

Liquid Interest Rate Markets through Olympus Bonds

Zeus, Indigo

Olympus DAO

Abstract

In this paper, we will highlight some disadvantages of the existing Olympus model on a long-term timeframe and the importance of resilience and the ability to thrive in a steady-state. We present potential changes to the model to better address periods of slower growth or contraction, namely an emphasis on the bond mechanism to drive not only treasury flows, but market structure and behavior as well. This paper proposes an evolution from the existing staking-centric model in favor of a bond-centric model, where incentives lie primarily in holding future-dated OHM (commonly known as bonds). This structure adds assurances for the network, enables the development of liquid interest rate markets to facilitate healthy credit markets, and promotes closer alignment with network participants.

Background

All currency and money, be it USD or BTC, are ultimately backed by collective faith. As a society, we choose to imbue resilient systems with our belief in value itself, and utilize that system to exchange within an economy. OHM's path to currency has been, and will continue to be, marked by trials of extreme conditions and behaviors; only through fire can it prove its resilience and antifragility.

Olympus, like any currency issuer, has two main tools at its disposal.

1. Liquidity
2. Interest Rates

We see the treasury's control over liquidity manifested most intuitively as ownership of liquidity pools and the ability to increase or decrease liquidity through reserves. The protocol can always generate or diminish liquidity, either by purchasing or selling liquidity shares on the open market with OHM or reserves, or by creating or destroying liquidity shares directly with the addition or removal of OHM and reserves. Though the treasury has been passive to date, it is worthwhile to note that it *can* take active control over not only market liquidity, but pricing as well – as the dominant owner of the market, the treasury has the power to set market price.

The treasury also influences liquidity through bond sales. However, unlike traditional supply/demand-agnostic automated market maker (AMM) liquidity, bond liquidity is isolated and directional. Bonds that accept reserves and pay OHM only provide supply and satisfy demand, whereas bonds that accept OHM and pay reserves only provide demand and satisfy supply.

We see the treasury's control over rates in the form of the staking reward rate, which creates a target rate of growth and baseline OHM-denominated rate of return for the network; opportunities yielding more than this will be undertaken, while opportunities yielding less will not. This has historically been manifested in bond yields, which will trade at some market-driven

premium to the staking risk-free rate. The premium paid to bonders over stakers is most easily explained as a cost of illiquidity (though transaction fees and capacity play a role as well).

The existing structure of the Olympus market is characterized by aggressive targeted growth denoted by the staking reward rate. High staking yields facilitate high bond capacities, allowing the treasury to rapidly bootstrap itself. This configuration was extremely successful for the early days of the network as it targeted zero-to-one growth. Given the strong macroeconomic tailwinds of 2021, the configured rates proved to be appropriate in matching network growth with supply growth; though volatility was experienced throughout, the protocol was successful in maintaining relatively consistent pricing throughout the year.

A few shifts diminished the efficacy of this model. First and foremost, a proliferation of copycat projects (“forks”) adopted it based on the belief that it inherently generates value. As early projects with no real purpose or long-term intentions, these forks maximized (and marketed) the native token yields, pulling liquidity and mindshare away as it appeared to work in the short term. This inappropriate utilization of the model led to suboptimal outcomes, which in turn cast doubt upon Olympus itself; as with any tool, success is derived from the way and the reason it is used. With these forks came a tidal wave of (often unsophisticated) new investors who, misunderstanding the dynamics of the model (in large part due to the marketing of these forks), overleveraged Olympus, leading to a debt unwind and liquidity crunch. Both of these dynamics were further exacerbated by a macroeconomic shift out of an easy money environment, creating an additional headwind that compounded the previous dynamics and led to a compression in the monetary premium of the network.

While we believe these to be transient, they have exposed areas of improvement in the existing model; mostly in regards to staking. Staking serves two purposes: to protect holders from dilution, and to keep supply out of the external market (where OHM is traded against other assets outside its economy) while the treasury conducts market operations. However, the generalized rebasing structure may not be the best suited to accomplish these goals.

Proposed Changes

We begin with the addition of **internal bonds**. These are similar to the **external bonds** we know today; a user trades a quote token with the protocol in return for an amount of OHM in the future. Internal bonds differ from external bonds in that the quote token (the token given to the protocol) is OHM rather than an external asset (i.e. DAI, ETH, etc). For each 1 OHM deposited, the depositor is given a note giving claim to X OHM at Y date, where X is assumed to be greater than 1.

