LUND UNIVERSITY FINANCE SOCIETY EST 1991

TRADING & QUANTITATIVE RESEARCH REPORT

What price will Bitcoin hit in 2025?

Modelling Crypto Prediction Markets as Barrier Options: Closed-Form and Monte Carlo Approaches

Analysts: Carl Nordahl, Alexander Ihrfelt, Alfred Hultgren
Supervisor: Oscar Näslund Cuesta

Abstract

This study examines whether cryptocurrency prediction market contracts can be valued using traditional derivatives pricing methods. We analyze Polymarket contracts tied to the event "What price will Bitcoin hit in 2025?", where payouts occur if Bitcoin reaches a specified level before year-end. These contracts closely resemble one-touch barrier options, linking prediction markets with option theory. Two valuation approaches are compared: a closed-form approximation and a Monte Carlo simulation with EWMA volatility and jump-diffusion dynamics using Deribit options and futures data together with observed Polymarket prices. Both models explain a substantial share of Polymarket price movements, suggesting these markets reflect economically meaningful probability assessments. However, the Monte Carlo framework performs more consistently across barrier levels and longer horizons, producing lower pricing errors, especially when jump risk is included. Delta-hedging tests also show materially lower P&L variance than unhedged positions. Overall, the findings show that crypto prediction markets can be analyzed with many of the same tools used in traditional derivatives markets. Despite liquidity and data limitations, one-touch option theory provides a credible benchmark for pricing, forecasting, and risk management in decentralized markets.

Table of Contents

Contents

1 Introduction & Theory **3**
1.1 Introduction . . . 3
1.2 Theory . . . 3
1.2.1 Prediction Markets . . . 3
1.2.2 Vanilla Option Theory . . . 4
1.2.3 Digital and One-Touch Options . . . 5
1.2.4 Monte Carlo Simulation . . . 5
1.2.5 RiskMetrics EWMA Volatility . . . 6
1.2.6 Jump Diffusion Model . . . 6
1.2.7 Δ-Hedging . . . 6

2 Data & Method **7**
2.1 Data . . . 7
2.2 Method . . . 7
2.2.1 Preprocessing . . . 7
2.2.2 IV calculation . . . 7
2.2.3 One-touch pricing: Method 1, Closed Form Approximation . . . 8
2.2.4 One-touch pricing: Method 2, Monte-Carlo simulation . . . 8
2.2.5 Visualization and Comparisons . . . 8

3 Results **9**
3.1 Method 1, Closed Form Approximation . . . 9
3.2 Method 2, Monte Carlo Simulation . . . 10
3.2.1 Initial Sample: November–December 2025 . . . 10
3.2.2 Extended Sample: January–December 2025 . . . 10
3.3 Delta-hedging . . . 11
3.3.1 Initial Sample: Closed-Form Delta . . . 11
3.3.2 Extended Sample: Monte Carlo Delta . . . 12

4 Analysis & Conclusion **14**
4.1 Analysis . . . 14
4.1.1 Method 1, Closed Form Approximation . . . 14
4.1.2 Method 2, Monte Carlo Simulation . . . 14
4.1.3 Delta-Hedge Performance . . . 15
4.1.4 Theoretical Implications & Limitations . . . 15
4.2 Conclusion . . . 16

5 References **17**

1 Introduction & Theory

1.1 Introduction

This study examines two valuation methods for Polymarket contracts as one-touch options, a class of exotic contracts defined by whether the underlying asset reaches a specific price level before expiration. Traditionally, instruments with these payoff structures were confined to professional financial markets. Prediction platforms have now pushed them into public view. Anyone with a crypto wallet can now trade event tokens that replicate payoffs once found only in structured products. Surrounded by political, cultural, and sports markets, however, it is unlikely most users recognise them as options at all. Markets built around questions such as 'What price will Bitcoin hit in 2025?' reveal real-money estimates of barrier outcomes and make the underlying probabilities visible in a straightforward way.

This matters because prediction markets convert individual expectations into prices that reflect the collective view of future events. Their prices can offer signals of practical relevance for both retail traders and institutions, and higher trading volume and greater participation from informed traders are likely to make these estimates more precise over time. Sentiment plays a central role here, since shifts in narrative and positioning can alter these probabilities in ways that differ from risk-neutral models and therefore shape how the results should be interpreted.

