Lyra

Sean Dawson, Nick Forster, Dominic Romanowski, Michael Spain

May 2021

Abstract

Lyra is an automated market maker for trading options on Ethereum. Successful market making depends on the ability to price options accurately and to effectively measure and manage risk.

Two recent developments have made it possible to implement a traditional approach on-chain. The first is the advent of layer two, which makes computing aggregate risk practical. The second is the maturity of liquid spot markets, which enables protocol-to-protocol, or "composable" hedging of risk.

1 Introduction

In this paper, we present Lyra, a novel automated market maker (AMM) for trading European options. Prior attempts to build an options AMM on-chain have faltered due to either blunt, high-fee mechanisms or unacceptable risk incurred by liquidity providers (LPs). The former leads to untradable prices and low volumes, the latter causes the LP's probability of ruin to approach 100% over a long-term time horizon.

Autonomously undertaking hedging actions based on the net Greek exposure of the LP's is the defining idea that drives the sustainability of Lyra over the long term. This is achieved by:

- *Market driven pricing*: Options are accurately priced using market based, strike-adjusted implied volatilities (IVs) applied to a Black Scholes model. This returns the theoretical value (W).
- *Vega Risk Management*: Charge an asymmetric spread around W based on whether a given trade increases or decreases the LP's vega risk.
- *Delta Risk Management*: Hedge the LP's net delta by trading the underlying asset on a spot market.

Section 2 discusses the fundamentals of options trading, its associated risks, and AMMs. Section 3 outlines Lyra's pool structure for LPs. Section 4 describes the mechanism for determining the IV that the AMM uses to price options. Section 5 describes how the AMM calculates and manages risk for LPs. Section 6 describes the final price offered to traders. Appendices A and B provide a mathematical description and examples of the mechanism and how it behaves in response to trading. Appendix C proposes a more efficient pricing mechanic for future versions of Lyra.

2 Risk and Options Fundamentals

In this section, we outline the basic principles of options trading by introducing the Black Scholes model alongside the well known "options Greeks". We then briefly discuss recent developments in decentralized finance (DeFi), namely AMMs and how they can be used to facilitate options trading.

2.1 Options

A European option is a contract between two parties that gives the purchaser the right, but not the obligation, to buy (calls) or sell (puts) the underlying asset for a certain price (strike K) on a certain date (expiry T). Options allow traders to manage risk, obtain leverage, and construct desired payoff structures. The price of an option, denoted W, is computed using the Black Scholes model. The model takes 5 key parameters:

1. Time to expiry ($T - t$), where t is the present time
2. Risk free interest rate (r)
3. Price of the underlying asset at present time t (S_t)
4. Strike price of the option (K)
5. Implied annualized volatility of the underlying asset (IV)

Options markets are driven by disagreements about the volatility of the underlying asset until the time to expiry. Assuming there is a liquid spot market, every options trade can be transformed into a bet on volatility by trading the underlying asset (delta hedging). It is therefore crucial for any options protocol to determine an equilibrium level for volatility per option.

Options require active risk quantification and management across multiple dimensions. These dimensions are defined mathematically by the *Greeks*, which quantify the sensitivity of an option's value to changes in the Black Scholes parameters (holding all others constant).

The two key Greek risks to manage are:

- *Delta* $\frac{\partial W}{\partial S}$: the change in the price of an option given a $1 move in the underlying asset.
- *Vega* $\frac{\partial W}{\partial IV}$: the change in the price of an option given a 1% point change in IV.

Other Greeks include:

- *Gamma* $\frac{\partial^2 W}{\partial S^2}$: the change in an option's delta given a $1 move in the underlying asset.
- *Theta* $\frac{\partial W}{\partial t}$: the amount by which an option's value declines per day that passes.
- *Rho* $\frac{\partial W}{\partial r}$: the change in the price of an option given a 1% point change in the risk-free interest rate.

These last three Greeks are secondary to delta and vega in risk terms, since by managing delta and vega we are respectively balancing our exposure to gamma and theta. Rho risk is typically of much smaller magnitude than the others as the risk-free rate tends to be fairly stable.

2.2 Automated Market Makers

AMMs are agents that pool liquidity and make it available to traders according to an algorithm [5]. The smart contract paradigm introduced by Ethereum - in which code can store value and transform it programatically - has enabled the rise of AMMs. LPs deposit funds trustlessly and participate as market makers, according to the functions defined in the smart contracts. Traders have permissionless access to buy and sell the underlying product. Decentralised AMMs have been able to efficiently aggregate liquidity, improving the trading experience across many types of assets and products. However, LPs incur the risks of market making, known as "impermanent loss" (IL). Trading fees must exceed IL, which occurs when the fees from uninformed flow exceeds losses incurred from toxic flow as described in [3].

2.3 Options and AMMs

Options are leveraged, volatile products that incur large risks along multiple dimensions. This is why current AMM designs for other markets, whilst instructive, cannot be directly mapped to options. Existing options AMMs attempt to deal with this complexity by increasing fees to levels which almost certainly protect LPs, but result in extremely high prices and little trader interest. This tension is resolved in traditional finance through risk management techniques, allowing market makers to quote competitive prices whilst ensuring that they maintain high risk-adjusted returns.

A successful options AMM should emulate this process, hedging risk to enable competitive pricing. *Delta hedging* smooths the exposure of the AMM to a large directional move in the underlying asset. This is combined with an asymmetric spread which incentivizes trades that hedge the AMM by charging a fee that is skewed according to its current *vega* risk exposure.

