] Coinbase AP

possible modifications to the program.

There is a remainder of 21M DYDX token that needs to be put to use after axing the LSM. Below, we will propose two possibilities for how dYdX can use to leftover token to bolster participation in the self-service Affiliate program and incentivize higher liquidity in dYdX's markets (Modifications 2 & 3). The first modification we propose does not require any additional token but we hope it illustrates some of the research that can be done on the Affiliate Program.

3.4.2 Modification 1 - Referrer trading fee discount

A modification that would not require allocating additional DYDX token to the Affiliate Program is having a tiered discount fee structure for referrers based on the number of referred users that meet a requirement(s). Table 2 gives an example of possible discount fee structure with tier names similar to the affiliate programs of other exchanges.

Tier	Number of Users Referred	Discount
Bronze	≤ 5	3%
Silver	≥ 5	5%
Gold	≥ 20	10%
Platinum	≥ 50	20%

Table 2: An example of a possible discount system with four tiers and referral requirement.

These trading fee discounts would stack on any existing trading fee discounts, such as the one from holding a balance of stkDYDX. Thus, for the example we gave in Table 2, the maximum trading fee discount that any user can have is 75%: 50% from holding more than 5 million stkDYDX (or more than 2.5 million stkDYDX and a Hedgies NFT), 20% from referring more than 50 qualified users, and 5% from being a referred user.

3.4.3 Modification 2 - Affiliate Rewards

Most dYdX program's emit DYDX rewards by distributing token *pro rata* amongst traders who provided some utility to the DEX. In this case, we explore the possibility of distributing the leftover DYDX from the LSM amongst traders that (1) hold some amount in stkDYDX, and (2) have *eligible* referees. Some scoring function, similar to Trading Rewards or LP Rewards, could be formulated to reward traders proportional to their contributions at the end of every epoch:

$$w_i = \frac{m^a g^b}{\sum_j^N w_j}, i = 1, ..., N \tag{4}$$

Where m is the number of eligible referees, g is the amount in staked DYDX, w is an individual trader's score, and a & b are constants. Modifying this Coubb-Douglas function and tuning constants can be done through governance. In order to bias in favor of more referrals, constants could initially be set to $a := 0.9$, and $b = 0.1$.

By *eligible* referees we are referring is a set of criteria that referred traders have to meet in order to count towards the referrer's score. These criteria are put in place to prevent the referral farming [15] and include: a minimum amount spent in fees, and a minimum average open interest held throughout the relevant epoch. Alternatively, m could be set to the amount referees have paid in fees throughout the epoch, although this might encourage end-of-epoch inorganic trading similar to that we have seen in the Trading Rewards program. By setting m to be the amount spent in fees, the cost of referral farming raises immensely, and perhaps will increase dYdX revenue.

This suggestion makes a careful use of the leftover DYDX from the LSM: it emits DYDX at the same rate as liquidity staking, except rewards are disbursed for referring more traders instead of staking USDC.

3.4.4 Modification 3 - Payout-as-you-refer

Referrers get a one-time payout for each new user they refer that meet certain requirements. For example, such a requirement could be that the new user has to trade at least $5,000 of volume in the first month of the account. If this requirement is met, then the referrer gets a one-time payout of N DYDX tokens. A higher N would mean a greater demand for referrals, but since supply is constant, there will be fewer referrals overall. The key design questions for this modification are then: (1) what criteria do referees need to meet, and (2) how much token should be awarded per referral?

However, despite a careful selection of referee criteria, a payout-as-you-refer program could still be more heavily subject to referral farming than the previous modifications. It further does not encourage referees to trade or use the platform past meeting the criteria for rewards. This program also runs out of DYDX token to provide at a non-fixed rate, unlike the suggestion in modification 2.

3.4.5 Deriving an Optimal Reward Structure for Modifications 1 & 3

In the appendix we explore a possible avenue for research on establishing parameters for modifications 1 and 3. With a presumably constant supply of DYDX to be allocated to the Affiliate Program, then there must exist an optimal way to structure the program to maximize the number of referrals.

