Cloudwall

Digital Asset Volatility Surfaces in Serenity

Youngsuk Lee, Quantitative Research, Cloudwall

July 2023

Abstract

This paper presents the methodology implemented in the Serenity software, by Cloudwall, for the construction of volatility surfaces applicable to cryptocurrency option markets. We provide comprehensive details of our approach, evidencing that the constructed volatility surfaces display appropriate and predictable behavior, thus establishing them as suitable tools for use in option pricing and risk measurement services.

1 Introduction

At Cloudwall, we provide state-of-the-art software to efficiently manage intricate portfolios of digital assets, alongside the derivatives associated with these assets. Recognizing their significance in the market landscape, we extend comprehensive support for plain/vanilla options. These options are deemed critical in financial markets due to their potential to mitigate risks, serve as expressions of perspectives on underlying assets, and act as a source of insight into the anticipations of market participants regarding future movements of these assets. The cryptocurrency options market is a vivid example of this, as it has witnessed an increasing demand for options and other derivative products linked with cryptocurrencies. Readers may refer to the growing option open interests in platforms like Deribit and CME). In response to our users' burgeoning interest in cryptocurrency options and derivatives, we have released real-time cryptocurrency option pricing and Greek calculation services within our Serenity software.

The purpose of this document is to present the methodology for constructing volatility surfaces for option markets using the Stochastic Volatility Inspired (SVI) parameterization. A well-behaved volatility surface is an essential component not only of the aforementioned services but also of risk analytics in general. Recently, we have also created time series of implied volatilities and integrated them into our risk measurement tools, which allow us to calculate risk metrics such as Value-at-Risk (VaR) and conditional VaRs for portfolios of cryptocurrency instruments, including options.

1.1 Acknowledgment

The author would like to thank Kevin Givens, who was a primary contributor to the development of the volatility surface construction process described in this document. The author would also like to thank Marco Marchioro for helping in the revision of the paper.

1.2 Structure of this document

We start with an overview of the system and later explain the details of the computations.

The diagram in figure 1 is a high-level overview of the workflows for constructing the volatility surfaces. The structure of this document closely follows this diagram. We begin with two preprocessing steps: (1) from the futures price quotes, we build the projection yield curve (see, section 2.3),

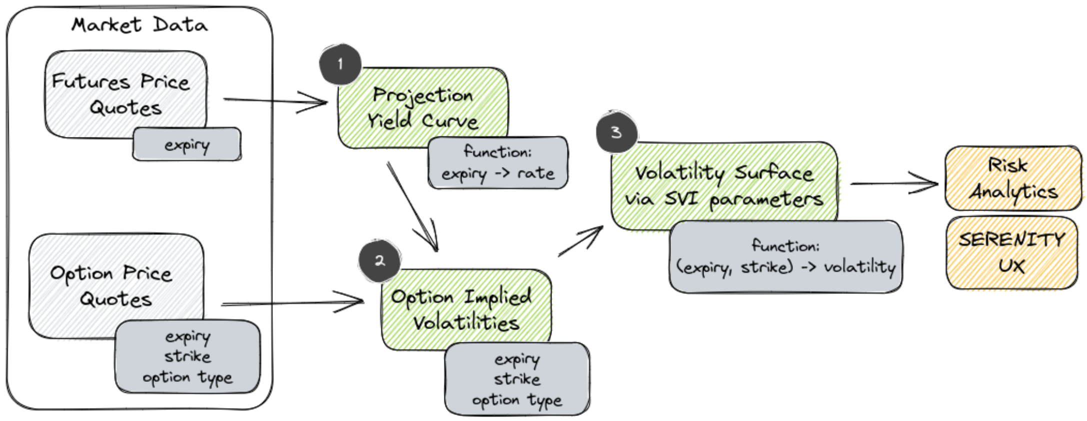

Figure 1: High-level overview of procedures to construct volatility surfaces.

which is used to project the forward price of the underlying asset at any option expiry; (2) We then derive the implied volatilities of the options quoted in the market dataset through the standard Black-Scholes-Merton formula as seen in reference [9] (see also section 2.4). We finally construct a volatility surface, and our approach to this end is to fit an SVI-parameterized volatility smile[1] to the implied volatilities of the options (see section 3).

The SVI parameterization, as shown in reference [1], is a popular choice for the smile structure of the volatility surface because it is flexible enough to fit the smile structure of the volatility surface and it is easy to implement. There are several use cases of the volatility surface, and we will discuss some of its straightforward applications (see section 4) before closing this document (see section 5).

1.3 Conventions

Since the cryptocurrency market operates continously 24/7, we simply measure the length of time intervals between two timestamps in annum, assuming that a year consists of 365.25 days. We do not use daycount conventions, such as those commonly used in the fixed income market. For example, the time-to-expiry of an option contract is calculated as the annum between the as-of-timestamp and the expiry timestamp.

2 Preprocessing for the volatility surface construction

In this section, we describe the market data (see section 2.2) used to prepare the two key ingredients for constructing the volatility surface: the projection yield curve (see section 2.3) and implied volatilities (see section 2.4). To this end, we first review the Black-Scholes-Merton formula (see section 2.1), which introduces key terminologies used throughout this document.