In the next section we begin to describe Lyra's architecture as an AMM, introducing the liquidity pool structure and the process by which new expiries are listed.

3 Pool Structure

Lyra will accept a stablecoin as collateral, and will offer options in *rounds*. A *round* is defined as a 28 day period, with options tradable for 4 discrete expiries within that time (7, 14, 21, and 28 days from the commencement of the round). The liquidity will be split into two sub-pools:

1. **Collateral pool**: Collateralizes options and pays/receives premiums.
2. **Delta pool**: Hedges the delta exposure of the AMM by trading the underlying asset.

In the next section we describe the AMM's mechanism for determining a theoretical price for an option (W) through dynamic (IV) and skew parameters.

4 Options AMM

The goal for an options market maker is to find an IV value at which demand roughly equals supply. In this situation the AMM can collect fees on trades without taking on risk itself, as it is buying and selling options in equal amounts. The AMM is designed to respond to supply/demand to reach this IV level efficiently. This market-derived IV value is then used to calculate the Black Scholes price of an option (W). The logical flow of the mechanism is outlined in Figure 1. A mathematical description of the volatility and skew impact mechanisms is provided in Appendix A, with a geometric visualization presented in Appendix B.

4.1 Initialization

When an expiry T_j is listed on Lyra, a baseline volatility value IV_j will be initialized along with ratios of the listed strike volatilities to IV_j. These initial values will be derived from current market data, with IV_j taken from the 50 delta (at-the-money) strike. Following initialization, both IV_j and the strike volatility ratios (skew) will be determined by the supply and demand for options for a particular strike and for its associated expiry.

4.2 Standard Size

It stands to reason that the price impact of a trade is proportional to its size. An individual buying 100 options will likely cause the price to move higher than had they bought a single contract. The AMM captures this effect through the notion of a standard size (SS) which allows it to contextualize each trade and alter its pricing parameters in proportion to a trade's size. Standard size and the number of contracts traded are linearly related with a constant of proportionality χ. That is:

$$N = \chi n \tag{1}$$

where N is the number of contracts in a trade, n is the number of standard sizes and χ is a constant initialized when options are listed.

We initialize the SS according to the vega of the at-the-money strike for the 7 day out expiration, given an initial volatility level derived from current market data. The SS will be proportional to vega, with assets supporting options with higher vegas resulting in a lower SS. The lower the SS, the more sensitive IV will be to a given trade, and vice versa.

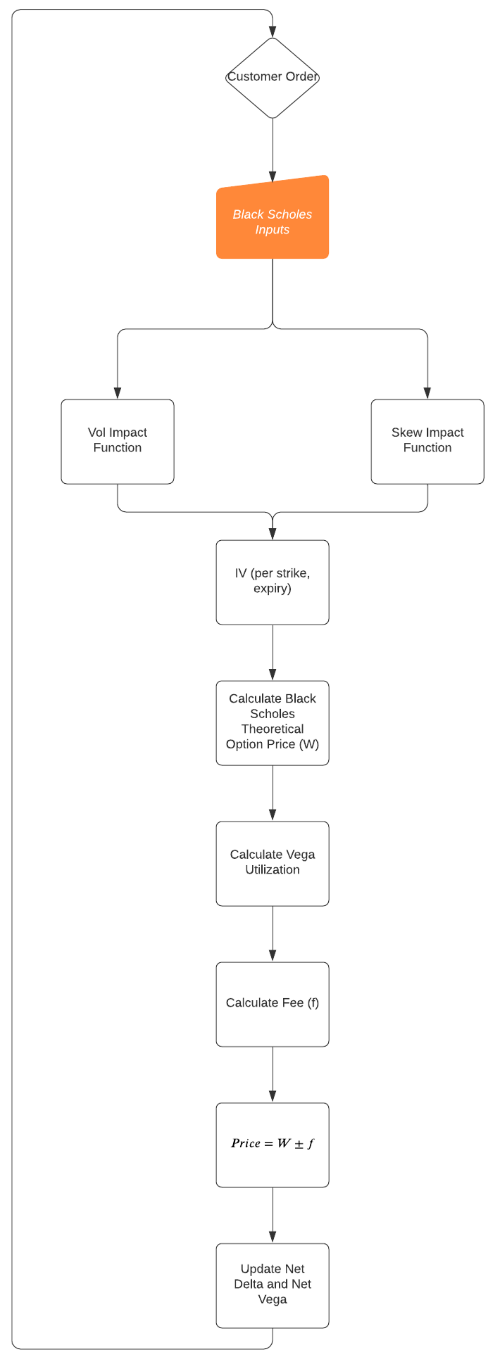

Figure 1: Lyra mechanism flowchart.

4.3 Volatility Impact

For every SS that the AMM buys or sells in a given expiry, the baseline IV will increase or decrease respectively for that expiry by 1 percentage point. That is:

$$IV_{\text{new}} = \begin{cases} IV_{\text{old}} + 1\% & \text{pool sells 1 SS} \\ IV_{\text{old}} - 1\% & \text{pool buys 1 SS} \end{cases} \quad (2)$$

In this way, the mechanism converges to an IV level which denies arbitrage opportunities as fast as possible. If the order flow which the AMM trades against leads to an IV level which contradicts other venues, arbitrageurs can trade against the AMM to bring IV (and prices) in line.