[15] Farming means one user creates multiple addresses, meets the reward criteria on each of them, and receives the DYDX rewards.

Thus, in Section 6.3 we define a linear program that assumes only Modifications 1 and 3 have been implemented. We use the linear program as a tool to find the optimal values for the variable components of our proposed modification to maximize our objective value of referrals. While out of the immediate scope of this paper, we believe that an in-depth analysis of the Affiliate Program with different modifications and objective functions would prove valuable to dYdX. We offer the analysis in Section 6.3 as a glimpse of how we would approach this research.

3.4.6 Properties

Benefit: A more popular referral programs means more user growth. Furthermore, enforcing a volume requirement to access rewards means more liquidity on dYdX (goal 1).

Drawback: Any referral program runs the risk of being farmed. We recognize that a direct payout system like that suggested in Modification 3 has the highest risk of being abused. A further research avenue is to explore the possibility of misuse for any modification proposed.

Benefit: Both **Modifications 1 or 2** encourage staking DYDX, either for rewards or fee discounts.

Drawback: If the referral program is popular and **modification 1** is implemented, many users with trading fee discounts could decrease profitability.

Benefit: Modification 2 proposes a use for the leftover DYDX that encourages: more referrals, more fees paid by referees, and more DYDX staking. This bolsters the profitability of the Affiliate Program, which could further help bring users to dYdX and reward referrers with dYdX's governance token, not USDC.

Drawback: Modification 2 could encourage wash trading by referees to reach some amount of OI and fees paid to optimize rewards for referrers. This more advanced form of farming could exploit a vulnerability in our reward function, causing DYDX to be paid out but not value is added to the DEX.

3.5 Funding Marketing Campaigns

Onboarding new traders onto dYdX is essential to the growth of the protocol. Efforts to broaden the top of the funnel to get these new traders is through increased spending to marketing programs. Although marketing programs are beyond the expertise of the authors, there is a marketing subDAO [16] that could use funding to pursue marketing campaigns. We met with this group to discuss their ideas for new spending areas, and they found that event sponsorships, media channel (e.g. YouTube) sponsorships, and influencer sponsorships may all be fruitful avenues for increased spending. These spending areas would re-

[16] Contact: `alexios@dydxgrants.com`

quire little organizational overhead to accomplish, and thus funding a marketing campaign is a feasible area to allocate funds.

Given our limited knowledge of marketing, we leave it to the community to decide if new token emissions should be allocated to marketing campaigns through the marketing subDAO.

4 Recommendations & Future work

First we highlight that the fundamental recommendation, and the object of this paper, is that the Liquidity Staking Module should be discontinued. The remaining DYDX token allocated for liquidity staking should be vested in the community treasury until an adequate proposal is passed on how it could best be put to use.

Below, we discuss what we find to be the best alternative use for the remaining DYDX in the module.

4.1 Recommendations

In the previous section we gave an overview of a number of alternative uses for the remaining DYDX allocated for liquidity staking. Here we would like to explain why we recommend a partnership with Hummingbot as the best next step. First, we would like to briefly reiterate why other alternatives are not being recommended:

1. There is little evidence to argue that a token burn will necessarily drive up DYDX price. Since that token will be permanently removed from circulation, it poses a high opportunity cost for community treasury when there are plenty of programs within dYdX that could benefit from additional funding.
2. Increasing trading rewards will not improve liquidity across dYdX markets and will mostly serve to inflate revenue from fees via further inorganic trading at the end of every epoch.
3. Increasing Liquidity Provider Rewards does not incentivize market makers to provide liquidity on perps other than ETH and BTC. Furthermore, a relatively high percentage of token supply is already allocated to Liquidity Provider Rewards, and so we believe that any more token allocated would provide marginal benefit compared to Hummingbot.
4. Modifying the Affiliate Program is an attractive next step for the community, but further research still needs to be conducted.
5. Since marketing strategy is outside the domain of expertise of the authors, we leave it up to the community to decide if increasing funding for marketing campaigns is a fruitful use of remaining token emissions.