2.1 Review: the Black-Scholes-Merton formula

Consider a European option on the BTC-USD price as an example to illustrate the Black-Scholes-Merton formula (see, e.g., reference [9]). According to the formula, the value V of the option can

[1] A volatility smile refers to a function that maps a strike price to a volatility for a fixed expiry.

be estimated to be

$$V(S,r,p,\sigma;K,T,\phi) = \phi\, e^{-rT} \left[F\mathcal{N}(\phi d_+) - K\mathcal{N}(\phi d_-)\right] , \tag{1}$$

where the market data parameters S, r, p and σ are the spot price, the discount rate, the projection rate, and the volatility of the underlying asset, respectively, and the contractual parameters K, T, and ϕ are the strike price, the time-to-expiry, and the option type (+1 for call and -1 for put), respectively. The forward price F for expiry T of the underlying asset is given by

$$F = S\, e^{pT} , \tag{2}$$

where d_+ and d_- are defined as follows:

$$d_\pm = \frac{\log(F/K) \pm \sigma^2 T/2}{\sigma\sqrt{T}} , \tag{3}$$

and $\mathcal{N}$ is the cumulative density function of the standard normal.

This Black-Scholes-Merton formula is slightly different from the usual *textbook* version[2] in that we use the projection rate p instead of the discount rate r to project the forward price F of the underlying asset. In traditional finance, p corresponds to the difference between the domestic (USD) and foreign (BTC) interest rates if the underlying asset is treated as a currency pair (say, BTC-USD), or the difference between the discount rate (USD) minus the dividend yield of the underlying asset if the underlying cryptocurrency (BTC) is treated as an equity. See, for example, references [11] and [12]. The projection rate is also known as the implied rate since it is *implied* from the relation (2) between the spot price and futures price of the underlying asset, and it is also referred to as the funding cost in the cryptocurrency market.

For the discount rate r, it should be set to the 'risk-free' interest rate of the domestic currency (USD) such as OIS rates in traditional finance. However, there is no clear candidates for this rate in the cryptocurrency market, so we set it to the projection rate for the purpose of constructing the volatility surfaces in the document:

$$r = p , \tag{4}$$

which is consistent with Deribit's Greek calculations.[3]

To be precise, both the forward price F and the projection rate p depend on the time-to-expiry T, and should be written as F_T and p_T, respectively. For example, (2) can be written as

$$F_T = S \exp\left(p_T \cdot T\right) , \tag{5}$$

to emphasize the dependence on T. However, we will suppress the explicit dependence on T for notational simplicity unless it is necessary to do otherwise.

2.2 Handling market data

As shown in figure 1, the procedure for constructing the volatility surface requires price quotes for futures contracts and European options as key inputs. Although there are several exchanges that offer futures and options contracts on cryptocurrencies, we use data from Deribit because it is a leading exchange for cryptocurrency options (see, for example, this article) and offers option contracts over a wide range of expiries and strikes for both BTC and ETH against USD.[4]

[2] See, for example, https://en.wikipedia.org/wiki/Black_model.

[3] See https://github.com/cryptarbitrage-code/deribit-position-greeks.

[4] The Deribit price quotes used in this document were sourced through a third-party market data provider, Kaiko (https://www.kaiko.com/).

#	expiry timestamp	futures	options			description
			count	min strike	max strike	
0	2023-05-02 08:00		38	26,000	32,500	1 daily
1	2023-05-03 08:00		44	25,500	33,000	2 daily
2	2023-05-04 08:00		38	25,000	32,000	3 daily
3	2023-05-05 08:00	o	92	16,000	51,000	1 weekly
4	2023-05-12 08:00	o	86	15,000	48,000	2 weekly
5	2023-05-19 08:00		88	15,000	49,000	3 weekly
6	2023-05-26 08:00	o	64	10,000	55,000	1 monthly
7	2023-06-30 08:00	o	74	5,000	100,000	2 monthly, June
8	2023-07-28 08:00		64	10,000	65,000	3 monthly
9	2023-09-29 08:00	o	78	5,000	80,000	September
10	2023-12-29 08:00	o	82	5,000	100,000	December
11	2024-03-29 08:00	o	72	10,000	120,000	March

Table 1: Futures and Option Contracts on BTC in Deribit as of 2023-05-01 20:00 (UTC). As a reference, the BTC-USD spot price was around 27,850. Futures contracts are available for the expiries indicated using circles.

The market structure

To illustrate the market structure of the futures and option contracts in Deribit, in table 1 we summarize the available price quotes on Deribit as of 2023-05-01 20:00 (UTC). While both futures and option contracts cover expiries up to about one year ahead, there are more option expiries than futures expiries. For options, there is a wide range of strikes available for each expiry timestamp, although options with strikes near the current spot price are expected to be more liquid than those far from it. While the range of strike prices is generally wider for longer-dated options, the smallest and largest strikes are not necessarily monotonic with respect to the time-to-expiry. For example, the range of strikes for June-expirying options spans is between 5,000 and 100,000, which is wider than that for July-expirying options (spanning 10,000 to 65,000). This is not surprising considering that June expirying options were first introduced about one-year ago. They have longer history than any other options on that day (see also figure 2), and thus June-expirying options have been traded over a wider range of spot prices than others.

As time passes, old contracts expire and new contracts are introduced to the market according to the contract introduction policy outlined by Deribit. Figure 2 illustrates this process, showing a downward-sloping pattern of the time-to-expiry over time.

Market data for price quotes

For each underlying crypto pair (i.e. for BTC-USD and ETH-USD), we use the following price quotes as inputs to the volatility surface construction procedure:

- futures: futures prices over expiry timestamps. The i-th futures price quote is represented as

$$(T_i, F_i) \tag{6}$$

  where T_i is the time-to-expiry and F_i is the futures price quoted in USD for the i-th contract.