4.4 Skew Impact

The Black Scholes model does not account for the effect of strike on IV, resulting in the volatility smile which typically arises in options markets. This effect is also known as the *skew* of an option. The AMM incorporates skew into its pricing using a similar mechanism to volatility. We define the skew ratio $SR_{i,j}$ as the ratio of IV of a given strike K_i ($IV_{i,j}$) to the baseline IV for the same expiry IV_j:

$$SR_{i,j} = \frac{IV_{i,j}}{IV_j} \quad (3)$$

Adjusting the skew ratios for each strike accounts for changes in strike volatility. The volatility input for the Black Scholes equation is specific to a strike (K_i), expiry (T_j) combination.

As an example, if the June base IV (IV_{June}) equals 120%, and the June 1900 call (30 delta) has an IV of 130%, the initial skew ratio is 130/120 = 1.0833. If the IV_{June} increases by 1 point, the June 1900 volatility will increase by 1.0833%. By adjusting these ratios up and down based on the supply/demand for a given strike, the AMM can accurately quote options prices using a Black Scholes model.

For every SS bought (sold) by a trader for strike K_i, the mechanism will increase (decrease) the skew ratio $SR_{i,j}$ by a constant c_r that will be initialized as 0.0075. That is:

$$(SR_{i,j})_{\text{new}} = \begin{cases} (SR_{i,j})_{\text{old}} + c_r & \text{Pool sells 1 SS} \\ (SR_{i,j})_{\text{old}} - c_r & \text{Pool buys 1 SS} \end{cases} \quad (4)$$

Continuing the above example, if the initial $SR_{1900,June}$ is equal to 1.0833 and a trader purchases 1 SS worth of 1900 calls or puts, this ratio will increase to 1.0833 + 0.0075 = 1.0908. This ratio will then be multiplied by the post-impact IV_{June} to obtain the IV used in Black Scholes equation. Rearranging (3) gives us the trade's volatility in terms of the skew ratio and baseline volatility:

$$IV_{i,j} = SR_{i,j} \times IV_j. \quad (5)$$

All option listings and expiries can be represented in matrix form, as shown in Appendix A, Definition 6.

4.5 Impact Illustration

As an example, Figure 2 plots the volatility versus the number of standard sizes bought/sold by a trader for an option with $K = 2400$ and $T = 7$. This impact is quadratic, as demonstrated in Appendix A, Lemma 1. This is because the AMM is increasing both the IV and skew of an option, which are then multiplied together to create the parabolic curves.

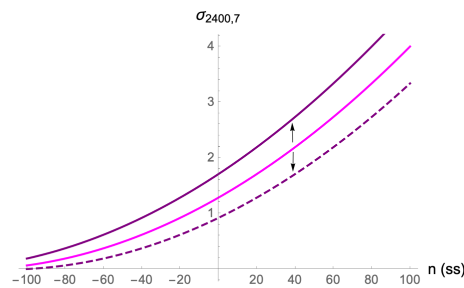

Figure 2: As the AMM sells options the volatility curve is shifted up (magenta to solid purple). Similarly, buying options will lower the volatility curve (magenta to dashed purple).

4.6 Fair Option Price

Now that we have an appropriate IV input to the Black Scholes equation, we calculate the fair price of the option (W). The Black Scholes price of the option for strike K_i and expiry T_j is given by:

$$W = \begin{cases} N(d_1)S_t - N(d_2)K_i e^{-r(T_j - t)} & \text{if call} \\ N(-d_2)K_i e^{-r(T_j - t)} - N(-d_1)S_t & \text{if put} \end{cases} \quad (6)$$

Where

$$d_1 = \frac{1}{IV_{i,j}\sqrt{T_j - t}}[\ln(\frac{S_t}{K_i}) + (r + \frac{(IV_{i,j})^2}{2})(T_j - t)] \quad (7)$$

$$d_2 = d_1 - IV_{i,j}\sqrt{T_j - t} \quad (8)$$

5 Calculating and Managing Pool Risk

We now have an AMM which determines IV, and can accurately price options using Black Scholes. However, without active risk management, LPs are subject to IL and a greatly increased probability of ruin. The net delta and vega positions define the risk profile, allowing for hedging actions to be undertaken when exposure is unacceptably high. This allows the AMM to continue to price options competitively over a long time horizon.

5.1 Delta Risk

Recall that the delta risk defines the exposure of an options position to moves in the underlying asset. Specifically, the delta risk for an AMM is defined as the dollar amount the AMM's position value increases (if long) or decreases (if short) given the price of the underlying asset increases by \$1.

5.1.1 Calculating Delta Risk

1. Given the current price of the asset S_t and the IV for each strike K_i in each expiration T_j, calculate the delta (δ) using the formula:

$$\delta_{i,j} = N(d_1^{(K_i,T_j)}) \tag{9}$$

where d_1 is as defined in (7) and $N(\cdot)$ is the cumulative standard normal distribution. The net position for the (K_i, T_j) pair, denoted by $\rho_{i,j}$ is the number of contracts the AMM is long (negative if the AMM is short).