They also create a natural point of comparison with traditional option markets, where models such as Black-Scholes and quotes from venues like Deribit express probability and pricing structure through a more formal framework. Setting these sources alongside one another makes it possible to compare Polymarket's pricing with traditional option valuations, identify which view aligns more closely with outcomes, and consider how prediction-market probabilities might support forecasting, hedging, or related strategies.

1.2 Theory

1.2.1 Prediction Markets

Prediction markets, or betting markets such as Polymarket and Kalshi, are (as of this writing) still a relatively new development, with Polymarket itself launching in 2020. These platforms allow individuals to speculate on uncertain future events, referred to as markets, by taking positions against other participants. This is best illustrated by an example: a bettor, A, thinks that JD Vance will be the Republican presidential nominee in 2028. A then posts a bid for a 'yes'-token at, say, 55 cents. If another bettor B is willing to be the counter-party for that bet, B sells their 'yes'-token or equivalently buys a 'no'-token. This allows for investors to speculate on uncertain events without the involvement of a traditional house setting the odds. On Polymarket, each token pays out $1 if the event occurs, which means that the prevailing price of a yes/no token can be interpreted as the market-implied probability of the event occurring. Prediction markets can be understood as a specialized form of crowd-sourcing, designed to collect a wide range of views and expectations about specific questions of interest. A crucial improvement over ordinary opinion aggregation is that participants must risk capital, which filters out noise, rewards informed reasoning, and forces continuous updating as new information becomes available. In October 2025, Intercontinental Exchange (ICE), the parent company of the New York Stock Exchange, announced a strategic investment of up to $2 billion in Polymarket. The explicit purpose of the partnership is to distribute Polymarket's sentiment data to institutional clients. This reflects that prediction markets are increasingly treated as legitimate sources of probabilistic insight, validated by some of the largest institutions in the financial sector [1][2].

The consensus around prediction markets is that they give fairer odds than traditional betting houses, since there is no house involved.

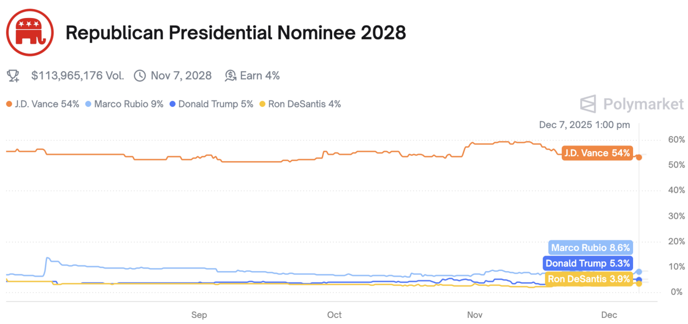

Figure 1: Example image of a market on Polymarket.

1.2.2 Vanilla Option Theory

This paper assumes that the reader has a basic knowledge of vanilla options and their payoff structures. However, pricing these options is not simple. The most accepted pricing model is the Black-Scholes formula. The model is named after economists Fischer Black and Myron Scholes, who first developed it. Black-Scholes is a method of pricing European call and put options, and works under the following assumptions:

- The return on the risk-free asset is constant and called the risk-free rate, r.
- The log-return of the underlying asset price follows a random walk with a drift that is equal to the risk-free rate. The volatility of the asset's returns is considered to be constant.
- The asset does not pay a dividend. This assumption is easily mitigated, although this is not necessary for the sake of the research in this project, since Bitcoin does not pay dividends.

Further, the model makes assumptions on the market that there are no arbitrage opportunities, and that there are no transaction costs.