Bonds are tokenized, both as unique ERC-1155's and as composable ERC20's. The NFT approach is appropriate for fixed-term bonds, where each deposit has a different expiration; the ERC20 approach is appropriate for fixed-expiration bonds, where each deposit expires at the same time. Internal bonds will focus on fixed-expiration, as ERC20's more easily enable liquid secondary markets in which users can pool OHM alongside a bond token (i.e. OHM-Mar'22) to allow trading to and from different expiries. This is a crucial component to the overall approach; we will explore the implications and benefits of this in a later section.

In concert with tokenized bonds, the base staking reward rate is reduced. This does not necessarily mean that previous yields are no longer available; rather, this means that to access them, a user must enter a bond position and take on risk of illiquidity. It also means that holders can compete with each other for higher rates, allowing the market to drive real interest rates. Whether OHM yields less or more becomes a decision made by the market.

Within this structure, we can think of the staking reward rate as the baseline rate of return within the Olympus economy; activities (bonds, lending, yield farming, etc) will only find participants if their yields are greater than sOHM. With internal bonds providing an alternative to staking, ideally available beyond just Ethereum to enable wide participation, lower sOHM yields are to the benefit of both holders and the protocol. Holders benefit from greater utility, as opportunities that were previously unable to compete with astronomical staking yields are now viable. Meanwhile, the protocol's cost of capital can be reduced, as staking's no-strings-attached expense to holders (where beneficiaries can be short-term players) is diminished.

Tokenization

Bond tokenization carries massive implications for the market structure and behavior of Olympus. Today, we see two contrasting dynamics:

- Staking is completely liquid (there is infinite liquidity and zero slippage between sOHM and OHM, which the protocol keeps liquid against external assets like DAI and ETH) and yields at a protocol-dictated rate.
- Bonds are completely illiquid and yield at a fixed but market-dictated rate (determined by the discount, which floats based on supply and demand, at the time of purchase).

The liquid nature of staking is beneficial, as illiquidity carries a premium in any market and we want to minimize cost for the protocol. At the same time, the inconsistency of bond yield is beneficial, as it leaves the ultimate decision of rates onto market participants who, with enough participation, will compete them down to their optimally efficient point. Bond tokenization and secondary markets combine the benefits of both while discarding their negatives.

With liquid secondary bond markets, we retain the illiquidity of bonds on a protocol level, but not an individual level; while that supply is locked, it may not be the bond originator who takes on the illiquidity. Instead, a bond holder can sell their exposure to LPs in exchange for liquid OHM. The protocol still gets what it wants while removing a burden from the holder; with this, the illiquidity premium should diminish.

By offloading the bulk of staking activity onto internal bonds, we also pull in market-dictated rates; interest rates on OHM are thrust into a competitive, protocol facilitated environment (as opposed to a protocol mandated environment, as with sOHM yields) in which players can extract and diminish yields by optimizing their individual strategies. Rates should fall naturally over time to efficient levels as the market matures.

Tokenization across multiple expiries also forgoes the infinite liquidity seen with sOHM today, which provides segmentation benefits that we will discuss in a later section.

Liquidity

The dynamic generally seen with external bonds is: a participant unstakes an amount of OHM, sells it on the market for the treasury's desired asset, and bonds that asset in exchange for a greater amount of OHM. It can be viewed as an arbitrage opportunity allowing holders willing to take on illiquidity to expand their position in the network.

With this, we see that the treasury's ability to issue bonds is capped by the downward pressure they produce. We will not be able to build highly liquid secondary bond markets so long as bonds are scarce as a result of an external capital constraint. If we were to aim for 30% of OHM supply paired against bond tokens, the treasury would need to sell at least 30% of supply. This would both realize a significant amount of market cap (and likely depress valuation as a result) and dilute existing holders. Not only that; this must be repeated each time the bond in the pair expires.