- options: option values over expiry timestamps and strikes for both call and put options. The

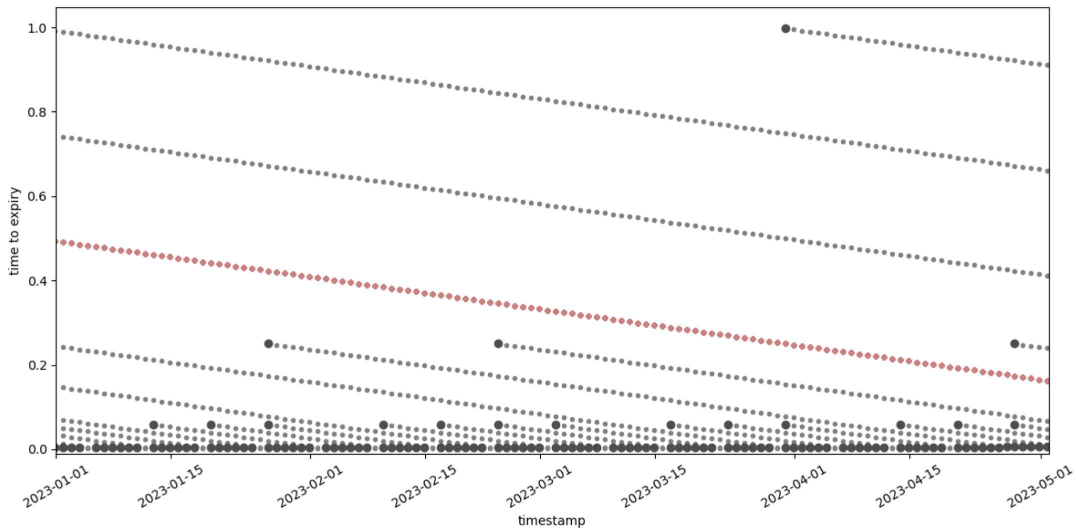

Figure 2: Time-to-expiries of BTC-USD options between 2023-01-01 and 2023-05-01, sampled daily, are plotted using grey dots for each timestamp in the interval. Newly introduced option expires are identified using black circles, and the June expiry is identified using red dots. Note that options with the June expiry have the longest historty as of 2023-05-01 20:00.

i-th option price quote is represented as

$$(T_i, K_i, \phi_i, V_i) \tag{7}$$

where T_i is the time-to-expiry, K_i is the strike price, ϕ_i is the option type (+1 for call or -1 for put), and V_i is the option value quoted in USD for the i-th option contract.

The data used in this document were sourced from our own database of Deribit price quotes, which were collected through Kaiko's API.

Note: Since Deribit's option *mark price* is equal to the undiscounted USD option value over the corresponding forward price.[5], the option value V_i is equal to the mark price multiplied by the spot price S, since we discount with the projection rate $r = p$, as discussion in section 2.1.

2.3 The projection rate term structure

In order to apply the Black-Scholes-Merton formula to options, we need to project the forward price F for the option's expiry timestamp. For options with expiry timestamps that match those of the futures contracts, the forward price F can be simply set equal to the futures price F_i quoted at the expiry timestamp T_i. However, for option contracts with no matching expiries, we need to construct a projection rate curve p_T so that the forward price F_T for any expiry timestamp T can be calculated using (5)

To derive the projection rate p_i for each of the futures price quotes (T_i, F_i), we solve equation (2) for $i = 1, \cdots, N$. We then construct p_T as a curve that interpolates/extrapolates the collection of projection rates (T_i, p_i) for $i = 1, \cdots, N$. Specifically, we use the flat-forward interpolation/extrapolation

[5]See `https://insights.deribit.com/dev-hub/deribit-position-greeks-python/` and links therein.

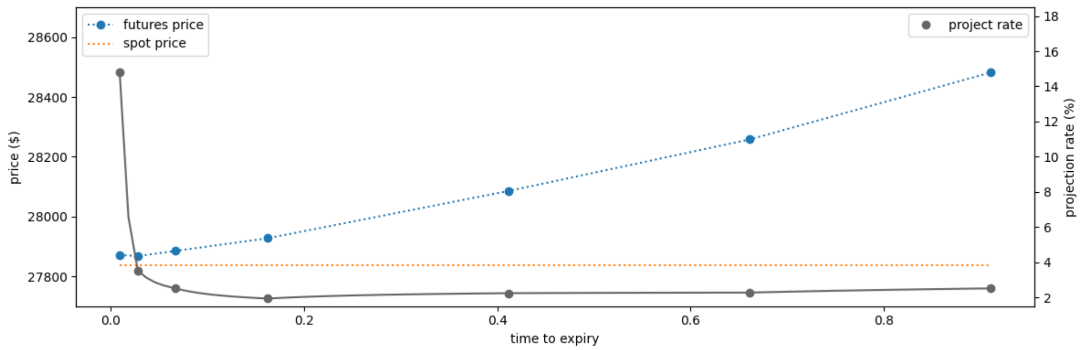

Figure 3: Projection yield curve (solid black line) is constructed using the BTC-USD futures price quotes (blue dots) as of 2023-05-01 20:00, which are summarized in table 1.

