

Risk Measurement of Digital Asset Options

Youngsuk Lee, Quantitative Research, Cloudwall

August 2023

Abstract

This document presents a comprehensive methodology for generating profit and loss (P&L) scenarios within the Serenity risk service, focusing on options in digital asset markets. We also showcase how this new service can be utilized as a tool for risk management by computing a number of distribution measures such as Value-at-Risk.

1 Introduction

Since its release in June 2023, the Serenity software has incorporated an advanced risk service that enables the calculation of various risk measures, including value-at-risk (VaR) and conditional value-at-risk (CVaR, also known as expected shortfall), for portfolios of digital assets and derivatives (perpetuals, futures and options).

This document presents a comprehensive methodology for generating the profit and loss (P&L) scenarios within the risk measure service, with a particular emphasis on options. The methodology encompasses historical simulations of the spot prices and the volatility surfaces, and it highlights the importance of conducting full revaluations to accurately capture the nonlinear effects associated with option positions. Furthermore, we demonstrate how this risk service can be effectively utilized as a tool for risk management and portfolio optimization.

Related publications This publication is part of a series on digital asset derivatives, which includes previous papers such as reference [3], on interest rate curves, and reference [4] on construction of volatility surfaces. Familiarity with the latter paper is assumed when discussing the method for volatility surface simulations.

Acknowledgment The author would also like to thank Marco Marchioro for insightful discussions provided during the development stages of the risk service and for helping in the revision of the paper.

2 Historical Simulations

First, we introduce the notation used to describe the historical simulation approach in this document. For a given scalar variable x such as the spot price, we define:

- $\bar{x}$, the reference value of x at the current *as-of* time;
- $x_0, x_1, x_2, \dots, x_N$, the historical data points for x , marked at regular time intervals, covering a look-back period up to the current time;
- $\tilde{x}_n$, the n -th historical scenario for x , generated by perturbing $\bar{x}$ using x_{n-1} and x_n , where $n = 1, 2, \dots, N$;

- Δt , the time interval between two consecutive historical data points x_{n-1} and x_n ;
- h , the risk horizon used to scale historical scenarios to represent returns over $h, \Delta t$

In Serenity, we generate $N = 24 \times 365 = 8,760$ hourly scenarios by default, using hourly data with a time interval of $\Delta t = 1$ hour over a one-year look-back period. We set $h = 24$ to represent scenarios equivalent to returns over a one-day period.

Note: Throughout this document, we assume this default setting unless otherwise stated.

2.1 The Spot Price Simulations

To simulate the n -th historical scenario $\tilde{S}_n$ for spot price S , we perturb the reference value $\bar{S}$ using the historical data points S_{n-1} and S_n as follows:

$$\tilde{S}_n = \bar{S} \left(\frac{S_n}{S_{n-1}} \right)^{\sqrt{h}}, \quad n = 1, 2, \dots, N. \quad (1)$$

Alternatively, we can express this as:

$$\ln \tilde{S}_n = \ln \bar{S} + \sqrt{h} \cdot (\ln S_n - \ln S_{n-1}), \quad (2)$$

which means that the n -th simulated change is calculated as the difference between the historical values in the log space, scaled by $\sqrt{h}$.

The square-root-of-time rule The scaling of the historical return by $\sqrt{h}$ is consistent with a fundamental principle of stochastic processes and mathematical finance, which is known as the *square-root-of-time* rule. This rule essentially states that the standard deviation of returns scales with the square root of time. Indeed, consider a random walk model where the changes in the spot price S are independent and identically distributed. If ΔS_t is the change in S in time period t , then the variance of the cumulative change over h periods is given by $h \cdot \text{Var}(\Delta S_t)$. The standard deviation, which is the square root of the variance, thus scales as $\sqrt{h}$. The implication of this scaling rule is that the potential dispersion of the spot price increases with time, but at a decreasing rate. In summary, the square-root-of-time scaling in the historical simulation method is crucial for accurately capturing the time-dependence of price volatility. This feature is consistent with the empirical observation that price fluctuations tend to grow with the square root of time, and it is fundamental to many models and methods in mathematical finance.