4.1.1 Partnering with Hummingbot

As we mention in section 3.1, partnering with Hummingbot will theoretically provide much-needed incentive for market making in smaller-cap markets on dYdX. Many exchanges have partnered with Hummingbot for precisely this reason, and we saw how Binance benefited immensely from its partnership with Hummingbot in providing liquidity for its XEM/USDC trading pair. Given the amount of low-liquidity, high-spread markets on dYdX, a partnership with Hummingbot is an appealing alternative to the Liquidity Staking Module.

This partnership would get rid of the need for USDC Staking in trying to incentivize market making on the exchange and would benefit lower-capital market makers instead of prioritizing larger, more established market makers with far larger sums of capital. This would also promote the decentralized nature of the exchange by decreasing the reliance on large market makers, who are currently rewarded by LP rewards.

Hence, a partnership with Hummingbot is most appealing if the predominant problem for dYdX is lack of liquidity on smaller markets and a lack of mechanism to reward low-capital market makers.

4.2 Future Work

4.2.1 Further work on the Affiliate Program

The Affiliate Program is a brand-new mechanism that dYdX has released. It would be interesting to examine the data of referred users to see if the Program has been accomplishing its goals. In section 3.4.1, we explore three possible modifications to the Affiliate Program: (1) a discount tier, (2) a rewards program, and (3) a payout-as-you-refer program. We believe that modification (2) presents an interesting use of the leftover DYDX from the LSM.

In Section 6.3, we use a linear program to model the modifications we proposed in Section 3.4.1. Some examples of further work that can be done on this model include analyzing the dual of the linear program and using shadow prices and complementary slackness to intuit the mathematical properties of our proposed modifications. With the goal of improving the objective of the Affiliate Program – increasing the number of users on dYdX – we can propose, model, and examine modifications to the Program. We can add secondary objectives, such as increasing new users' trading volume or decreasing the number of 'farmed' referrals, and explore how weighing each objective differently would affect the number of users.

5 Conclusion

Since its inception, the liquidity staking module (LSM) has been heavily underutilized by its approved borrowers, and has seen an incredibly volatile supply of staked USDC. This paper was designed with the following objectives: (1)

perform analytics on the LSM's historical utilization, (2) understand in what ways the LSM achieved or did not achieve its stated goals, and (3) propose alternative uses for the remaining DYDX allocated for liquidity staking.

In Section 2, we analyze the prior utilization rates of the Liquidity Staking Module. We find that a majority of the pool has remained largely unused, such that most rewards paid out for Liquidity Staking have been awarded for unborrowed USDC. We show that two market makers functionally never touched their liquidity allocation, and two others have never used more than 50%. We then explore how the downward-trending, highly volatile behavior of DYDX token price has driven changes in staked USDC supply, and why this might be of concern for stakers and market makers alike. Notably, Wintermute has gone over its allocation twice, most recently in June of 2022, due to sudden declines in staked USDC supply. *From this, we conclude that the LSM has not accomplished either of its stated objectives and should be discontinued.* The remaining DYDX allocated for liquidity staking would be better used elsewhere.

In Section 3, we propose four alternatives for how the foundation might better employ the remaining DYDX token. These mechanisms include partnering with Hummingbot, burning the remaining DYDX token, increasing trading rewards, modifying the current Affiliate program, and funding a marketing campaign. We analyze the benefits and drawbacks of each proposal.

In Section 4 we consolidate our proposals into a set of recommendations, and justify why we find partnering with Hummingbot to be the best next step for the community.

If liquidity staking is axed then determining the best alternatives for the remaining DYDX is up to governance. In this paper we outline what we find to be the best next steps, but our main goal is to spur further community discussion on what the priorities are within the dYdX ecosystem, and how the remaining DYDX can be put to better use.