2. Calculate the delta exposure $E_{i,j}$ for a given strike K_i within a given expiry T_j as:

$$E_{i,j} = \delta_{i,j} \times \rho_{i,j} \tag{10}$$

3. Sum the delta exposures for each strike (n_X total), expiry (n_Y total) combination to calculate net delta (Δ):

$$\Delta = \sum_{i=1}^{n_X} \sum_{j=1}^{n_Y} E_{i,j} \tag{11}$$

To calculate the delta exposure of the AMM in \$ terms ($E_\$$), multiply net delta by the current asset price S_t:

$$E_\$ = \Delta \times S_t \tag{12}$$

5.1.2 Managing Delta Risk

The AMM will hedge a given net delta position Δ by buying, selling or short selling the underlying asset on a spot exchange. This action is triggered by external actors who are incentivized to make valuable calls to the protocol, with minimum delta risk thresholds implemented.

5.2 Vega Risk

Recall that the vega risk defines the exposure of an options position to moves in the IV of the underlying asset. Specifically, the vega risk for an AMM is defined as the dollar amount the AMM's position value increases (if long) or decreases (if short) given the IV of the underlying asset increases by 1% point.

5.2.1 Calculating Vega Risk

Having a metric for vega risk is more involved than for deltas. Simply summing up the vega across expiries ignores the time-dependent risk profile of vega. In practice, the risk from an option where $IV = 120\%$ expiring in 2 years is very different from one expiring in 24 hours. This motivates the following definition for a StandardVega (Ω) which uses a normalization factor (N_j) per expiry T_j:

$$\Omega_{i,j} = \text{vega}_{i,j} \times N_j \tag{13}$$

where $N_j = \sqrt{\frac{30}{T_j - t}}$, $T_j - t$ is the number of days to expiry. The parameter N_j normalizes vega to the same expiry chosen as 30 days away [1], allowing for a valid comparison of vega across expiries. To calculate the overall risk of the AMM in Ω terms, we now repeat an algorithm similar to the delta risk summation:

1. Given the current price of the asset S_t and the IV for each strike K_i in each expiration T_j, calculate the vega:

$$\text{vega}_{i,j} = S_t N'(d_1^{(K_i,T_j)})\sqrt{T_j - t} \tag{14}$$

where $d_1^{(K_i,T_j)}$ is defined in (7).

2. Multiply $vega_{i,j}$ by the net position of the AMM ($\rho_{i,j}$) and the normalization factor N_j to calculate the standard vega exposure of option $E_{i,j}$:

$$E_{i,j} = vega_{i,j} \times \rho_{i,j} \times N_j \tag{15}$$

3. Sum the standard vega exposures for each strike (n_X total), expiry (n_Y total) combination to get

the net standard vega Ψ:

$$\Psi = \sum_{i=1}^{n_X} \sum_{j=1}^{n_Y} E_{i,j} \tag{16}$$

To hedge vega risk we first need to contextualize the AMM's net standard vega exposure by comparing it to the AMM's available liquidity. During a trade, this is calculated as follows:

Define normalized vol as the net standard vega (Ψ) of the AMM multiplied by the post impact IV value of a trade ($IV_{i,j}$) for a strike K_i and an expiry T_j.

$$NormVol = \Psi IV_{i,j} \tag{17}$$

For example, if the post impact $IV_{i,j}$ is 150 vol, and we are net short 500 vega the AMM's NormVol for the trade is equal to -75,000. A 10% increase (15 percentage point increase) in $IV_{i,j}$ would lead to NormVol increasing to

$$165 \times -500 = -82,500$$

and the AMM will approximately incur $7,500 in impermanent losses. This impermanent loss can be realized if the LPs have options sold to them at this new $IV_{i,j}$ level, all other variables held equal. For example if the AMM sells a call with $2 of vega at 150 vol and buys it back for 165 it has realized a loss of:

$$\$2 \times (165 - 150) = \$30$$

We can now define vega utilization at time t as a means to quantify in dollar terms the risk of changes in IV to the AMM. Vega utilization (VU_t) is a 20% change in NormVol as a percentage of the size of the collateral pool (C_{total})

$$VU_t = \frac{0.2 \times \text{NormVol}}{C_{\text{total}}}. \tag{18}$$

This will factor in the change in collateral from the price of options in the proposed trade.

Continuing the above example, if the LP has $800,000 in the collateral pool:

$$VU_t = 15,000/800,000 = 1.875\% \tag{19}$$

5.2.2 Managing Vega Risk

The AMM manages its vega risk by incorporating its vega exposure (defined in vega utilization terms) into the fee charged to a trader. This fee is defined as:

$$f = A \times W + B \times H \times VU + C \times S_t \tag{20}$$

where A, B and C are coefficients that will be initialized closer to launch and eventually governed by the community. The parameter H is equal to 0 if the trade brings the absolute value of the AMM's net standard vega closer to 0, and 1 otherwise. S_t is the price of the underlying asset, and C represents the % fee associated with collateralizing and delta hedging on a spot exchange. Breaking this fee down into its components we have:

- A flat fee based on the option price.
- A dynamic vega risk component.
- A flat fee for exchange costs.

6 Final Price

The final price per option (F) offered to the trader is:

$$F = W \pm f \tag{21}$$

where W is the Black Scholes price of the option obtained in section 4.6, and f is the fee computed in section 5.2.2. We add the fee if the trader buys the option and vice versa. The fee (f) mechanism assesses the current risk of the liquidity pool's position in vega terms (outlined in section 5.2.1) and determines whether the trade a trader is proposing will increase the AMM's risk, and charge a higher fee if so.