Internal bonds fix the former; by allowing holders to give their OHM to the treasury directly (instead of selling for a different asset first) we are able to generate infinite bond liquidity relative to the supply of OHM. In the previous hypothetical, 30% of market cap no longer needs to be realized. Rather, 30% of supply must simply be willing to take on a pseudo-illiquid position. Those who do also become the beneficiaries of the dilution that these bonds produce.

Segmentation

A bond-centric approach creates enormous benefits around risk management and supply segmentation. To understand this, we first need to understand what the end result of this system looks like. An illustrative state in a bond-centric approach, where OHM supply is ~10 million:

- 1.0 million OHM is paired against external assets in protocol-owned liquidity.
- 1.0 million OHM is staked as sOHM yielding 100% APY.

- 1.0 million OHM is staked or deployed in various opportunities earning, on average, greater than 100% APY paid in external assets (this could be lower if depositors place a premium on non-OHM yields).
- 5.0 million OHM is locked into bonds, originating from OHM and other assets (once created, they are indistinguishable). That supply is broken down into:
 - 1.0 million in March 3rd expiration;
 - 1.0 million in June 3rd expiration;
 - 1.0 million in September 3rd expiration;
 - 1.0 million in December 3rd 2022 expiration;
 - 1.0 million in December 3rd 2023 expiration;
- ~2.0 million OHM is paired with bond tokens to create secondary markets. The ratio of ohm pooled against a given expiry sets the implied price of the bond as dictated by the market. Those pools are:
 - 0.61 million OHM : 0.73 million OHM-Mar'22
 - 0.50 million OHM : 0.73 million OHM-Jun'22
 - 0.41 million OHM : 0.73 million OHM-Sep'22
 - 0.33 million OHM : 0.73 million OHM-Dec'22
 - 0.12 million OHM : 0.73 million OHM-Dec'23



Illustrative scenarios:

- 1) Alice wants to quickly purchase Dec'22 bonds with 100k OHM. There is only 340k OHM paired against bonds representing 730k OHM at maturity in December 2022. In other words, the market is valuing the December maturity at a 54% discount, implying a yield of 115% in the one year interim period. Suppose Alice purchases bonds with all 100k OHM. This action "skews" the prevailing yield curve, changing the discount from 54% to 22% and lowering implied yield from 115% to 28% (an extreme scenario). As Alice bids, liquidity is potentially unlocked from two sources. Bob, who currently holds Dec'22 bonds, sells them to buy Jun'22, Sep'22, or Dec'23 bonds, capturing the yield curve imbalance that the buyer has created. Additionally, the protocol can begin to sell more Dec'22 bonds and fewer other expirations.
- 2) Alice wants to quickly sell Dec'22 bonds for 100k OHM. There is only 340k OHM in the liquidity pool, and the pool trades in line with the yield curve. As Alice sells, liquidity can be unlocked from other players, who sell their short term bonds and purchase Dec'22 to capture the skew in the yield curve. However, other players could instead choose to cooperatively play PvP, refusing to provide new liquidity unless the Dec'22 market is highly discounted. Alice is forced to accept the discount or slow/stop her selling.
- 3) Alice has taken out \$50m debt against Dec'22 bonds worth \$100m. OHM falls 10% against DAI, leading the market to view this position as at risk. The market has three options: they can bid the December expiry, increasing the value of Alice's collateral and potentially averting her liquidation; they can build up liquidity, reducing the impact of, and capitalizing on the volume generated by, the potential liquidation; they can remove liquidity. Removing liquidity will hurt the credit rating of the market, deterring future lenders and raising the future borrow cost of all players. However, in this situation, lenders have over-provided against weak collateral, building bad debt with insolvency risk. If liquidity providers for Dec'22 remove their liquidity, Dec'22 bonds become illiquid and are thus unavailable for use by liquidators to recoup the loan principal.

It is important to understand that liquidity for bonds is guaranteed at expiration, but can only be expected prior to maturity. We can expect responsive liquidity during normal conditions, as profit motivated actors fill market inefficiencies and maintain the yield curve to maximize their individual returns. However, segmentation allows the market to circuit break itself in extreme circumstances. Liquidity and quality of execution may diminish during large directional flows, increasing volatility of a specific market segment but diminishing volatility in the market as a whole.

Yield Curves

This system is in large part reliant on the formation and analysis of yield curves. By allowing independent pricing of the same asset along multiple timeframes, the market can price in varying expectations manifested as variable yields and visualized by the yield curve.