The scaling factor $\sqrt{h}$ is our approach for creating a large number of return scenarios at any time scale of h hours. For example, figure1 shows hourly BTC-USD spot price time series data over a one-year period ending on 1st July 2023, in the log space, overlaid with hourly historical changes. The cumulative distribution of the hourly historical changes is shown in Figure2 (a), together with those of the daily ($h = 24$) changes and the hourly changes scaled by $\sqrt{24}$. The plot demonstrates that the scaling approach is reasonable, as the scaled hourly changes closely track the daily changes. The Q-Q plot in figure 2 (b) provides another view of the same data, showing that the scaled hourly changes are distributed similarly to the daily changes.

Alternatively, we may use overlapping h -hour returns, i.e., $(\ln S_{n+h} - \ln S_n)$ in the left-hand side of (2). However, the resulting scenarios are structurally autocorrelated and thus less meaningful from a statistical point of view. While this approach is also used in the industry, we prefer our approach for its simplicity and statistical soundness.




Figure 1: BTC-USD spot price time series data (blue line, left y -axis) over a one-year period ending on July 1, 2023. The hourly changes in the log space are also shown as grey lines on the right y -axis.




Figure 2: (a) Cumulative distributions of the changes in the BTC-USD spot price time series data in Figure 1. The changes are calculated as the difference between the historical values in the log space. The blue solid line represents the hourly changes, the blue dashed line represents the hourly changes scaled by $\sqrt{24}$, and the orange solid line represents the daily changes. (b) Q-Q plot of the hourly changes scaled by $\sqrt{24}$ against the daily changes.

2.2 Volatility Surface Simulations

A volatility surface is represented as a two-dimensional function, denoted as $\sigma_T(K)$. In this function, K represents the strike price and T signifies the time-to-expiry. The method used to simulate this entire surface entails selecting certain points on the surface, referred to as the *risk drivers*. The scenarios for these *risk drivers* are then simulated based on their respective historical variations. For remaining points on the surface, the scenarios are extrapolated or interpolated from the risk drivers.

The *initial* version of the Serenity software employs a single risk driver for simplicity. Specifically, we use the 3-month at-the-money (ATM) volatility. This is understood in terms of delta-moneyness, i.e., a 50% delta. While this approach effectively captures risks associated with parallel shifts in the surface, it fails to account for risks due to alterations along the strike or expiry dimensions, such as tilting or bending.

We acknowledge these limitations and plan to enhance future versions of the model. The improvements will involve incorporating additional risk drivers to encapsulate the risks currently not accounted for. This refinement will better enable us to capture a more comprehensive range of risks present in the volatility surface. Concretely, the n -th volatility surface scenario $\tilde{\sigma}_{T,n}(K)$ is simulated as

$$\tilde{\sigma}_{T,n}(K) = \bar{\sigma}_T \left(\frac{\bar{F}_T}{\tilde{F}_{T,n}} K \right) \left(\frac{\sigma_n^{\text{atm3m}}}{\sigma_{n-1}^{\text{atm3m}}} \right)^{\sqrt{h}}, \quad (3)$$

where $\bar{\sigma}_T(K)$ is the reference volatility surface, and $\bar{F}_T$ and $\tilde{F}_{T,n}$ are the reference and simulated forward prices, respectively.

Note Recall from the work of reference [4] that the forward price F_T can be defined in relation to the spot price S and the projection rate p_T as follows:

$$F_T = S \exp(p_T T) \quad (4)$$

Consequently, the reference forward price $\bar{F}_T$ can be expressed using the reference spot price and the reference projection rate:

$$\bar{F}_T = \bar{S} \exp(\bar{p}_T T).$$

On the other hand, considering that only the spot price is simulated, the simulated forward price can be represented as

$$\tilde{F}_{T,n} = \tilde{S}_n \exp(\bar{p}_T T).$$

Equation (3) consists of two factors: The first factor models the *sticky moneyness*, which shifts the reference surface horizontally while preserving its smile shape with respect to the moneyness K/F_T (see, e.g., reference [5] for more details) This approach ensures that the at-the-money volatility of the simulated surface remains the same as that of the reference surface. The second factor represents the volatility perturbation modeled using historical changes of the volatility risk driver with a scaling factor based on the risk horizon h . Here, σ_n^{atm3m} is the n -th historical value of the 3-month at-the-money volatility. For a visual illustration of the effect of each factor and their combined effect, see figure 3, which shows three different combinations of 50% increase in the forward price and 20% increase in the volatility driver, and the resulting simulated volatility smiles.