For example, consider if the AMM is net short options, and a trader comes in to trade an option where $W = \$125$. If the trader was a buyer, the trade would increase the risk of the AMM, and vice versa if they are a seller. As such, the AMM would charge a larger fee if they were a buyer. This would manifest in an asymmetric market around $125, with the trader paying $140 ($f = \$15$) for the option and selling it at $120 ($f = \$5$), making it relatively more expensive to buy than sell. Combined with IV slippage, this penalizes users who continue buying with a higher fee relative to users who sell. The AMM either achieves more balanced flow (i.e. having both buyers and sellers trade with the protocol) or charges a large enough fee to compensate LPs for the increased risk.

Appendices

A Mathematical Formalism

In the remainder of this paper we present the basic framework which we will eventually use to rigorously prove the features and advantages of the protocol. Future work will expand on what is shown here. For simplicity, we consider the case with only one asset (say, ETH) offered by the liquidity pool since the pricing mechanism treats all assets independently. Suppose Lyra offers n_X total strike prices $K_i, i = 1, \ldots, n_X$ with $K_\alpha < K_\beta$ if $\alpha < \beta$ and n_Y possible expiries T_j, $j = 1, \ldots, n_Y$. Here T_j corresponds to the date of a particular expiry. For instance T_1 could be the 1st of May 2021. Each expiry does not necessarily offer the same number/type of strikes.

Example 1. The May 1 expiry may offer 2000, 2200 strikes on ETH but the May 28 expiry can offer these in addition to a 2500 strike.

Definition 1. We call $K \in \mathbb{R}^{n_X}$ the **strike vector** and $T \in \mathbb{R}^{n_Y}$ the **expiry vector**. Each pair (K_i, T_j) is called a **valid listing** if Lyra offers the strike K_i with expiry T_j, otherwise it is an **invalid listing**.

To each (K_i, T_j) we assign a scalar $R_{i,j} \in \mathbb{R}$ that accounts for the effect strike price has on the implied volatility.

Definition 2. We call $R_{i,j}$ the **skew ratio** for the listing (K_i, T_j).

Suppose that the expiry vector T is ordered in increasing size, i.e. if $T = (T_1, \ldots, T_{n_Y})$ then $T_1 < T_2 < \cdots < T_{n_Y}$. To each expiry we assign an implied volatility b_j which is initialized as the implied volatility for an approximate 50 delta option with the same expiry.

Definition 3. The implied volatility b_j assigned to an expiry T_j is called the **baseline volatility** for that expiry.

Without loss of generality, suppose $n_X \geq n_Y$ (i.e. there are more strikes than expiries). Let R be the $n_X \times n_Y$ matrix whose (i,j) entry is the skew ratio for the listing (K_i, T_j), i.e.

$$R \coloneqq \begin{pmatrix} R_{1,1} & \ldots & R_{1,n_Y} \\ \vdots & \ddots & \vdots \\ R_{n_X,1} & \ldots & R_{n_X,n_Y} \end{pmatrix}.$$

If the listing corresponding to the entry $R_{i,j}$ is invalid, then $R_{i,j} = 0$. Let $\mathcal{V}$ be the $n_Y \times n_Y$ diagonal matrix $\mathcal{V} \coloneqq \mathrm{diag}(b_1, \ldots, b_{n_Y})$.

Definition 4. We call R the **ratio matrix** and $\mathcal{V}$ the **baseline matrix.**

Example 2. Suppose there are 2 possible expiries on ETH: May 1 and May 7 and the former offers strikes 2000 and 2200 while the latter offers these same strikes in addition to 2500. A possible ratio matrix R is

$$R = \begin{pmatrix} 1.05 & 1.06 \\ 1.1 & 1.12 \\ 0 & 1.4 \end{pmatrix}. \tag{22}$$

Note that $R_{3,1} = 0$ since there is no 2500 strike for the May 1 expiry.

Definition 5. The **trading volatility** (referred to as the volatility) $\sigma_{i,j}$ of the listing (K_i, T_j) is given by

$$\sigma_{i,j} = b_j R_{i,j}. \tag{23}$$

The trading volatilities are used to compute the fair price of an option using Black Scholes. We encode all trading volatilities in the volatility matrix.

Definition 6. The **volatility matrix** V corresponding to ratio matrix R and baseline matrix $\mathcal{V}$ is the $n_X \times n_Y$ matrix $V \coloneqq R\mathcal{V}$, i.e.

$$V = \begin{pmatrix} \sigma_{1,1} & \ldots & \sigma_{1,n_Y} \\ \vdots & \ddots & \vdots \\ \sigma_{n_X,1} & \ldots & \sigma_{n_X,n_Y} \end{pmatrix}.$$

Example 3. Suppose that the baseline volatilities for the May 1 and May 7 expiries are 1.2 and 1.4 respectively. The baseline matrix is $\mathrm{diag}(1.2, 1.4)$. Combining this and (22) gives the volatility matrix

$$V = \begin{pmatrix} 1.26 & 1.484 \\ 1.32 & 1.568 \\ 0 & 1.96 \end{pmatrix}. \tag{24}$$

For instance, the trading volatility of a 2200 strike with expiry May 7 is given by the $(2,2)$ entry of (24).