6 Appendix

6.1 Granger Causality Testing

In this subsection we explore whether or not DYDX token price Granger-causes the supply of staked USDC in the Liquidity Staking Module. Granger causality doees not necessarily establish a strict causal relationship between two variables X and Y. Instead, if variable X is said to Granger-cause variable Y, then this can be better understood as "variable X is useful in forecasting variable Y", or "variables X and Y are temporally-related". In this section, we will show that changes in DYDX token price Granger-cause the supply of staked USDC in the Liquidity Staking Module. More formally:

Changes in DYDX token price provide statistically significant information about future staked USDC supply.

Our notebook and corresponding data can be found on this repository.

To test for Granger causality we must first assume that the two variables we are testing are stationary. We use the Augmented Dickey-Fuller test to test for stationary. We find that, while DYDX token price is not a stationary timeseries itself, its first derivative is. Recall that the first derivative of a timeseries is simply the changes (or delta) of that timeseries at each timestep. Using the first difference of a non-stationary timeseries is commonly known as a first differences approach. Staked USDC, on the other hand, is itself a stationary timeseries.

Now that we have a bivariate set of stationary timeseriees, we can deploy a VAR model to test for Granger causality. We use the standard Wald (Chi-squared) test to test for Granger causality, and set the maximum lag of our model to 29 days (note that a DYDX epoch is 28 days). The Wald test establishes the following null hypothesis, and attempts to reject it at a 5% significance level:

H_0: The first difference of DYDX token price does not Granger-cause the supply of staked USDC in the Liquidity Staking Module.

After running our tests, we find that we **reject the null hypothesis at a 5% significance level**. The test statistic, critical value, and corresponding p-value are displayed in the table below:

Test Statistic	Critical Value	p-Value
5.176	3.841	0.023

We conclude that changes in DYDX token price Granger-cause the supply of staked USDC. In other words, we expect that changes in DYDX token price help us forecast the supply of staked USDC in the LSM.

6.2 A Numerical Approach to Estimating Fees

For our previous grant, Xenophon Labs used numerical methods to determine a Nash equilibrium for the optimal fees each trader would pay when gaming the

Trader Rewards mechanism. We derived a pure-strategy, unique Nash equilibrium for the optimal set of fees traders would pay using a modified Newtonian Gradient Descent and compared it to historical data. In this section, we leverage the same code to observe how modifying the parameter R would affect the gross amount paid in fees, and therefore how much revenue would be generated for dYdX Trading Inc.

For more information on how we developed our Newton's method approach, how it compares to historical data on dYdX markets, and how we generate distributions of open interest and staked DYDX for our simulations, please refer to the linked paper.

Below we display the relevant parameters in our simulation. Note that all these parameters with the exception of $whale_alpha$ are determined empirically from data in the latest epochs. We welcome community members to clone our repository and explore how tuning these parameters might affect our results.

Parameter	Description	Value
D	Total average open interest throughout epoch	100M
n	Number of eligible traders	1000
R	Total DYDX to be disbursed	Varying
p	DYDX token price	2
alpha	Newton's method learning rate	0.01
G	Total average staked DYDX throughout epoch	25M
num_whales	Number of market whales (large share of OI)	10
whale_alpha	Controls whale size in Dirichlet PDF	10

Table 3: Parameters for Newton's method simulation

We vary the total DYDX to be disbursed from its original value $R = 3,835,616$. As of the writing of this paper, there are 54 remaining DYDX epochs and roughly $21M$ remaining DYDX to be disbursed from liquidity staking. Hence, if we distribute the remaining $21M$ DYDX amongst all the trading reward epochs, we would increase R to $R = 3835616 + \frac{21000000}{54} \approx 4,224,505$. This reflects increasing R by 388889.

In our review of Trading Rewards, we observed a quasi-linear relationship between R and the sum of fees paid. We noted that the linearity of this relationship is contingent on a mostly uniform distribution of open interest; when market whales are introduced (ie. traders that represent a large share of open interest) this relationship is weakened. However, we still expect that increasing R would increase the total amount paid in fees at the end of the epoch. Using the parameters defined in table 3, we run our simulation for $R \in [3,835,616, 4,224,505]$ and display the results in figure 6.