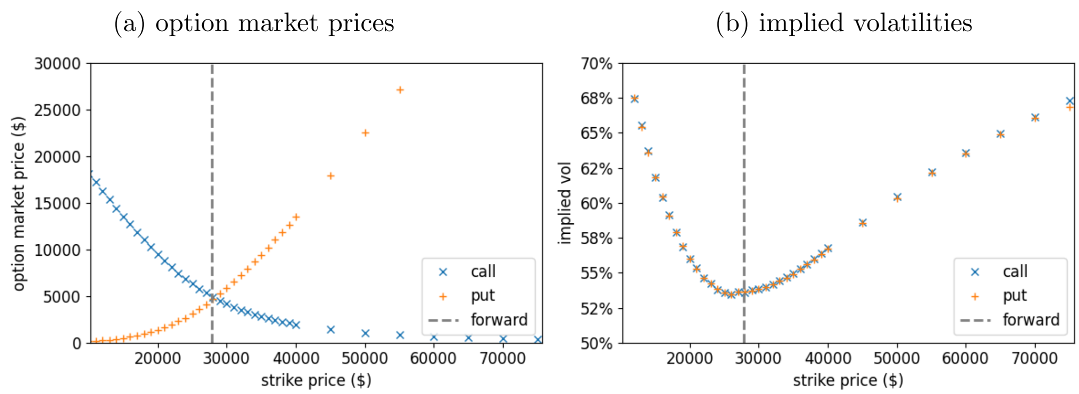

Figure 4: (a) option market prices (denoted by V^{mkt} in (8)) for BTC-USD call (blue crosses) and put options (orange dots) over a range of strike prices as of 2023-05-01 20:00 for expiry 2023-12-29 08:00. (b) Implied volatilities (denoted by σ^{mkt} in (8)) for the same options as in (a).

method, which is a standard method used in the interest rate market. For example, figure 3 illustrates the projection yield curve together with the futures price quotes that are used as inputs to the curve construction.

For further discussions on this and related topics, we refer readers to our companion paper in reference [10].

2.4 Computation of implied volatilities

We detail here the computations of implied volatilities. Consider an option contract with strike K, expiry T, and type ϕ. First, derive the forward price F using (2) together with the projection curve described in section 2.3. Then, use the Black-Scholes-Merton formula (1) to solve for the implied volatility of the option, denoted by σ^{mkt}, by matching the theoretical value of the option with its market price V^{mkt}:

$$V(S, p, \sigma^{\mathrm{mkt}}; K, T, \phi) = V^{\mathrm{mkt}}(K, T, \phi)\,. \tag{8}$$

We use the standard Newton-Ralphson method to solve this equation.

Showing all dependent variables explicitly, the implied volatility can be expressed as a function of contractual parameters K, T, and ϕ, with market data S, p as input parameters:

$$\sigma^{\text{mkt}}(K, T, \phi; S, p)\,. \tag{9}$$

For example, figure 4 shows the market prices of BTC-USD call and put options over a range of strike prices for a given expiry timestamp and the corresponding implied volatilities. The forward price is located in both figures to indicate the strike price where the put and call prices are expected to be equal due to put-call parity.

Recall that, for the same strike K and expiry T, the implied volatilities of the call and put options are *supposed* to be identical because of put-call parity. However, with noise in market data (e.g. differences in time for collecting option prices), there can be differences between the implied volatilities of the call and put options (e.g. small differences can be noticed in Figure 4 (b)). Nevertheless, in section 3, we create a single volatility surface for both types of options, which is a common practice in the industry, enforcing put-call parity.

3 The volatility surface construction

The purpose of this section is to describe how we construct a volatility surface, which is a 2-dimensional continuous function that maps strike price K and the time-to-expiry T to implied volatility σ:

$$K, T \mapsto \sigma(K, T)\,. \tag{10}$$

At a high level, we fit an SVI-parameterized *smile* curve to the implied volatilities of each expiry T_i in the dataset. The SVI parameterization is a popular choice in the industry due to its simplicity and flexibility, and has also been used in studying option volatilities for the cryptocurrency markets as well (see, e.g., references [3, 4]). The SVI smiles are then interpolated to create a continuous volatility surface for any expiry T.

In the rest of this section, we first briefly review the SVI parameterization in section 3.1, which is a prerequisite to describe how we fit the SVI smiles to the implied volatilities of each expiry in section 3.2. We then collect the fitted SVI smiles over all expiries to form a volatility surface, which we example in section 3.3.

3.1 The stochastic volatility inspired (SVI) parameterization

We introduce the log-moneyness k (also known as log-strike) and the implied variance w for expiry T, which are convient variables for discussing volatility surfaces:[6].

$$k \;:=\; \ln\left(K/F_T\right)\,, \tag{11}$$

$$w(k,T) \;:=\; T\sigma^2\left(F_T e^k, T\right)\,, \tag{12}$$

where the := symbol is used to introduce a definition. The SVI parameterization, again see reference [1] for more details, describes the implied variance w as a function of the log-moneyness k for a given expiry T. Specifically it takes the form:

$$w(k,T) = a + b\left\{\rho(k-m) + \sqrt{(k-m)^2 + s^2}\right\}\,, \tag{13}$$

[6] For example, when measured in the foreign currency, the undiscounted call option value can be written in the two variables:

$$N(d_1) - e^k N(d_2),$$

where $d_1 = -k/\sqrt{w} + \sqrt{w}/2$ and $d_2 = d_1 - \sqrt{w}$. See also reference [7] for more discussions.

where the coefficients $\{a, b, \rho, m, s\}$[7] are parameters that depend on the expiry T and satisfy the following constraints:[8]

$$a \in \mathbb{R}, \quad b \geq 0, \quad |\rho| < 1, \quad m \in \mathbb{R}, \quad s > 0, \quad a + b\,\sigma\sqrt{1-\rho^2} \geq 0.\,. \tag{14}$$

In terms of parameters names, a is the level, b is the slope at-the-money (ATM), ρ is the skew, m is the moneyness at which the smile is centered, and s is the smile curvature.

