Impermanent Loss in Extreme Market Conditions

Sean Dawson* Nick Forster†

July 16, 2021

Abstract

We analyse the impermanent loss (IL) Lyra's liquidity providers (LPs) will experience in the event of a sudden spike in volatility. As the recent May 2021 crash demonstrated, these events pose a significant risk to automated market markers (AMMs). A detailed preparation for these scenarios is essential to ensure minimal risk to LPs. In this report we study the analytic properties of Lyra's mechanism alongside numerical experiments conducted in Mathematica with a focus on options on Ethereum (ETH). This analysis is done in the absence of fees and assumes perfect optimization on the part an arbitrageur. These insights inform our choice of various parameters in our model, in particular the skew update parameter α and the number of contracts $\mathcal{S}$ that define one standard size. In the interest of full transparency we have also released the Mathematica code used in this paper. Interested members of the community are encouraged to investigate other possible scenarios themselves.

1 Lyra's AMM Model

In this section, we briefly recap the essential workings of Lyra's pricing mechanism. Consider an underlying asset (Ethereum) with spot price S and an option on this underlying with strike K, time to expiry τ, implied (trading) volatility σ and risk free rate r. In the Black Scholes pricing model, the cost of a European call option is given by

$$C = N(d_1)S - N(d_2)K\exp(-r\tau) \tag{1}$$

where

$$d_1 = \frac{1}{\sigma\sqrt{\tau}}\left[\ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)\tau\right], \tag{2}$$

$d_2 = d_1 - \sigma\sqrt{\tau}$ and $N(\cdot)$ is the standard normal cumulative distribution function.

We now recall Lyra's mechanism for pricing options and volatility. For more details, see our white paper[1]. Consider z strikes K_i $(i = 1, \ldots, z)$ all with expiry τ. For each expiry there is a baseline volatility b that all listings in this period share. To account for the effect strike has on volatility, we assign a scalar $R_i \in \mathbb{R}^+$ to each K_i which we call the skew ratio. The trading volatility σ (used in (2)) for a listing (K_i, τ) is given by $\sigma_i = R_i b$. When a trade of n contracts is made of the listing (K_i, τ), we update skew and baseline volatilities as follows:

$$(R_i, b) \to \left(R_i + \alpha\frac{n}{\mathcal{S}}, b + \beta\frac{n}{\mathcal{S}}\right). \tag{3}$$

In (3), $\alpha, \beta \in \mathbb{R}^+$ are constants called the skew and baseline update parameters respectively, while $\mathcal{S} \in \mathbb{R}^+$ determines the number of contracts in one standard size. When the pool sells contracts n is positive and

*sean@lyra.finance

†nick@lyra.finance

[1] https://www.lyra.finance/files/whitepaper.pdf

vice versa. In this paper we only consider call options sold by the protocol since analagous statements can be made for puts. In the following analysis we assume $r = 0$, no trading/vega utilization fees, and the fact that an arbitrageur can execute the most optimal trading strategy. This both simplifies our computations and ensures we are considering the worst case scenario facing LPs.

2 Impermanent Loss

In periods of extreme market upheaval, the volatility of all listings should spike dramatically. Since Lyra uses an AMM to price volatility, such an environment exposes liquidity providers to impermanent loss. We now wish to quantify this loss under different adverse scenarios. For simplicity, we assume only one strike (K) and one expiry (τ) with trading volatility σ_0 and true market volatility g.

Definition 1. The make up $\eta(\sigma_0)$ is the number of contracts of that listing that must be sold by the AMM to increase σ_0 to g.

For a given number of contracts n sold by the AMM, the new trading volatility is given by

$$\sigma(n) = (R + \alpha \frac{n}{\mathcal{S}})(b + \beta \frac{n}{\mathcal{S}}) \tag{4}$$

where $\sigma_0 = \sigma(0)$. The make up η is found by setting (4) to g and solving for n. This yields

$$\eta = \frac{\mathcal{S}}{2\alpha\beta} \left(\sqrt{4\alpha\beta g + (\beta R - \alpha b)^2} - \alpha b - \beta R \right). \tag{5}$$

Define $X(n; g)$ as the differential between two options priced using these two volatilities, i.e.

$$X(n; g) \coloneqq C(g) - C(\sigma(n)).$$

The maximum impermanent loss occurs when an arbitrageur continually buys an infinitesimal number of contracts until the AMM volatility equals g. This gives us the following.