The resulting closed-form evaluation results from solving the Black-Scholes Equation:

$$\frac{\partial V}{\partial t} + \frac{\sigma^2}{2} S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} - rV = 0 \tag{1}$$

where V: value of the option, S: underlying asset price, t: time. The method of solving the equation is left as an exercise to the reader. The solution of the equation results in the following closed-form evaluation of a European call option:

$$C = SN(d_1) - Xe^{-rT}N(d_2) \tag{2}$$

where d_1 and d_2 is calculated as:

$$d_1 = \frac{\ln\frac{S}{X} + (r + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}} \tag{3}$$

$$d_2 = d_1 - \sigma\sqrt{T}. \tag{4}$$

The variables in equations (2) through (4) are:

- C: theoretical price of the option.
- S: price of the underlying asset.
- $N(*)$: cumulative normal distribution function.
- X: strike price of the option.
- r: risk-free rate.
- T: time to expiry. Note: this is different from t in equation (1).
- σ: volatility of the underlying asset.

The theoretical price of a European put option, P, can then be derived from the Put-Call Parity:

$$P = Xe^{-rT}N(-d_2) - SN(-d_1). \tag{5}$$

The biggest issue with using the Black-Scholes in practice is estimating the volatility of the underlying asset. The volatility in the Black-Scholes formula is the volatility of the asset during the options lifetime. Since there is no known way of predicting the future, this cannot be observed. One can use historical price data to estimate realised volatility, but that gives limited results since volatility fluctuates with time. Because of this fact, investors often use the term "implied volatility", or IV. Given the price of a put or call option on the market, one can work their way backwards through the Black-Scholes formula to get the IV of the asset. This method is called the inverse Black-Scholes and is conceptualized in Figure 2:

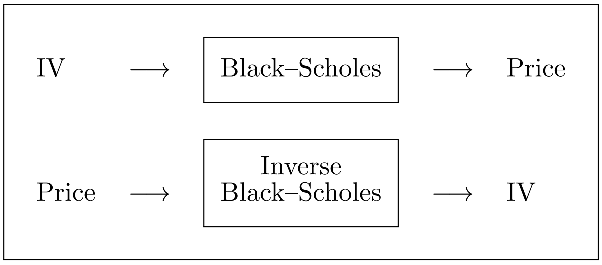

Figure 2: Conceptual illustration of Inverse Black–Scholes.

So, if an investor wishes to price another type of option and they need the volatility of the underlying asset, they can calculate the IV from options already trading on the market. If all the assumptions in Black-Scholes held in practice, the IV for all options with the same time to expiry would be the same. But this isn't the case. Like previously mentioned, Black-Scholes makes the assumption that log-returns of asset prices are normally distributed. This is empirically known to be a simplified reality. In reality, log-returns on equities have a negative skew and log-returns on Bitcoin can shift between negative and positive skews. Also, both classes have fatter tails than a normal distribution. This means outlier events, especially on the negative side, are more likely to happen than a normal distribution would suggest. Investors know this, and as a result, options that are far out of the money often trade at higher implied volatilities than at-the-money options [3][4].

1.2.3 Digital and One-Touch Options

The next step in the world of financial derivatives is to introduce digital options. Digital options are like European options, with the exception that their payoff is a set value, \$100 for example. Using theory from the Black-Scholes formula, a theoretical price for a digital call and put option is:

$$C_D = Qe^{-rT}N(d_2)$$

$$P_D = Qe^{-rT}N(-d_2)$$

where Q is the fixed payoff on the option, and d_2 is the same as in equation (4). These equations are simpler to intuitively understand than (2) and (5): the price of these options is the payoffs discounted to the present-day value, times some probability that the strike price, X, will be met or exceeded at the time of expiry.

The reason for introducing these digital options is because, with some assumptions, there is a simple, linear relationship between digital European options and one-touch options. One-touch options are options which give the holder an immediate, fixed payout if the underlying asset touches a fixed barrier, the option's strike, at any time during the option's lifetime. The formal payoff function for a one-touch option with a barrier above the underlying asset price is:

$$\text{Payoff} = \begin{cases} Q, & \text{if } \exists\, t \in [0,T] : S \geq X \\ 0, & \text{otherwise} \end{cases}$$

For the case where the barrier is below the underlying asset price, the payoff function is:

$$\text{Payoff} = \begin{cases} Q, & \text{if } \exists\, t \in [0,T] : S \leq X \\ 0, & \text{otherwise} \end{cases}$$

Under the assumption that the risk-free rate is zero, and therefore the drift of the underlying asset is zero, a one-touch option is worth 2 times the worth of a digital (call if barrier is above underlying asset price, else put) European option. This is because the options have the same expected payoffs. The next paragraph explains this relation for a one-touch option with a strike above the underlying asset price.