As discussed in reference [2], the SVI parameterization is free of calendar spread arbitrage if the total variance $w(k, T)$ is a non-decreasing function of T for each log-moneyness k, or equivalently,

$$\partial_T w \geq 0\,. \tag{15}$$

To be free of butterfly arbitrage (see reference [2]), the probability density function $f(k)$ implied from the SVI parameterization must be non-negative (absence of butterfly arbitrage), which implies that call prices are decreasing and convex in the strike price. The density can be computed from the undiscounted call option price function $C(K)$ (see reference [5]):

$$f(k) = \left.\frac{\partial^2 C(K)}{\partial K^2}\right|_{K=F_T e^k} \geq 0\,. \tag{16}$$

An explicit formula for $f(k)$ can be found in reference [2], and it can be illustrated using probability historigrams as shown in Figure 6.

In literature, there are sophisticated approaches to ensure the SVI parametrization is free of arbitrage. For example, the paper in reference [2] proposed a class of artibrage-free SVI volatility surface parametrizations, similarly reference [6] proposed methods to resolve the arbitrage problem using quadratic programming algorithms or imposing sufficient conditions that guarantee an arbitrage-free SVI. However, in this document, we take a simple approach of fitting the SVI parameterization to the implied volatilities of each expiry using a standard least-squares method with minimal constraits. As demonstrated in section 3.2, this simple approach provides a good fit to the market data, and the resulting SVI smiles reflect the arbitrage-free nature of the market data.

3.2 Parameter calibration

For a given expiry T, we fit the SVI parameterization to the implied volatilities of the option contracts that share the same expiry. This fitting is achieved by minimizing the sum of weighted squared errors (SSE) between the implied total variances and the SVI parameterization over the option contracts:

$$\mathcal{L}(a, b, \rho, m, s) = \sum_{i=1}^{N} \eta_i \left(w_i^{\text{mkt}} - w(k_i)\right)^2. \tag{17}$$

Here, k_i is the log-moneyness, $w_i^{\text{mkt}} := T\left[\sigma_i^{\text{mkt}}\right]^2$ is the implied total variance, and η_i is the weight of the i-th option contract. During the minimization process, we impose the constraints described by the inequalities (14) together with the following additional constraint to avoid butterfly arbitrage:

$$b\,(1.0 + |\rho|) \leq 4.0\,, \tag{18}$$

which a necessary condition to avoid butterfly arbitrage, as suggested by reference [1].

[7] In the original SVI parameterization, the parameter s is denoted by σ. However, we use s to avoid confusion with the implied volatility σ.

[8] The last constrait is to ensure the variance remains non-negative.

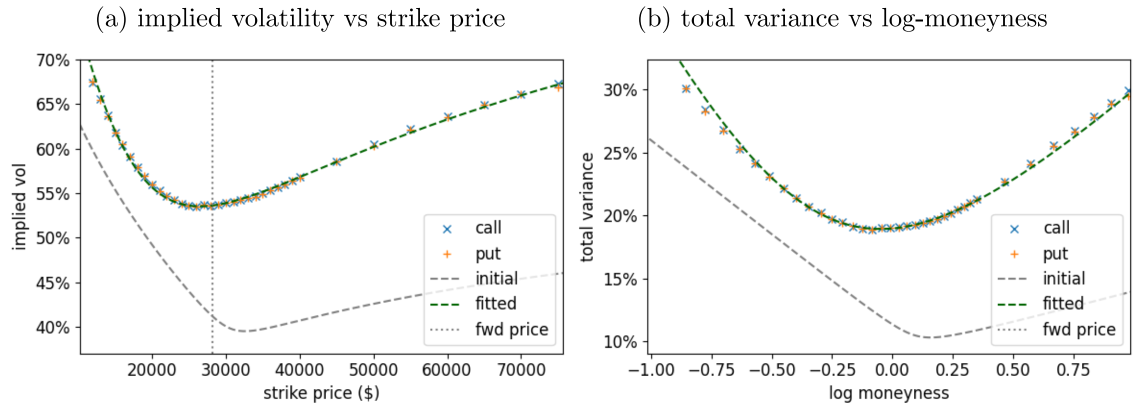

Figure 5: (a) SVI-parameterized volatility smile (dashed green line) fitted to the implied volatilities (from both calls and puts) of BTC-USD call options over a range of strike prices for expiry 2023-12-29 08:00. The initial starting smile is shown in dashed black line. (b) The same plot but in the total variance space over the log-moneyness.

After considering various options of the weights η_i, initial parameter guesses and optimization algorithms, we have chosen the following approaches for our current production implementation.

- The *vega*-weighting scheme is used where the weight of the i-th option contract is proportional to the Black-Scholes-Merton vega and is given by

$$\eta_i = \mathcal{N}'(d_1)\,. \tag{19}$$

  This scheme puts more weight on options near the money than on far in-the-money or out-of-the-money options, reflecting the typical liquidity characteristics of option contracts in the market. For further motivations on this choice, see [13].

- We apply a trust region algorithm (see reference [14]) to solve the constrained optimization problem of minimizing the objective function (17) subject to the constraints (14) and (18).

- For the initial parameter guesses, we follow the suggestions in reference [6].

When applied to the implied volatilities shown in figure 4 (b), our approach produces the SVI parameterized volatility smile shown in figure 5 (a) along with the initial starting smile. Figure 5 (b) shows the same plot but in the total variance space over the log-moneyness, which is the coordinate system where the SVI parameterization is defined and where the error terms in the objective function (17) are measured. Upon visual inspection of the plots, it appears that our approach provides a good fit to the market data. The fitted SVI smile is free of butterfly arbitrage since the implied probabilities are positive as shown in figure 6. Specifically, we show the implied probability density function (16) is bucketed into equal-sized bins over the strike prices in figure (a), while the same function bucketed over the log-moneyness values in figure (b).