Definition 2. The maximum impermanent loss $\ell(\sigma_0, g)$ is the integral of $X(n; g)$ from $n = 0$ to $n = \eta$, i.e. $\ell(\sigma_0, g) = \int_0^\eta X(n; g) dn$. Equivalently, we have

$$\ell(\sigma_0, g) = \eta C(g) - \int_0^\eta C(\sigma(n)) dn. \tag{6}$$

Remark 1. One can make sense of (6) by thinking of the first term $\eta C(g)$ as the market value of the η contracts at the true market volatility g. The integral is the "discounted" cost of these options (since it takes η contracts to "pull up" the AMM volatility to the true market value). We stress that (6) is a strong upper bound for impermanent loss, since the presence of fees will significantly dampen this effect. Further, this also assumes a single arbitrageur. In reality, there will be multiple such actors and their competition will result in less IL for the pool. We now illustrate an example of (6) in market conditions similar to those of the May 2021 crash.

Example 1. Consider an asset (Ethereum) with spot price $S = 2000$ and a listing with strike/expiry combination $(K, \tau) = (2100, 28/365)$. Let the baseline volatility and corresponding skew ratio be $b = R = \sigma_0 = 1$. Suppose the number of contracts in one standard size is $\mathcal{S} = 20$ and the parameters α, β are $(0.0125, 0.01)$ respectively. These are the current values Lyra uses on testnet. If the true market volatility spikes to $g = 3$, then the number of contracts required to update σ to this value is found using (5). This yields $\eta \sim 1305$ contracts. Using (6) we compute the impermanent loss to be $\ell(1, 3) \sim \$304000$. It is important to realise that had an arbitrageur, Alice, bought 1305 contracts in a single trade when the trading volatility was originally $\sigma = 1$, then the pool would experience no impermanent loss. The most profitable strategy for Alice is to to buy 1305 total contracts in immediate succession dn at a time where dn is as small as possible. In reality, this is highly unlikely to be achieved for reasons mentioned earlier. Again we stress this level of impermanent loss is a clear upper bound to what the pool would face.

In Figure (1) we demonstrate how this impermanent loss changes when varying the true market volatility g. We also show how the IL curve is deformed under different choices of τ and spot price S. As expected, shorter time to expiry and strikes out of the money result in smaller IL. Note that the parameters set here are calibrated for ETH, so for higher volatility assets parameters like standard size will likely be more sensitive to a given sized trade. It is worth noting that for an asset like ETH, market volatilities for options with a time horizon of > 7 days that far exceed $g = 3$ are unlikely. Nevertheless, we include these unlikely scenarios in the Appendix for completeness. This example shows that at least for this one listing, IL is generally capped around $\sim \$300,000$ for $g = 3$. In Section 4 we discuss the more general scenario of multiple listings. In the next section we investigate how the parameters α, β and $\mathcal{S}$ affect the impermanent loss.

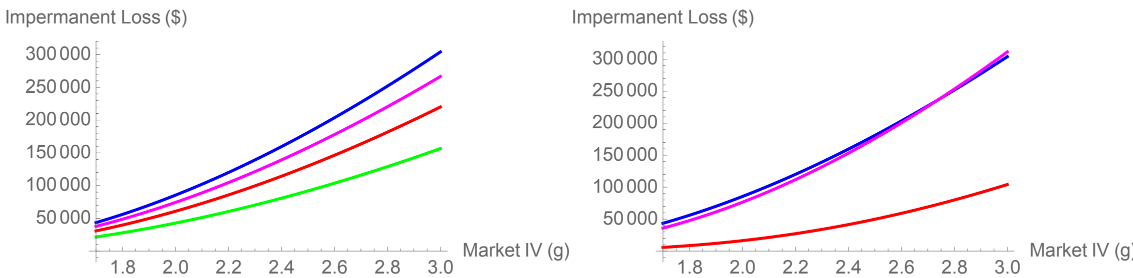

Figure 1: Impermanent loss versus true market volatility in Example (1) with varying a) expiry $(7, 14, 21, 28)$ day expiries (green, red, magenta, blue) (with spot 2000) b) spot $(1000, 3000, 2000)$ (red, magenta, blue) (with expiry 28 days).