Notice that the mostly linear trend we expected between R and the sum of fees paid is observed when we run Newton's method. This linearity, as discussed in

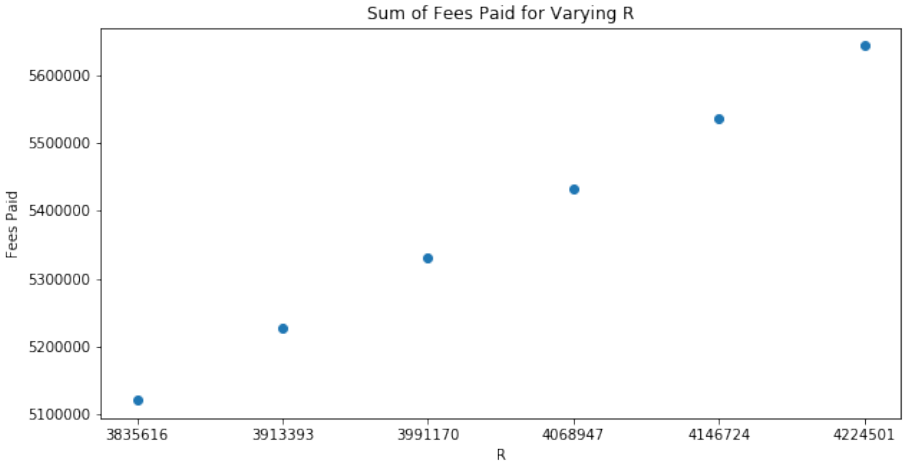

Figure 6: Sum of fees paid for varying R

our previous review, is mostly moderated by the distribution of open interest. In fact, in our review we derive a closed-form equation that approximates the relationship between the sum of fees paid, denoted as F, and R:

$$F = 0.7Rp \tag{5}$$

In our plot above, we find a gradient of roughly 1.35. Implicit in this approximation would be all of the parameters defined in table 3, which can vary greatly from epoch to epoch and are hard to accurately estimate. An important takeaway is that in a given epoch, the increase in fees resulting from diverting more DYDX to Trading Rewards will *not exceed* the current value of the DYDX itself (note that p is the price of DYDX at the end of the epoch).

Hence, we expect that diverting all the remaining DYDX to Trading Rewards will significantly increase the amount traders will pay in trading fees, and therefore will result in an increase in revenue for dYdX Trading Inc. The dollar amount by which F will increase is contingent on all of the parameters in table 3, however it is predominantly a function of the price, p. We expect the increase to be roughly approximated by $\Delta F \approx 0.7p \times \Delta R$. With $p = 2$, this equates to an additional $\$500,000$ in fees every epoch.

6.3 Optimal Affiliate Program Structure

First, we make a strong assumption: for a rational referrer, she would choose to stake her DYDX reward if and only if the benefits gained from staking is greater than what she would get from liquidating her tokens. From the exchange's point

of view, it wants to incentivize referrals while minimizing the amount of reward it has to pay out per referral. We need to set values for the required volume to qualify a referral, the amount of reward given out per referral, and the tiered discount structure for a user having multiple referrals. To accomplish this, we will define a linear program.

Assume that the referred user has to complete their volume requirements in 30 days, and so the rational referrer has a decision-time-horizon of 30 days as well. This implies that whenever the rational referrer needs a price for DYDX token, a logical value would be the 30-day average price. Thus, we will use the 30-day rolling average ($\mu_{[05.03.2022,06.02.2022]} = \$2.31$) of dYdX in our model.

We will also set the monthly unique new referrals to be 1,000 and the average trading volume of a referrer to be \$5,000. Both of these values are admittedly rather arbitrary; further work on the Affiliate Program would involve finding better estimates for both of these values. However, for now, these placeholder values will suffice. Now, we will define the linear program in algebraic notation:

First, consider the following sets:

1. $A \triangleq$ the tiers of the percentage of the referred user's trading fees earned by the referrer, defined by the current Affiliate Referral Bonus program.
2. $B \triangleq$ the tiers of the percentage discount that the referrer gets on their own trading fees.