Both the ratio and baseline matricies evolve with each trade conducted. We use superscripts to label the evolution of a quantity, i.e. $R^{(0)}_{\alpha,\beta}$ indicates the initial skew ratio of the (K_α, T_β) listing while $R^{(1)}_{\alpha,\beta}$ denotes the same quantity after one trade and so forth.

Definition 7. A trade $\mathcal{T}_{i,j}(n)$ of n standard sizes of the listing (K_i, T_j) is the map

$$\mathcal{T}_{i,j}(n)\left(R^{(0)}_{\alpha,\beta}, b^{(0)}_\beta\right) = (R^{(1)}_{\alpha,\beta}, b^{(1)}_\beta) \tag{25}$$

where

$$(R^{(1)}_{\alpha,\beta}, b^{(1)}_{\beta}) = \begin{cases} (R^{(0)}_{\alpha,\beta} + c_r n, b^{(0)}_{\beta} + c_b n) & (\alpha,\beta) = (i,j) \\ (R^{(0)}_{\alpha,\beta}, b^{(0)}_{\beta} + c_b n) & \alpha \neq i, \beta = j \\ \left(R^{(0)}_{\alpha,\beta}, b^{(0)}_{\beta}\right) & \beta \neq j. \end{cases} \tag{26}$$

Here c_r is a parameter initially set to 0.0075 and $c_b = 0.01$. The sign of n is positive if the AMM is selling an option and negative otherwise.

From (26), it is easy to see that we have the following.

Lemma 1. *Let $V^{(0)}$ be the original volatility matrix with entries $\sigma^{(0)}_{\alpha,\beta}$. Suppose a trade of n standard sizes of the listing (K_i, T_j) is conducted. This will transform $V^{(0)} \to V^{(1)}$ as follows*

$$\mathcal{T}_{i,j}(n)\left(\sigma^{(0)}_{\alpha,\beta}\right) = \sigma^{(1)}_{\alpha,\beta}(n)$$

where

$$\sigma^{(1)}_{\alpha,\beta}(n; \sigma^{(0)}_{\alpha,\beta}) = \sigma^{(0)}_{\alpha,\beta} + Z^{(0)}_{\alpha,\beta}(n) \tag{27}$$

with

$$Z^{(0)}_{\alpha,\beta}(n) = \begin{cases} \theta^{(0)} n + \phi n^2 & (\alpha,\beta) = (i,j) \\ \xi^{(0)} n & \alpha \neq i, \beta = j \\ 0 & \beta \neq j \end{cases} \tag{28}$$

and $(\theta^{(0)}, \phi, \xi^{(0)}) = (c_b R^{(0)}_{i,j} + c_r b^{(0)}_j, c_b c_r, c_b R^{(0)}_{\alpha,j})$. Note that (27) also takes in parameters $R^{(0)}_{\alpha,\beta}$ and $b^{(0)}_{\beta}$, but this is implicit when writing $\sigma^{(0)}_{\alpha,\beta}$.

Definition 8. We call $\sigma^{(1)}_{\alpha,\beta}(n)$ the **volatility curve** of the listing (K_α, T_β) when a trade of n standard sizes is made of the (K_i, T_j) listing.

Example 4. Continuing the previous examples, suppose Alice purchases 10 standard sizes of the 2200 strike with May 7 expiry. This updates the ratio and baseline matricies to

$$R^{(1)} = \begin{pmatrix} 1.05 & 1.06 \\ 1.1 & \mathbf{1.195} \\ 0 & 1.4 \end{pmatrix} \qquad \mathcal{V}^{(1)} = \begin{pmatrix} 1.2 & 0 \\ 0 & \mathbf{1.5} \end{pmatrix}$$

(updated values bolded for emphasis). Recomputing (24) gives

$$V^{(1)} = \begin{pmatrix} 1.26 & \mathbf{1.59} \\ 1.32 & \mathbf{1.7925} \\ 0 & \mathbf{2.1} \end{pmatrix}. \tag{29}$$

Note that all entries of the second column (i.e. May 7 expiry) have been increased but the second entry more so. This agrees with (28); the volatility of the listing (K_i, T_j) transforms quadratically with n while all other volatilities sharing the same expiry (i.e. in the same column) change linearly. All entries outside of column j are unaffected by the trade.

Remark 1. We can interpret the changes in volatility geometrically as follows. Suppose Alice buys n_1 standard sizes of the (K_i, T_j) listing with original baseline volatility $b^{(0)}$, skew ratio $R^{(0)}$ and trading volatility $\sigma^{(0)} = b^{(0)} R^{(0)}$. Her volatility curve is given by

$$\sigma^{(1)}(n_1; \sigma^{(0)}) = \sigma^{(0)} + \theta^{(0)} n_1 + \phi n_1^2 \tag{30}$$

where $\theta^{(0)}$ and ϕ are as defined earlier. If Bob then buys n_2 standard sizes of the same listing, his volatility curve would be

$$\sigma^{(2)}(n_2; \sigma^{(1)}) = \sigma^{(1)} + \theta^{(1)} n_2 + \phi n_2^2 \tag{31}$$

where $\theta^{(1)} = c_b R^{(1)} + c_r b^{(1)}$. Thus, Alice's trade transforms the volatility curve from (30) to (31). We represent this in Figure (3). The volatility curve is originally the red parabola and the current trading volatility $\sigma^{(0)}$ of the (K_i, T_j) listing is the y intercept of said curve. After h er purchase of n_1 standard sizes, the volatility curve transforms to the blue line and $\sigma^{(0)}$ is updated to $\sigma^{(1)}$ (i.e. the y intercept is shifted up). Note that the red and blue curves are not simple vertical shifts of one another. Similarly, Bob's purchase of n_2 standard sizes transforms the blue to the green curve and also updates the trading volatility to $\sigma^{(2)}$. This result could also have been achieved if Alice originally bought $n_1 + n_2$ standard sizes. That is, the volatility curve is path independent. Similar arguments apply when Alice and Bob trade listings different strikes and/or expiries.