3 Volatility Parameters

We now investigate the effect the parameters α and $\mathcal{S}$ have on the impermanent loss. To do this, we consider the listing traded in Example 1 with $(S, \tau) = (2000, 28/365)$ but we vary these two parameters. We present our results in Figure (2) below. Figures a) and b) represent the effects of varying α and $\mathcal{S}$ respectively. We do not investigate changes in β since $\mathcal{S}$ has the same effect and is more general. In a) we plot the IL for $\alpha = (0.0075, 0.0125, 0.0175)$ corresponding to the red, blue and magneta lines respectively. Similarly, b) is done for $\mathcal{S} = (30, 20, 10)$. In all examples, decreasing the sensitivity of volatility to a trade results in greater impermanent loss. We note that changes in $\mathcal{S}$ have a greater impact on impermanent loss than changes in α. This is because α is localised to a single strike while $\mathcal{S}$ extends over both the skew and the entire expiry. This effect becomes more significant when each expiry has multiple strikes. We now address this and other concerns in the final section.

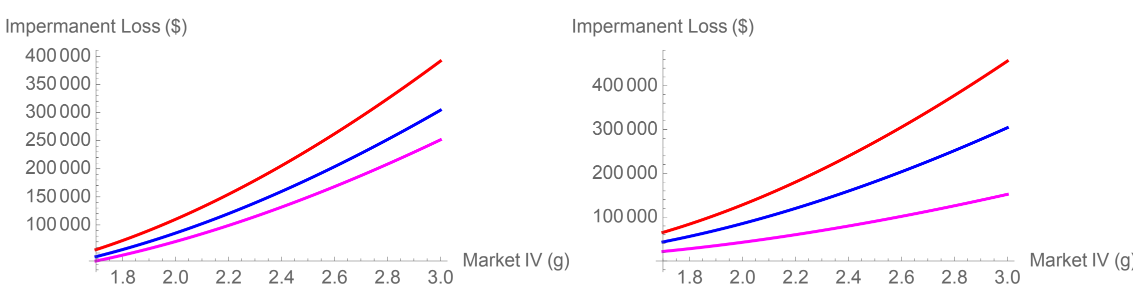

Figure 2: Plots of impermanent loss resulting from varying α (a) and $\mathcal{S}$ (b).

4 Impermanent Loss over an Expiry

We now consider a more realistic scenario where each expiry consists of multiple strikes. Without loss of generality, suppose there are z strikes $\boldsymbol{K} := (K_1, \ldots, K_z)$ corresponding to the expiry τ with trading volatilities $\boldsymbol{\sigma} := (\sigma_1, \ldots, \sigma_z)$ and true market volatilities $\boldsymbol{g} := (g_1, \ldots, g_z)$. The most optimal strategy for Alice the arbitrageur would be to compute the difference between the AMM and market price for all strikes offered for a single contract. She then purchases an infinitesimal amount of the contract with the largest difference in her favour and repeats the process until the market volatility of all listings equals those which the external market offers. With each (infinitesimal) contract Alice purchases, the baseline volatility across all strikes increases, meaning the impermanent loss of each subsequent contract becomes progressively less until it reaches zero. We illustrate this in the following example.