Consider a one-touch option with a payoff of \$1 and a strike price (barrier) at \$100, owned by Charlie. Another investor, Warren, owns two digital call options with the same payoffs, strikes, and expiration. One of two things can happen: 1. The asset price never reaches the asset price. Both Charlie and Warren receive a payoff of zero. 2. At some point, the asset reaches the strike price. Charlie immediately gets a payoff of \$1. Warren's two digital call options now have a 50% chance of finishing in the money (assuming zero-drift), so each is worth 50 cents. Warren sells his two digital options for 50 cents each and receives a payoff of \$1. Since Charlie's and Warren's portfolios have the same expected payoffs, they must be worth the same. For one-touch options with a strike *below* the underlying asset's current price, the same logic can be applied by modifying Warren's portfolio to consist of two digital *put* options. With this theory, there exists a closed-form valuation formula for one-touch options:

$$P_{OT} = 2Qe^{-rT}N(d_2), \text{ if } X > S_0, \tag{6}$$

$$P_{OT} = 2Qe^{-rT}N(-d_2), \text{ if } X < S_0. \tag{7}$$

Reminder that S is the underlying asset price, and X is the strike price, or barrier, of the one-touch option. Also, the assumption of $r = 0$ is only needed to motivate the above formulas, when using them for valuation, a proper value of r should be used. In reality, $r \neq 0$, which makes the estimations from these formulas worse as the options' expiries are further away in time. Far away is ambiguous, but a guideline that is set in this paper is that anything more than two weeks will significantly harm the precision of this evaluation [5].

1.2.4 Monte Carlo Simulation

Rather than directly relying on the closed-form Black–Scholes solution, an alternative approach is to simulate the underlying price dynamics using Monte Carlo methods. The Black–Scholes framework assumes that asset prices follow a geometric Brownian motion, implying that log-returns evolve according to a Brownian motion with drift. Over an infinitesimal time interval, the log-price process evolves as:

$$\ln S_{t+\Delta t} = \ln S_t + \left(\mu - \tfrac{1}{2}\sigma^2\right)\Delta t + \sigma\sqrt{\Delta t}\, Z,$$
$$Z \sim \mathcal{N}(0,1).$$

Notice that the log-return dynamics consist of two components: a deterministic drift term governed by the parameter μ, and a stochastic diffusion term driven by a standard normal random variable $Z \sim \mathcal{N}(0,1)$, scaled by the volatility parameter σ. The stochastic term captures the random fluctuations of returns and is proportional to the asset's implied volatility. Exponentiating the log-price expression yields the corresponding evolution equation for the asset price:

$$S_{t+\Delta t} = S_t \exp\left((\mu - \tfrac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}\, Z\right)$$

After simulating a large number of price paths, the probability of the asset reaching a barrier X is approximated by the fraction of paths that crossed the barrier at least once before maturity. With discounting for the time to maturity, the Monte Carlo estimate of the fair price becomes:

$$P = e^{-rT}\frac{1}{N}\sum_{i=1}^{N} I_i$$

where

$$I_i = \begin{cases} 1, & \text{if } \max_{0\le t\le T} S_t^{(i)} \ge X, \\ 0, & \text{otherwise.} \end{cases}$$

and N is the number of simulated paths.

1.2.5 RiskMetrics EWMA Volatility

To account for volatility changes over time in the simulations, this project uses a RiskMetrics EWMA approach. This involves updating the volatility parameter inside a simulation path using the past returns inside that path. The variance used in time step t is calculated as:

$$\sigma_t^2 = \lambda\sigma_{t-1}^2 + (1-\lambda)r_{t-1}^2,$$

where r_{t-1} is the log-return of the previous timestep, and λ is a decay factor. This ensures that volatile paths are updated with a higher volatility inside the simulation, and the opposite for stable paths. The decay factor ensures that past shocks have an exponentially decreasing effect on volatility. For daily time steps, the RiskMetrics standard uses $\lambda = 0.94$. σ_0 can be set as historical realized annual volatility [6].