3.3 The volatility surface

By collecting the SVI smiles over all expiries, we can illustrate a volatilty surface as in figure 7. Broadly speaking, the smiles are U-shaped and centered around the forward price. When plotted

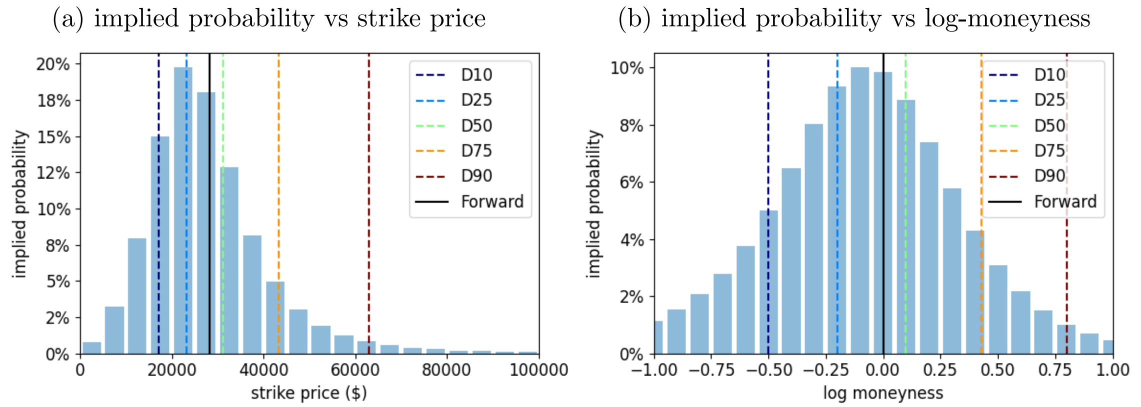

Figure 6: (a) Probability histogram implied from the SVI-parameterized volatility smile shown in figure 5. (b) The same plot but over the log-moneyness.

in the implied volatility space, see figure (a), one can notice that the shape is more pronouced for shorter expiries than longer ones. When plotted in the total variance space, see figure (b), one can observe that the level of the total variance increases with the time-to-expiry for each moneyness, indicating that the surface is free of calendar spread arbitrage, i.e. satisfying equation (15).

Figure 7 also shows that the typical range of strike prices where the option market is active (thus quotes represented using 'dots' are available) is wider (narrower) for longer (shorter) expiries. This is expected since the uncertainty of the underlying asset price is higher (lower) for longer (shorter) expiries. To incorporate this uncertainty in the comparison of volatilities from different expiries, one would stretch the strike axis for shorter expiries, or equivalently, squeeze it for longer expiries. A standard ways to achieve this is to express the strike axis in terms of delta-moneyness as demonstrated by the lines along constant delta-moneyness values in figure (b). Specifically, the delta-moneyness $\Delta(K)$ for strike price K is defined as the negative of the Black-Scholes-Merton delta of the put option with strike K (see, for example, reference [15]):

$$\Delta(K) = -\mathcal{N}(-d_1(K)) \,. \tag{20}$$

Since the delta-moneyness function $\Delta(K)$ maps $(0, \infty)$ to $(0, 1)$ monotonically, it can be used to standardize the strike axis for all expiries to the unit interval as illustrated in figure 8. It is visually clear that the smile shapes are more consistent across expiries when plotted in the delta-moneyness space than, say, in the log-moneyness space, allowing use to compare the volatilities across different expires on a more equal fotting. For example, the implied volatilities across expires, see figure (a), have compariable magnitudes for the same delta-moneyness and the monotonic pattern in the total variance space, see figure (b), is easier to observe.

4 Use cases

Once a volatility surface is constructed, it can be used for various purposes. In this section, we illustrate two simple use cases: summarizing a volatility surface into a table consisting of standard volatility metrics that are frequently quoted in the market (see section 4.1), and constructing volatility time series data which can be used as inputs to various applications, including risk measurements (see section 4.2)

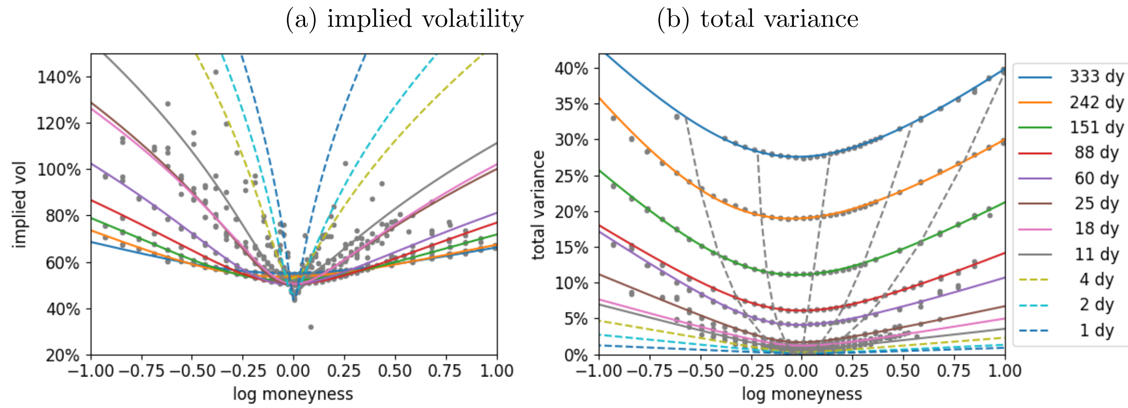

Figure 7: Volatility smiles in terms of the log-moneyness across expiries as of 2023-05-01 20:00. The implied volatities from quoted options are represented using dots, while the fitted SVI-parametrized smile curves are shown using solid or dashed lines depending on expiries. They are shown in two different spaces: (a) in the implied volatility space and (b) in the total variance space. In (b), the five grey 'cones' (dashed lines) represent the log-moneyness values corresponding to delta-moneyness 10%, 25%, 50%, 75%, and 90%, respectively. These cones highlight the range of log-moneyness values where the option market is active with available quotes.