Example 2. We continue Example (1) but with strikes $(1800, 2000, 2100, 2300, 2500)$, expiry 28 days, spot 2000, $\beta = 0.01$ and $\mathcal{S} = 20$. We assume Alice follows the optimal strategy outlined above. In Figure (3) a) we show the impermanent loss the AMM experience for each trade Alice makes by varying α as before.[2] Consider for instance the blue curve ($\alpha = 0.0125$); as the attack goes on, the baseline volatility rises from $\sigma_0 = 1$ to a final level $\sigma = 2.275$ after 2550 contracts have been exploited for impermanent loss. This means that each subsequent contract yields less profit for Alice. In b) we show the effect of varying $\mathcal{S}$ as before. Clearly, changing $\mathcal{S}$ has a more pronounced effect since it extends over the entire expiry. We note that the exploit is over when the AMM volatilities of all strikes equals $g = 3$. Since we initialised the skew ratios of all strikes to be the same at $R = 1$, the final skew ratios of all strikes are also equivalent at $R = 1.31875$ with new baseline $b = 2.275$.

In Figure (4) we repeat the computations done in Figure (2) but now considering all strikes in the given expiry (with optimal strategy). We note that the effect of $\mathcal{S}$ on IL is substantially more pronounced for the same reasons outlined above. The maximum impermanent loss is approximately \$866,000 for $\mathcal{S} = 30$. For the parameters we have chosen to initialise Lyra (blue curve), the maximum impermanent loss does not exceed \$577,000. It is worth noting that this IL does not scale with an increase in pool liquidity, which means that as total value locked (TVL) increases, the relative threat of such a move decreases. For shorter expiries, the IL is naturally smaller, meaning the other listings in this round will have IL less than these examples (for the same g).

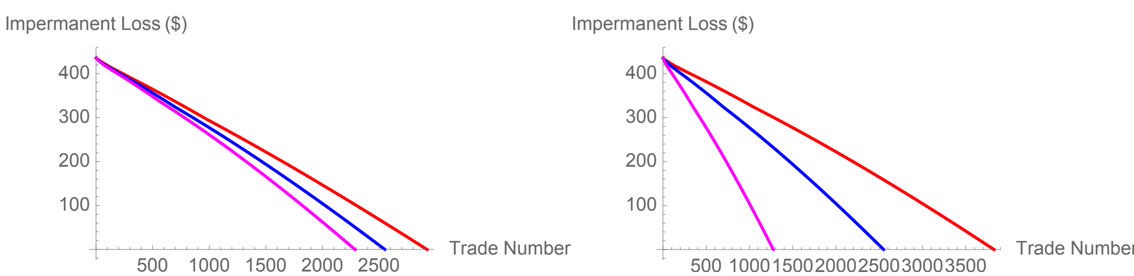

Figure 3: Impermanent loss per trade as we vary a) α and b) $\mathcal{S}$.

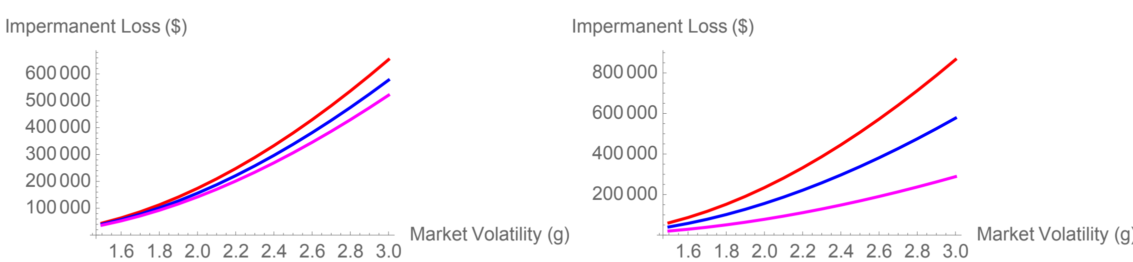

Figure 4: Total impermanent loss over the entire expiry with varying a) α and b) $\mathcal{S}$.

[2] We should mention that in this and the following computations, we assume that Alice buys one single contract at a time (instead of an infinitesimal). This was done to speed up the run-time of these experiments. The difference between a summation one contract at a time and an integral in these examples is exceedingly small with error ($< 0.5\%$), so we assume the two methods are essentially the same.