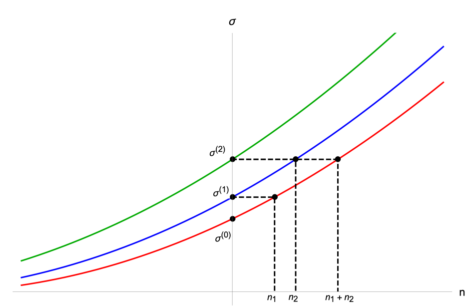

Figure 3: As the AMM sells options the volatility curve is shifted up (red to blue, blue to green). Similarly, buying options will lower the volatility curve.

Example 5. From the previous examples we know that a trade of 10 standard sizes of the $(2200, \text{May } 7)$

listing will impact the volatilities of all listings that share that same expiry. In Figure (4) we show in red/blue the original volatility curves before Alice's purchase for the 2200 and 2500 strikes. After her purchase, these are all shifted up to the solid orange/cyan lines as described by (31). If Alice were to then sell 20 standard sizes, these curves would shift down to their respective dashed lines. This same result could be achieved had she originally sold 10 standard sizes (volatility is path independent). Note that a) are all parabolic curves while b) are all linear. A similar plot to b) can be drawn for the 2000 strike.

B Geometric Connections

Lyra's mechanism appears significantly different to other well known AMMs (Uniswap, Balancer, etc). This is no surprise, since these examples facilitate the exchange of assets (tokens, coins) instead of derivatives of those assets. In the case of Uniswap, the dynamics from each trade are encoded in the curve $xy = k$ where x and y are the quantities of tokens X and Y in the pool. This begs the question for Lyra: what is analogous to Uniswap's curve?

In this section we observe that there is a geometric meaning to Lyra's Black Scholes pricing mechanism encoded in what we call the Black Scholes curves. These are simply related to the volatility curves discussed in the previous section. Due to this, we believe that some results about AMMs known in the literature can be adapted to our circumstances, see e.g. [4]. In future work we hope to rigorously prove such connections exist.

For simplicity, we focus on call options since analagous statements can easily be made for puts. If Alice trades n standard sizes of the (K_i, T_j) listing with original volatility $\sigma_{i,j}^{(0)}$, then the price she pays for one contract is given by the Black Scholes formula

$$C_{i,j} = N(d_1^{(i,j)})S_t - N(d_2^{(i,j)})K_i \exp(-r(T_j - t)). \quad (32)$$

Here

$$d_1^{(i,j)} = \frac{1}{\sigma\sqrt{T_j - t}}\left[\ln\left(\frac{S_t}{K_i}\right) + \left(r + \frac{\sigma^2}{2}\right)(T_j - t)\right], \quad (33)$$

$d_2^{(i,j)} = d_1^{(i,j)} - \sigma\sqrt{T_j - t}$, S_t is the spot price at time t, r is the interest rate (assume fixed, small) and $\sigma := \sigma_{i,j}^{(0)} + Z_{i,j}^{(0)}(n)$ is the volatility for a trade of n standard sizes. We observe that when Alice comes to buy from the AMM, all parameters other than n are fixed. Thus, we can treat (32) as a function only of n and parameterised by the current volatility $\sigma_{i,j}^{(0)}$, i.e. $C_{i,j} = C_{i,j}(n; \sigma_{i,j}^{(0)})$.

Definition 9. We call the curve $C_{i,j}$ the **Black Scholes Call (BSC) curve** for the listing (K_i, T_j). The BSC curve for an invalid listing (K_α, T_β) is trivial, i.e. $C_{\alpha,\beta} = 0$. We encode the BSC curves for all listings in the **call matrix**

$$C := \begin{pmatrix} C_{1,1} & \cdots & C_{1,n_Y} \\ \vdots & \ddots & \vdots \\ C_{n_X,1} & \cdots & C_{n_X,n_Y} \end{pmatrix}. \quad (34)$$

A similar matrix can be constructed for puts and we call this the **put matrix**.

Let $C^{(0)}$ be the original call matrix and $C^{(1)}$ the updated matrix after a trade of n standard sizes worth of contracts of the listing (K_i, T_j). By a similar argument used for the volatilites, we have the following.

Lemma 2. *Suppose n_1 standard sizes of the listing (K_i, T_j) is traded by the AMM and the call matrix $C^{(0)}$ has entries $C_{\alpha,\beta}(n; \sigma_{\alpha,\beta}^{(0)})$. Then the entries of the updated call matrix $C^{(1)}$ are given by $C_{\alpha,\beta}(n; \sigma_{\alpha,\beta}^{(1)}(n_1))$ where*

$$\sigma_{\alpha,\beta}^{(1)}(n_1) = \sigma_{\alpha,\beta}^{(0)} + Z_{\alpha,\beta}^{(0)}(n_1). \quad (35)$$

It is easily shown that like the volatility curves, each trade serves to shift the BSC curve up or down as described in Remark 1. Specifically, we observe that $C_{\alpha,\beta}(0; \sigma_{\alpha,\beta}^{(0)})$ is the present cost of one option for the (K_α, T_β) listing and $C_{\alpha,\beta}(0; \sigma_{\alpha,\beta}^{(1)})$ is the price after one trade.