Consider the following parameters:

1. $f_a \triangleq$ For affiliate tier $a \in A$, the percentage of the referred user's trading fees given to the referrer.
2. $s_{30} \triangleq$ the referrer's 30-day trading volume
3. $\mu_{30} \triangleq$ the 30-day average of the price of DYDX token
4. $u \triangleq$ the estimated number of unique referrals by unique referrers in a 30 day span.
5. $\xi \triangleq$ the trading fee, prior to any discounts, paid by the referrer

The decision variables are:

1. $R \triangleq$ the amount of DYDX token given out as a one-time payment for referring a user who qualifies
2. $v \triangleq$ the amount of volume that the referred user needs to trade in 30 days to earn the referrer R tokens
3. $d_b \triangleq$ For tier $b \in B$, the percentage discount that the referrer gets on their own trading fees

Consider the objective function:

$$\min \quad u[(R \cdot \mu) + (s_{30} \cdot \xi \cdot d_2)] \tag{6}$$

Note that $(R \cdot \mu)$ is the amount of USDC value that the dydx exchange loses per qualified referral and $(s_{30} \cdot \xi \cdot d_2)$ is the amount of USDC value that the exchange loses in discounted trading fees for the referrer in a 30-day span. We assume that the average referrer is in the second tier of the trading fee discount system. Thus, multiplying the sum of these two terms gives us our **objective function: minimizing the total USDC that the dYdX exchange loses in a 30-day span**. Conversely, $R \cdot u$ gives us the amount of dydx staked back into dydx's protocol, if each of the u referrers acts rationally. Finally, $v \cdot u$ can give us an estimate of the new liquidity provided by the new referred users and can also be used to estimate the extra fees paid to the dydx exchange.
Now, consider the following constraints:

1. **Rational Staking Constraint:**

$$\mu_{30} R \leq (f_a v) + (s_{30} d_b) \qquad \forall a \in A, b \in B$$

2. **Trading fees constraint:**

$$d_b \leq 0.4 \qquad \forall b \in B$$

   This constraint enforces the stacking principle of the discounted trading fees. Since users can currently get a maximum 55 percent discount on their own trading fees, setting $d_b \leq 0.4$ ensures that at least some fees are paid on all taker orders. Alternatively, adjust this constant to set $0.45 - \max d_b \quad \forall b \in B$ to a percentage that is okay.

3. **Limit on Required Volume**

$$v \leq 10000$$

4. **Bounds on R**

$$50 \leq R \leq 5000$$

5. **Bounds on Trading Fee Discount**

$$d_1 >= 0.05$$

$$d_b + 0.05 <= d_{b+1}$$

   Here, we constrain the lowest tier of trading fee discount for the referrer to be at least 5 percent and for each tier to be at least 5 percent higher than the last.

6. **Nonnegativity:**

$$f_1, s_{30}, R, v, d_b \geq 0$$

We used MINOS 5.51 on AMPL to solve our linear program with resulting decision variable values reported in Table 4. As discussed above, certain parameters were initialized with an educated guess, but a curious reader can challenge our assumptions by inserting different values.

The .mod and .dat files can be found on our repository.

Decision Variable	Value
R	50
requiredVolume	577.125
Bronze Discount	0.05
Silver Discount	0.1
Gold Discount	0.35
Platinum Discount	0.4

Table 4: Values of decision variables.

6.3.1 Sensitivity analysis

We observe that changing the average price of dydx changes the required volume to be traded by the referred user. This seems intuitive because the referrer needs her referral to trade more in order to offset not selling her DYDX at a higher price (assuming the referrer's trading volume is constant).

Further work: Recall that the objective of our linear program is to *minimize* the losses to dYdX. The interpretation of the dual of our linear program is the referrer is trying to *maximize* her gain from the referral program. In linear programming, this examination of the dual is called shadow pricing. We can look at the shadow pricing of our model along with the slackness of certain variables to gain a better idea of the mathematical foundations of our recommendation.