1.2.6 Jump Diffusion Model

An advantage of the numerical Monte Carlo approach is that it's highly flexible. The randomwalk dynamics can be adjusted or extended without changing the overall pricing logic. In particular, Bitcoin often experiences sharp discrete jumps, to which one-touch options are highly sensitive. To account for this, an additional random jump mechanism à la Merton is added to the model:

$$S_{t+\Delta t} = S_t \exp\left((\mu - \tfrac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}\, Z + \sum_{i=1}^{N_{\Delta t}} J_i\right)$$

Here, N is a Poisson process with intensity λ, so $N_{\Delta t} \sim \text{Po}(\lambda\Delta t)$, where λ denotes the annual jump frequency. J denotes the logarithmic jump size, which is assumed to be normally distributed with a mean μ_J and deviation σ_J given by historical Bitcoin price data. The estimation procedure and data treatment are described in detail in the next chapter. A reasonable assumption is that these discrete jumps should be neutral on average, occurring just as often upward as downward. A common but incorrect simplification in the literature is therefore to set $\mu_J = 0$. However, since

$$\mathbb{E}\left[e^J\right] = e^{\mu_J + \frac{1}{2}\sigma_J^2}$$

the expected jump effect will be non zero whenever $\mu_J + \frac{1}{2}\sigma_J^2 \neq 0$. To keep the risk neutral pricing, $\lambda(\mathbb{E}\left[e^J\right] - 1)$ is subtracted from the drift. This leaves the final formula for the jump diffusion model:

$$S_{t+\Delta t} = S_t \exp\Big((\mu - \lambda\kappa - \tfrac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}\, Z + \sum_{i=1}^{N_{\Delta t}} J_i\Big),$$
$$\kappa = \mathbb{E}[e^J] - 1.$$

1.2.7 Δ-Hedging

A fundamental part of options-trading is reducing the overall risk for a portfolio. One way to do this, is with Δ-hedging (read: delta-hedging). An option's Δ is the option price's sensitivity to price changes in the underlying asset. By holding positions with opposite deltas (e.g. short call option and long underlying), one can eliminate the first-order directional risk of the underlying asset. Formally, the Δ of a one-touch option is calculated as:

$$\Delta = \frac{\partial P_{OT}}{\partial S}.$$

The following explicit formula for Δ can be obtained by differentiating equations (6) and (7):

$$\Delta = 2Qe^{-rT}\frac{n(d_2)}{S_0\sigma\sqrt{T}}, \text{ if } X > S_0, \tag{8}$$
$$\Delta = -2Qe^{-rT}\frac{n(d_2)}{S_0\sigma\sqrt{T}}, \text{ if } X < S_0.$$

Where $n(*)$ is the probability density function for a $N(0,1)$ stochastic variable. Note that $n(-d_2) = n(d_2)$.

The goal of delta-hedging is to make the total portfolio delta, $\Delta_{portfolio}$, zero. Note that the underlying asset has a delta of 1 by definition. So, to delta-hedge a portfolio that consists of short options, one should at all times own Δ_{option} number of the underlying asset.

For Model 2, no closed-form expression for Δ is available. The delta is therefore approximated numerically using a central finite difference:

$$\Delta \approx \frac{P_{\text{hit}}(S_0 + h) - P_{\text{hit}}(S_0 - h)}{2h}.$$

Rather than re-simulating paths for shifted spot values, the same Monte Carlo paths are reused. Due to the multiplicative structure of the model, a change in S_0 can be implemented by rescaling the barrier X,

$$X_{\text{up}} = X\frac{S_0}{S_0 + h}, \quad X_{\text{down}} = X\frac{S_0}{S_0 - h}.$$

The corresponding hitting probabilities are then recomputed using these adjusted barriers.

In the implementation, $h = 1000$ is used. This value is small relative to the spot level, while remaining sufficiently large to reduce numerical noise.

2 Data & Method

2.1 Data

The dataset combines cryptocurrency derivatives market data with prediction market prices. All timestamps are recorded in UTC.

Minute-level OHLCV price data for European Bitcoin call options were obtained from Deribit using the *public/get_tradingview_chart_data* API. In the initial test sample, the period covered is 26 November 2025 to 26 December 2025. The specific instruments analyzed have strikes of 90,000, 95,000, and 100,000 USD, all expiring on 26 December 2025. The close price from each minute bar is used as the observed option market price.