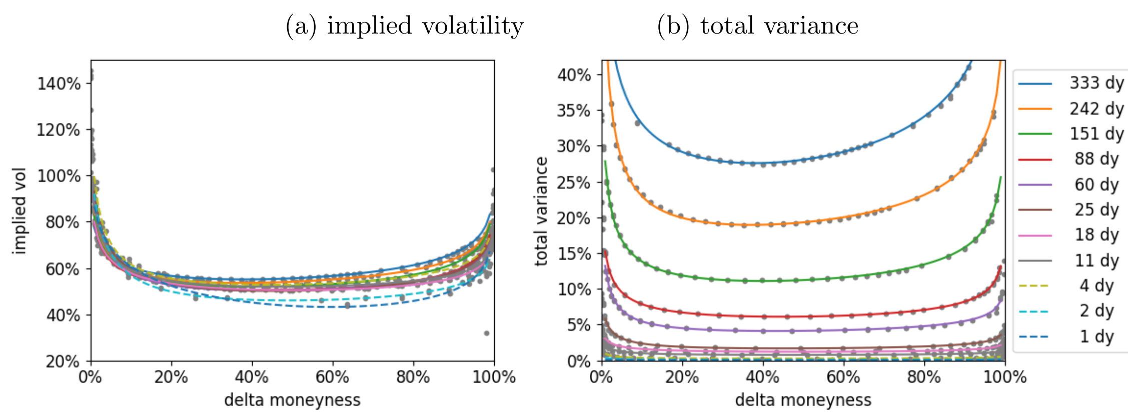

Figure 8: Volatility smiles in terms of the delta-moneyness across time-to-expiries. Otherwise, the same as figure 7.

	expiry datetime	10P	25P	ATM	25C	10C	25RR	25BF	10RR	10BF
0	2023-05-02 08:00	61%	50%	44%	44%	50%	-5.2%	-3.5%	-10.3%	-12.0%
1	2023-05-03 08:00	57%	48%	46%	48%	52%	-0.4%	-2.1%	-4.8%	-8.5%
2	2023-05-05 08:00	62%	54%	53%	55%	60%	1.1%	-2.0%	-1.4%	-8.1%
3	2023-05-12 08:00	59%	53%	52%	54%	57%	0.2%	-1.6%	-2.1%	-6.2%
4	2023-05-19 08:00	58%	52%	50%	52%	56%	-0.3%	-1.7%	-2.2%	-6.3%
5	2023-05-26 08:00	59%	52%	50%	53%	58%	0.9%	-2.2%	-1.4%	-8.6%
6	2023-06-30 08:00	58%	51%	50%	53%	58%	1.6%	-1.9%	0.2%	-7.6%
7	2023-07-28 08:00	57%	52%	51%	53%	58%	1.3%	-1.9%	1.1%	-7.4%
8	2023-09-29 08:00	58%	53%	52%	56%	62%	2.8%	-2.0%	3.9%	-7.7%
9	2023-12-29 08:00	59%	54%	54%	58%	64%	3.6%	-1.9%	4.8%	-7.4%
10	2024-03-29 08:00	60%	56%	55%	60%	66%	3.8%	-2.2%	6.0%	-7.8%

Table 2: Volatility table summarizing the volatility surface in figure 8.

4.1 Volatility metrics and smiles

The volatility metrics that are commonly used in the FX market include the at-the-money volatility, the risk reversal, and the butterfly. They are defined in terms of option deltas (see, e.g., reference [15]):

- $\sigma_{x\mathrm{P}}$, known as x-delta put, is defined as $\sigma(K)$ such that $\Delta(K) = x/100$ for $0 < x < 50$.
- σ_{ATM} is the at-the-money volatility, defined as $\sigma(K)$ where $\Delta(K) = .5$.
- $\sigma_{x\mathrm{C}}$, known as x-delta call, is defined as $\sigma(K)$ such that $\Delta(K) = 1 - x/100$ for $0 < x < 50$.

Then, the risk reversal (RR) and butterfly (BF) are defined as

$$\sigma_{x\mathrm{RR}} := \sigma_{x\mathrm{C}} - \sigma_{x\mathrm{P}} \quad \text{and} \quad \sigma_{x\mathrm{BF}} := \sigma_{\mathrm{ATM}} - \frac{1}{2}(\sigma_{x\mathrm{C}} + \sigma_{x\mathrm{P}})\,. \tag{21}$$

Using the above definitions provided above, we can summarize the volatility surface shown in figure 8 for two typical values of x in the FX market, namely, $x = 25$ and 10. The resulting summary is presented in table 2.