Similar to the spot price simulation model in equation (1), the factor $\sqrt{h}$ in equation (3) is to scale hourly changes to scenarios equivalent to those over a period of h -hours. For example, figure 4 shows hourly historical time series of the 3-month at-the-money volatility and their changes in the log-space. The time series data are obtained by interpolating the SVI-parameterized volatility surfaces (see, e.g., reference [1]) as described in reference [4]. By comparing the cumulative distribution of daily changes with that of hourly changes scaled by $\sqrt{24}$ in figure 5, we see that the scaling approach is reasonable: The scaled hourly changes track the daily changes closely, although the former produces a wider (thus more conservative) distribution.

2.3 Other pricing variables

In the current version of the risk service, pricing variables other than spot prices and volatility surfaces are not simulated *directly*.

Perpetual and futures prices Perpetual or futures prices are simulated using historical changes in the corresponding spot prices while keeping the multiplicative basis unperturbed. Specifically,




Figure 3: (a) Illustration of volatility simulation model described in equation (3). The reference curve is the 3-month BTC-USD volatility smile as of 2023-07-01, and the simulated smiles are produced by applying each of the corresponding scenarios in the table. In this panel, the smiles are expressed in terms of the strike price. The dashed vertical lines indicate the forward prices before and after the perturbation. (b) Contains the same information, but the smiles are expressed in terms of the log value of the strike price.

the n -th simulated perpetual price $\tilde{P}_n$ is given by

$$\tilde{P}_n = \bar{P} \left(\frac{S_n}{S_{n-1}} \right)^{\sqrt{h}}, \quad n = 1, 2, \dots, N \quad (5)$$

where $\bar{P}$ is the reference perpetual price of P and S_n is the n -th historical value of the spot price. The same approach is used for simulating futures prices.

Incorporating the basis risk into the historical simulation framework presents a practical challenge in terms of sourcing quality historical price quotes. For instance, perpetual contracts on some underlying assets may not be liquid, and the basis is expected to vary across different exchanges, requiring the sourcing of historical price data by exchange. Therefore, incorporating the perpetual basis into the simulation framework is a topic for future research.

Interest rate curves The current simulation framework does not perturb interest rate curves. However, in the next version of the model, we may consider enhancing it by including projection curves derived from Deribit futures quotes. We have access to their historical data as part of our ongoing process of constructing volatility surfaces.

3 Generating P&L scenarios: option contracts

Having laid out the methods for generating historical scenarios for pricing variables, we can now apply them to calculate P&L scenarios of a portfolio of digital asset contracts. In this section, we use




Figure 4: Same as figure 1 but for 3-month at-the-money BTC-USD volatility time series data.




Figure 5: Same as figure 2 but using 3-month at-the-money BTC-USD volatility time series data in figure 4.

an option contract to illustrate how simulated scenarios affect the P&L distribution of the contract. Specifically, in addition to the default settings in §2, we assume the as-of-time is 2023-07-01 UTC and the option contract is a BTC-USD call option with strike price 1.3 times of the spot price. All reference and historical data are sourced from Cloudwall’s internal database.

3.1 P&L scenarios by pricing variables

Figure 6 illustrates the probability density distributions of the option contract’s profit and loss (P&L) based on different simulated pricing variables. As expected, the density distribution resulting from simulating both spot prices and volatility surfaces exhibits the widest spread. It is followed by the density distribution for spot simulation alone, and then by the density distribution for volatility simulation alone. The same order is observed in the cumulative distribution plots presented in Figure (b).

While these plots indicate that simulating spot prices holds greater significance than simulating volatility surfaces, they also confirm that the impact of simulating volatility surfaces cannot be overlooked. The influence of volatility simulation is further evident in the scatter plot shown in Figure 8, which will be discussed in the subsequent section.