5 Conclusion

In this report we defined and derived an analytic expression for impermanent loss under Lyra's mechanism (6). Using this we constructed a numerical model inspired by the May 2021 crash and simulated how much impermanent loss liquidity providers would incur in such an environment. With the currently chosen parameters, impermanent loss seems to be relatively modest even in the worst case scenario. The addition of fees and the unlikelihood of an arbitrageur performing a perfect strategy means true losses will most likely be substantially smaller. In future investigations we hope to analyse data from the recent testnet launch in more intricate models that will give more insight into the most likely case fro impermanent loss.

A Simulating Larger Volatilities

In the interest of full transparency, we replicate all previous figures up to $g = 5$. Even though volatility above $g = 3$ is highly unlikely, we want LPs to be assured we are aware of and have considered these possibilities.

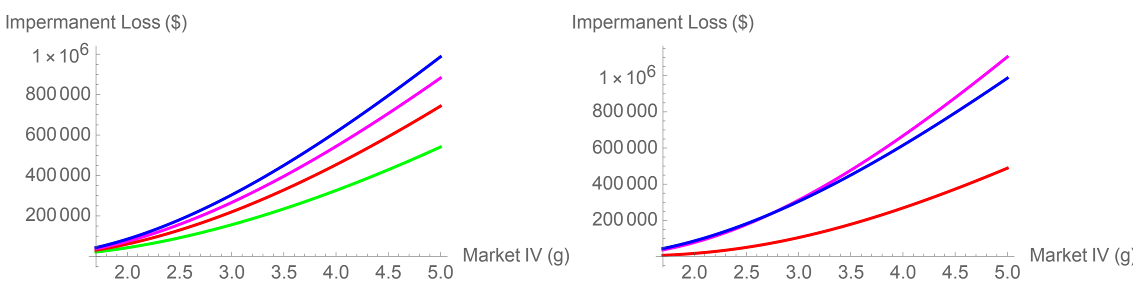

Figure 5: Extended version of Figure 1.

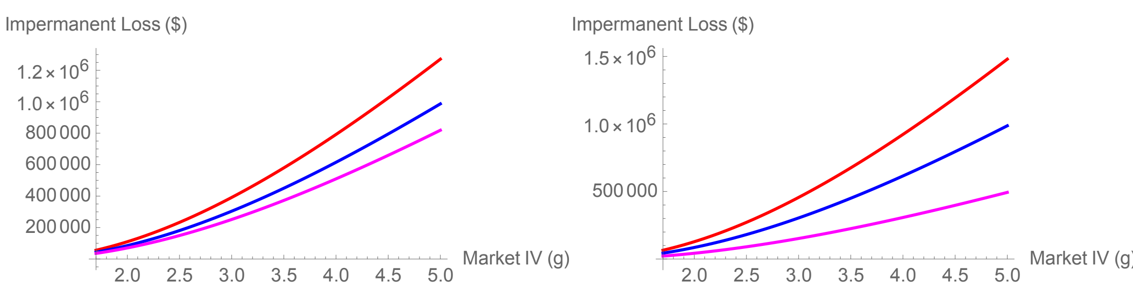

Figure 6: Extended version of Figure 2.

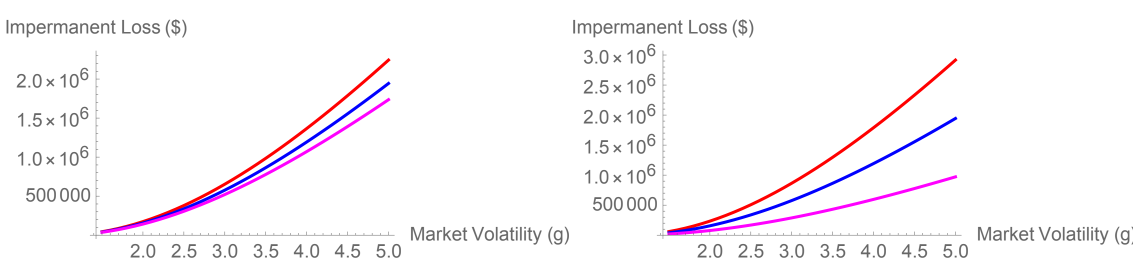

Figure 7: Extended version of Figure 4.