Yield curves should be driven by the expected return of compounding short term bonds (i.e. one quarter out). A six month bond should yield approximately three month yield squared. A twelve month should yield three month⁴ or six month². With this, expectations on future yields become important, as it drives the outcome of compounding over some time frame.


We can extrapolate yields from bond prices with the following equation:

$APY_{bond} = p^{-1/n} - 1$, where $p = \frac{\text{present value}}{\text{future value}}$ and $n = \text{bond tenor in years}$. For example, a six month bond discounted 50% implies an annualized yield of 300%.

With this in mind, let's go back to the previous example:

- ~2 million OHM is paired with bond tokens to create secondary markets. It is December 3, 2021. Those pools are:
 - 0.61m OHM : 0.73m OHM-Mar'22
 - 0.50m OHM : 0.73m OHM-Jun'22
 - 0.41m OHM : 0.73m OHM-Sep'22
 - 0.33m OHM : 0.73m OHM-Dec'22
 - 0.12m OHM : 0.73m OHM-Dec'23


Here is that yield curve:



This is the most intuitive shape. This curve should form in environments where rates are expected to remain constant across the measured timeframe. As expirations extend further out, yields increase because there is greater illiquidity to compensate for. On this curve, an arbitrageur could sell December 2022 bonds and buy September 2022 bonds, correcting the slight deviation from the curve and securing a greater position when they reach expiration. Thus,

we should expect to rarely see significant deviations from the curve, given the bond token is adequately liquid for rate-arbitrage.

Here is a different scenario:



In this case, the market is pricing in diminishing rates over time. The expectation is that short term yields will be lower in 2023 than they were in 2022. Thus, December '23 bonds should trade at a lower *annualized* yield because the nominal yield earned will still be greater than a short-term compounding strategy.

There is a reasonable expectation that early on, these markets operate inefficiently. Many expirations may deviate from the correct yield curve, sometimes by significant margins. They may even form yield curves that do not accurately reflect the state of the market. However, this inefficiency creates opportunity for sophisticated actors to capitalize on, generating new activity and bringing new participants into the ecosystem. Over time, as competition over rate-arbitrage opportunities picks up, we can expect yields to compress to their optimal levels based on the present and future risk-free rate of return (the staking reward rate).

Implications

One of the most significant components of this system is the market power it imbues to the protocol. In the current environment, 100% of OHM supply is effectively liquid at any given time, and thus the protocol must account for scenarios in which 100% of supply fights against it. This feeds the concept of risk-free value; the point at which the protocol has greater capital than the market has capitalization. The bond-first structure shifts this dynamic in favor of the protocol. Now, not only is a significant portion of total supply illiquid and unable to interfere with the treasury's objectives; the treasury now knows exactly when supply will become liquid, and can influence the rate at which supply becomes illiquid.

This structure creates a massive sink and utility for supply – bond token liquidity. This is a relatively risk-free endeavor; there is impermanent loss involved, but it can be quantified and is not particularly significant until deep discounts (i.e. an 0.8 : 1 pool would experience 1.1% IL upon its' return to peg). Trading fees should be more than enough to compensate for IL, and if not, conservative incentives will fill the gap. Bond token LP is a productive activity that generates value for the ecosystem, unlike vanilla staking, and thus even with an incentive expense this can be viewed as a step forward.

We can expect yields to depress over time as competition ramps up. This is to the detriment of an individual player but to the benefit of the protocol.

Vaults

We can expect to see the formation of bond vaults. These protocols will provide open automated bond strategies to market participants, abstracting away the market competition to offer a simple interface akin to staking today. A vault would accept user deposits of OHM, and deploy that OHM into various expiries and strategies along yield and risk curves.

For example, a vault may choose to hold 10% of its deposited capital in sOHM, generating the lowest yield but offering the greatest liquidity. This can be thought of as liquid cash-on-hand, used to honor redemptions when users withdraw from the vault. Under normal circumstances, this will be more than enough to ensure users can withdraw without notice whenever they desire. An additional 40% of deposits might be used to purchase near-dated bonds, which should intuitively be the most liquid on secondary markets. Even in the case where secondary markets dry up for these bonds, they will expire within a short period of time and thus are relatively safe. The remainder would be deployed into longer dated bonds, which carry the greatest illiquidity risk but also offer the greatest yields. The yield for depositors in the vault becomes the aggregate return of all of these positions.