Figure 6: (a) Density distribution of P&L scenarios from the option contract in §3 under different scopes of simulated pricing variables. (b) Same data but in terms of cumulative distributions. The horizontal dashed lines represent commonly used confidence levels for risk measurement purposes.




Figure 7: (blue) Number of P&L scenarios between 0.5th and 1.5th percentiles, grouped by weekly time interval. (red) Same but for scenarios between 98.5th and 99.5th percentiles.

One useful feature of historical simulations is the ability to identify the specific historical time stamps that contribute to P&L at a given confidence level. To illustrate this, figure 7 presents the time stamps that contribute to the P&L scenarios near the 1st and 99th percentiles. The bar chart for the 1% (99%) confidence interval is generated by counting the hourly historical time stamps that result in P&Ls falling between the 0.5th (98.5th) and 1.5th (99.5th) percentiles within each weekly time interval.

Notably, it is unsurprising to observe that there is a week in November 2022 (March 2023) with the highest number of P&L dates for the 1% (99%) confidence interval. These occurrences can be attributed to specific events, such as the FTX fraud event in November 2022 and the SVB default in March 2023.¹

3.2 Full revaluations on P&L scenarios

When dealing with nonlinear products like options, it is essential to conduct full revaluations of the portfolio for each scenario to obtain accurate P&L outcomes. Linear approximations, such as delta approximation, are insufficient for this purpose. Figure 8 depicts the P&L distributions of the option contract based on whether full revaluations are performed or not.²

¹November 2022: FTX fraud event, March 2023: SVB default.

²Options are valued using the Black-Scholes-Merton formula. See the seminal Black-Scholes paper in reference [2].




Figure 8: (a) Comparison of the P&L distribution using the full revaluation approach (orange) with that using the first-order delta approximation (grey) assuming only the spot price is simulated. The complete version (blue) of using the full revaluation and simulation of both the spot price and the volatility surface is also shown as a benchmark. (b) Scatterplots of P&L scenarios in terms of simulated spot price changes.

Comparing the density distributions, the distribution for the delta approximation is narrower than that obtained from full revaluation with spot price simulations alone. This indicates that the full revaluation approach captures the nonlinear effects associated with changes in spot price more accurately, which are not accounted for by the delta approximation. The scatterplot presented in Figure 8 provides clear evidence of how full revaluation captures the nonlinear impact in relation to spot price fluctuations, which is not adequately captured by the delta approximation. Additionally, the plot highlights the influence of volatility surface simulations (depicted by the blue dots), resulting in a more scattered and wider P&L distribution, as observed in the previous section.

4 Risk measures and application to risk management

Once a set of P&L scenarios is generated, we can calculate various distribution measures that offer insights into specific characteristics of the P&L distribution. For example, risk measures, such as Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR, a.k.a. expected shortfall), focus on potential losses at a given, or above, confidence level. Appendix A provides additional details on the definition of commonly used distribution measures, including VaR and CVaR. Investors and portfolio managers employ a combination of these measures to make informed decisions for effective risk management.

To illustrate the utilization of distribution measures in the context of risk management, let's consider a portfolio comprising 1 BTC and 10 ETH tokens. Figure 9 (a) displays the risk measure output for this portfolio, computed on July 18th. Suppose our objective is to manage portfolio risk by minimizing potential losses without significantly impacting the potential for gains. Achieving this goal would be challenging if only linear products were available, as they often affect both loss and gain potentials simultaneously. However, this can be accomplished by holding a put option position. For instance, figure 9 (b) demonstrates the output after incorporating 4 units of BTC-USD put options with a strike price of 28,000 and an expiration date of 2023-08-23. The loss measures (CVaR and VaR) are substantially reduced, while the gain measures (CGaR and GaR) are not significantly affected. It's important to note that this risk reduction comes with the cost of the option premium,

portfolio (a)

	Risk Measures		Dispersion Measures			Gain Measures	
Assets ↑	CVaR 95%	VaR 95%	Down Dev	Std Dev	Up Dev	GaR 95%	CGaR 95%
▼	\$3,050	\$1,688	\$1,012	\$1,354	\$1,217	\$1,775	\$3,406
Bitcoin (BTC)	\$1,711	\$936	\$575	\$804	\$779	\$966	\$1,994
Ethereum (ETH)	\$1,339	\$752	\$437	\$550	\$438	\$809	\$1,413

portfolio (b)