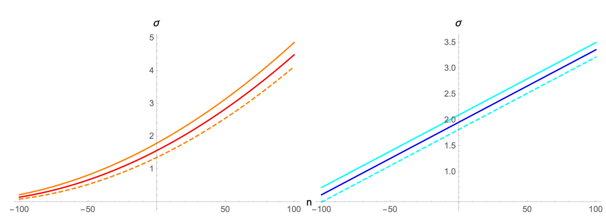

Figure 4: Trading volatilities for the May 7 expiry with strikes (a) 2200 and (b) 2500.

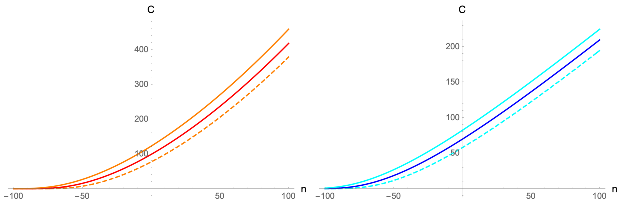

Figure 5: BSC curves for the May 7 expiry with strikes (a) 2200 and (b) 2500.

Example 6. Let us continue the previous example where Alice buys 10 standard sizes of the $(2200, \text{May } 7)$ expiry. The BSC curves for the 2200 and 2500 strikes are shown as solid red/blue curves in Figure (5) a) and b) respectively. After Alice's purchase, all curves are shifted up to the solid orange/cyan lines. Note, however, that the red curve is shifted by more than its blue counterpart. In both cases, when another individual, Bob, comes to trade with the AMM, he will be paying more to buy the same number of contracts as Alice. If Alice were instead to sell 10 standard sizes to the AMM then the original curves would shift down to their dashed counterparts. We plot for 100 standard sizes in both directions to highlight the curvature of each $C(n)$.

We now have a geometric meaning behind Lyra's Black Scholes mechanism; trades transform the volatility curve (Figure (4)) and this translates to an evolution of the Black Scholes curves (Figure (5)). The deformations of these curves seems analogous to the concept of reachable sets described in [2]. Further research is needed to make this connection more rigorous. In the final section of this paper we look at a method for improving the Black Scholes pricing method.

C Pricing Efficiency

We now propose an improvement to the Black Scholes pricing mechanism. Without loss of generality, we assume the AMM is selling options. A similar argument applies if the AMM were to buy instead. Suppose Alice purchases n_f standard sizes of the listing (K_i, T_j) and the Black Scholes curve at the present moment is $C(n)$. Since the number of contracts is a scalar multiple of the standard size (i.e. $N = \chi n_f$), then the total price Alice pays for her N contracts is given by

$$\mathcal{F} = \chi n_f C(n_f). \tag{36}$$

Geometrically, the price Alice pays is proportional (up to a factor of χ) to the area of the dotted rectangle depicted in Figure (6) a). This is a sub optimal pricing since (in the absence of fees and gas) to minimise $\mathcal{F}$, one would ideally trade an infinitesimal amount dn until one has purchased the desired n_f standard sizes. The cost of this optimal strategy is

$$\mathcal{F}_{\text{Optimal}} = \chi \int_0^{n_f} C(n)dn. \tag{37}$$

In Figure (6) a) this is the shaded area under the curve. The difference between the crude (36) and optimal fees (37) becomes significant for large trades and so leads to unnecessarily high fees.

Due to computational constraints, the idealised situation in (37) is not feasible. Instead, a compromise can be reached by performing an upper Riemann integral (a lower sum would lose the AMM money) with η rectangles (Figure (6) b)). Specifically, we have

$$\mathcal{F}_{\text{Approx}} = \chi \sum_{k=1}^{\eta} C(k\Delta n)\Delta n$$

where $\Delta n = \frac{n_f}{\eta}$. Preliminary testing leads us to believe $\eta = 3$ is sufficient for most purposes.

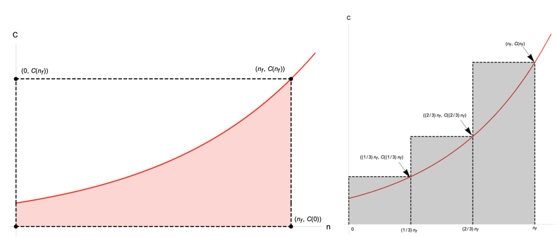

Figure 6: a) The optimal price Alice pays is proportional to the shaded area. b) An upper Riemann sum gives a much better price.

References

[1] *Quantifying the exposure of short-dated option volatility calendar spreads.* 2019, Zurkowski, D.

[2] *Improved Price Oracles: Constant Function Market Makers.* 2020, Angeris, G., Chitra T.

[3] *Toxic Flow: Its Sources and Counter-Strategies* 2020, Zhu, S., Kullander, T

[4] *Growth rate of a Liquidity Provider's wealth in $XY = c$ automated market makers* 2020, Tassy, M., White, D.

[5] *Automated Market Making: Theory and Practice. Ph.D.* 2012, Othman A.