Vault systems do carry the risk of bank run. Since vault deposits are converted into positions with varying levels of illiquidity, in the case where all depositors seek to withdraw, they may find themselves unable. This is not necessarily the case; if secondary markets remain liquid for bond tokens at fair pricing, a vault will remain perfectly solvent even in a mass withdrawal. However, in a more systemic event, secondary markets may begin to break down and this becomes a real concern.

The main counter to this concern is the persistent availability of native staking. Those who find this illiquidity risk unacceptable are more than welcome to simply remain in the protocol facilitated environment, where liquidity back to OHM is guaranteed. The tradeoff for this safety is, of course, a lower expected return. While vaults can have some benefit of increased liquidity over an individually run bond strategy, this should not be considered the main purpose of participation. Rather, vaults should serve the purpose of automating strategies for less active participants, or those looking to leverage economies of scale.

The rise of these service protocols, which cannot be guaranteed but is to be expected, will be interesting at the very least. In a perfect world, they will generate greater participation in bond markets, contributing large-scale competitive forces that keep markets liquid and efficient.

Repurchase Markets

The vault structure may create demand for OHM repo markets. At a high level, these repo markets would allow a participant to borrow liquid OHM against an illiquid bond (or vice versa), likely for a short period of time.

For example, suppose a vault is fielding mass withdrawals, and has depleted its liquid reserves (sOHM and closest-dated bonds). It is now faced with a short list of choices. It can:

- Provide depositors with bond tokens.
- Remove bond token liquidity to provide liquid OHM, but reduce liquidity in the overall bond market and risk systemic yield curve imbalances.
- Refuse to honor withdrawals for a period of time.
- Borrow liquid OHM against bond tokens to honor short-time liabilities, to be paid back with interest at a future date.

Refusal to honor withdrawals, or servicing of withdrawals with illiquid or pseudo-illiquid positions, is likely to be accompanied by dissent from the vault users. Removing bond token liquidity may be viable, but it risks creating systemic issues that are harmful to the network and, by extension, the vault. Repo markets provide a viable alternative at the cost of interest.

As with vaults, it remains to be seen whether such services will ultimately be offered. There are abundant implementation questions when it comes to these topics (though there are promising projects addressing these questions already). This should, however, highlight the expansive potential of this structure.

Drawbacks and Tradeoffs

This proposed system adds a great amount of complexity. With that complexity comes productivity and value, but it is reliant on efficient markets maintained by competition among network participants. Without competition, markets will likely trade at less-than-optimal rates, participation in the network becomes more difficult, and the entire thing does more harm than good. While this does not seem like the most likely outcome by any means, it is something that

should be strongly countered through thorough preparation and the continued maturation of Olympus market prior to large-scale introduction.

This is also self-referential to a moderate degree. There is far more circular activity generated than within the current structure. Over a long time frame, this is important; crucial, really. We want circularity; activity should flow from one place to another without leaving the economy. However, self-referential dynamics are something that must be earned.

In Conclusion

Through the proposed structure of this paper as well as its natural second-order effects (to be discussed in an additional paper), the mental model of the network changes. Yields go from perpetual projection to fixed sum and fixed time-period. The concept of risk-free value and a singular floor erode in favor of continuous treasury utilization. Behavior around staking and participation shift from “what can the network do for me?” to “what can I do for the network?” Though in the case of the latter they more so become one and the same, it's meant to highlight a point. The way in which we think about Olympus must evolve if and when this is implemented.

Most of these concerns should be remedied with time and maturity. The intent of this paper is not to push it through tomorrow. This is a marathon, not a sprint. This can act as a guiding light for months and potentially years to come. It's not going to come together soon, nor is it all going to come together at the same time. What we seek to build with this system is a well-regulated interest rate market for OHM, providing guidance for bond discounts and lending markets and an outlet for healthy market competition. We can make progress prior to the items in this paper, and we can make progress after the items in this paper. We should focus on the end result we are seeking, determine whether this is a viable and constructive path toward those results, and begin executing if the answer is yes.