	Risk Measures		Dispersion Measures			Gain Measures	
Assets ↑	CVaR 95%	VaR 95%	Down Dev	Std Dev	Up Dev	GaR 95%	CGaR 95%
▼	\$1,290	\$826	\$410	\$1,229	\$1,716	\$1,487	\$3,430
Bitcoin (BTC)	\$782	\$642	\$221	\$272	\$147	\$798	\$1,291
Ethereum (ETH)	\$1,318	\$787	\$404	\$498	\$495	\$814	\$1,350
OPT BTC 04Aug23 280	-\$809	-\$603	-\$215	\$459	\$1,075	-\$125	\$789

Figure 9: Screenshots of the output tables from the Serenity risk measure service UX. The first row of each table provides a classification of the contribution measures in the second row. The third row shows the portfolio-level metrics for each measure and the contributing P&Ls from each asset in the portfolio to the portfolio-level metrics are shown from the fourth row and below. Portfolio (a) consists of two token assets, while portfolio (b) contains put options additionally.

which is not reflected in the risk measure output. Nonetheless, the risk measure output serves as a starting point for further analysis, such as identifying the optimal option hedge that minimizes risk while maximizing gain potential.

5 Conclusions

This document provides a comprehensive methodology³ for generating P&L scenarios within the Serenity risk service for digital assets, specifically focusing on options. The current framework for historical simulation encompasses spot prices and volatility surfaces, both of which are demonstrated to have significant effects on the P&L of option contracts. Furthermore, the importance of employing full revaluations is emphasized when dealing with nonlinear contracts. Lastly, we have illustrated the identification of hedging option positions as a means to reduce potential losses while minimizing the impact on potential gains. This highlights the capability of our risk measure service as an effective tool for risk management.

Disclaimer

This document is not financial advice, solicitation, or sale of any investment product. The information provided to you is for illustrative purposes and is not binding on Cloudwall Capital. This does not constitute financial advice or form any recommendation, or solicitation to purchase any financial product. The information should not be relied upon as a replacement from your financial advisor. You should seek advice from your independent financial advisor at all times. We do not assume any fiduciary responsibility or liability for any consequences financial or otherwise arising from the reliance on such information.

³as of July 2023

You may view this for information purposes only. Copy, distribution, or reproduction of all or any portion of this article without explicit written consent from Cloudwall is not allowed.

A Distribution measures

In this appendix, we present the definitions of commonly used distribution measures.

For a given set of profit and loss (P&L) scenarios and a percentile value $50 < p < 100$,

- Value-at-Risk (VaR) $p\%$: The $(100 - p)$ -the percentile of the P&L scenarios, where the sign is reversed to express the final value as a positive number.
- Conditional Value-at-Risk (CVaR) $p\%$: The average of the loss scenarios larger than VaR $p\%$.
- Gain-at-Risk (GaR) $p\%$: The p -the percentile of the P&L scenarios.
- Conditional Gain-at-Risk (CGaR) $p\%$: The average of the profit scenarios larger than GaR $p\%$.
- Standard Deviation (Std Dev): The standard deviation of the P&L scenarios.
- Downside Deviation (Down Dev): The standard deviation of the profit scenarios.
- Upside Deviation (Up Dev): The standard deviation of the loss scenarios.

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 4
- [2] F. Black, M. Scholes, *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy. 81 (3): 637–654, 1973 7
- [3] K. Givens, *Risk Neutral Discounting for Crypto Options in Serenity*, Research, Cloudwall, <https://cloudwall.tech/research>, 2023 1
- [4] Y. Lee, *Digital Asset Volatility Surfaces in Serenity*, Research, Cloudwall, <https://cloudwall.tech/research>, 2023 1, 4
- [5] D. Reisch, U. Wystup, *FX Volatility Smile Construction*, CPQF Working Paper Series, No. 20, <http://hdl.handle.net/10419/40186>, 2009 4