4.2 Time series construction

The volatility surface described in section 3 is specified at the expiry dates where the market quotes are available. As time goes by, these expiry dates become shorter, eventually reaching zero at the expiry date. If time series were to be constructed by simply collecting volatilities at a fixed expiry over time, it would not be an appropriate for meaningful time series analysis. To avoid this issue, we construct time series data at constant time-to-expiry points (e.g. 1M, 3M, etc) by interpolating the volatility surface in the time-to-expiry direction for a fixed moneyness. Specifically, we use linear interpolation in the toal variance space to preserve the calendar spread arbitrage-free nature of the surface.

Figure 9 shows a sample of daily volatilty time series for a 3M time-to-expiry and 50% delta-moneyness (ATM), interpolated from the fitted BTC-USD implied volatility surfaces over the same period as in figure 2. In Serenity, we actually have hourly time series data for various moneyness and time-to-expiry points for both BTC and ETH. In fact, some of these time series are used in our risk measurement calculations such as VaRs and conditional VaRs.

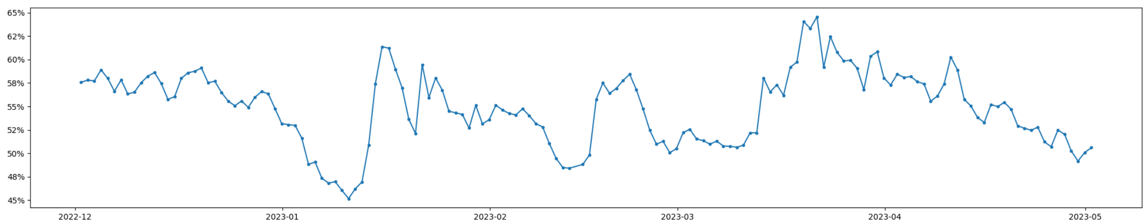

Figure 9: Daily BTC-USD implied volatility time series for 3M time-to-expiry and ATM delta-moneyness over the same period as in figure 2.

5 Conclusions

This document details the process by which we construct volatility surfaces applicable to the cryptocurrency options market, utilizing the SVI parametrization. The resulting surfaces exhibit appropriate and predictable behavior as evidenced by the illustrative examples presented herein, thus demonstrating their efficacy. These surfaces have practical applications in the Serenity software by Cloudwall, specifically as inputs for real-time computations related to option valuation and sensitivities.

Furthermore, we have established a series of volatility time series data encompassing various moneyness levels and time-to-expiry points. Some of these time series are actually used as inputs for Serenity computations of risk measures.

We anticipate that our volatility surface data, along with the derived time series, will serve as valuable resources for comprehending the risk associated with the cryptocurrency options market. Furthermore, these tools are expected to provide us with deeper insights as we integrate implied volatility shocks into the scenario analysis and stress testing methodologies that are available in Serenity.

The methodology detailed in this document, in conjunction with the entirety of the accompanying documentation, underscores the sophistication of the quantitative mechanisms driving the Serenity software. Additionally, we would like to highlight how Serenity endows its users with robust tools for effectively navigating the derivative markets in the digital-asset world.

References

[1] J. Gatheral, *A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives*, http://faculty.baruch.cuny.edu/jgatheral/madrid2004.pdf, May 2004 2, 7, 8

[2] J. Gatheral, A. Jacquier *Arbitrage-free SVI volatility surfaces*, Quantitative Finance, https://www.tandfonline.com/doi/abs/10.1080/14697688.2013.819986?journalCode=rquf20, 2014 8

[3] M. Echenim, E. Gobet, A.C. Maurice *Unbiasing and robustifying implied volatility calibration in a cryptocurrency market with large bid-ask spreads and missing quotes*, Kaiko Research. https://hal.science/hal-03715921v1/file/main.pdf, 2022 7

[4] J. Matic, N. Packham, W. Härdle *Hedging cryptocurrency options*, https://link.springer.com/article/10.1007/s11147-023-09194-6, 2023 7

[5] D.T. Breeden, R.H. Litzenberger, *Prices of state-contingent claims implicit in option prices*, Journal of Business, 621-651. https://www.jstor.org/stable/2352503, 1978 8

[6] T. Ferhati, *Robust Calibration For SVI Model Arbitrage Free*, https://hal.science/hal-02490029/file/Robust_Calibration_For_SVI_Model_Arbitrage_Free_v4.pdf, 2020 8, 9

[7] J. Gatheral, *The Volatility Surface: A Practitioner's Guide*, Wiley Finance. ISBN 978-0-470-01877-6, 2006 7

[8] P. Hagan, D. Kumar, A. Kesniewski, D. Woodward, *Managing Smile Risk*, Wilmott. Vol. 1. pp. 84–108, January 2002

[9] F. Black, M. Scholes, *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy. 81 (3): 637–654, 1973 2

[10] K. Givens, *Risk Neutral Discounting for Crypto Options in Serenity*, Research, Cloudwall, https://cloudwall.tech/research, 2023 6

[11] A. Castagna, F. Mercurio *Consistent Pricing of FX Options*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=873788 3

[12] J. Hull, *Options, Futures, and Other Derivatives*, 9th Edition, 2014 3

[13] L. Nagy, M. Ormos, *Volatility Surface Calibration to Illiquid Options*, The Journal of Derivatives. Spring 2019, 26 (3) 87 - 96. https://doi.org/10.3905/jod.2019.26.3.087, 2019 9

[14] A. Conn, N.I. Gould, P.L. Toint, Trust region methods. 2000. Siam. pp. 19. 9

[15] D. Reiswich, U. Wystup, *FX Volatility Smile Construction*, CPQF Working Paper Series, No. 20, http://hdl.handle.net/10419/40186, 2009 10, 12