To represent the underlying asset, minute-level OHLCV price data were also collected for the December 2025 Bitcoin futures contract (`BTC-26DEC25`) from Deribit over the same period. Although Deribit BTC options are European, with final settlement determined by a 30-minute time-weighted average price (TWAP) of the Deribit index, they are quoted relative to the corresponding futures contract during trading. The matching-dated future therefore serves as the underlying price input in the implied volatility calculations.

Additionally, price data on perpetual futures were downloaded, also from Deribit. These prices serve as a proxy for Bitcoin's price, which was not available at a 1-minute resolution from Deribit. These values are later used in estimating one-touch option values, and in determining parameter choices for simulations. Prediction market data were collected from Polymarket for the event "What price will Bitcoin hit in 2025?", where each submarket corresponds to a specific price barrier. Historical price data for both YES and NO tokens were retrieved at approximately second-level resolution over the full event window (1 January 2025 to 31 December 2025). However, in the analysis presented in this report, Polymarket prices are only used over the overlapping period for which corresponding Deribit option and futures data are available.

2.2 Method

2.2.1 Preprocessing

The option prices gathered from Deribit were quoted in BTC. To prepare the data for IV calculation, these quotes were translated into USD using the current Bitcoin price (proxied by a perpetual future).

For the drift parameter in the pricing models, the futures-implied funding rate was used, derived from spot and futures prices:

$$r = \ln(F/S)/T,$$

where F, S and T are the future price, spot price, and time to expiry of the future contract. Strictly, this corresponds to the risk-neutral drift rather than a pure risk-free rate, but it is the appropriate input for both discounting and the drift of the simulated process under the risk-neutral measure.

2.2.2 IV calculation

To calculate the implied volatility of the call options from Deribit, the following algorithm was implemented:

1. Collect S_0, C, T, r, X for the specified option.
2. Make an initial guess for σ. Our algorithm used the guess $\sigma = 0.5$.
3. Calculate $C_{implied}(S_0, T, r, X, \sigma)$ using equation (2) and compare this to C using the squared difference, $(C_{implied} - C)^2$.
4. Use the *fmin* function from *scipy.optimize* package to minimize the squared difference by varying σ.

Table 1: Summary of Data Sources

Instrument	Source	Frequency	Sample Period
BTC Call Option (70K)	Deribit	1-minute	1 Jan–16 Dec 2025
BTC Call Option (80K)	Deribit	1-minute	1 Jan–28 Feb 2025
BTC Call Option (95K)	Deribit	1-minute	26 Nov–26 Dec 2025
BTC Call Option (100K)	Deribit	1-minute	26 Nov–26 Dec 2025
BTC Call Option (110K)	Deribit	1-minute	1 Jan–21 May 2025
BTC Call Option (150K)	Deribit	1-minute	1 Jan–16 Dec 2025
BTC Futures (BTC-26DEC25)	Deribit	1-minute	1 Jan–26 Dec 2025
BTC Perpetual Futures	Deribit	1-minute	1 Jan–31 Dec 2025
Polymarket Tokens	Polymarket	1-second	1 Jan–31 Dec 2025

The value of σ that this algorithm converges to is the implied volatility, IV, of the option. These values are then used in both one-touch option pricing methods.

2.2.3 One-touch pricing: Method 1, Closed Form Approximation

The valuation of the one-touch options using the first method described in the theory section was straightforward. The values of S_0, T, r, Q and σ were inserted into equations (6) or (7), depending on whether the strike price was above or below Bitcoin's current price, as explained earlier. The valuations were implemented and performed in Python.

2.2.4 One-touch pricing: Method 2, Monte-Carlo simulation

For each valuation timestamp, the Monte Carlo simulation is initialized with a set of time-dependent input parameters ($S_0(t)$, $\sigma(t)$, $\mu(t)$, $\tau(t)$). The remaining time to expiry, $\tau(t)$, is defined as the time (in years) between the current timestamp and the Polymarket contract's settlement time. The drift parameter, $\mu(t)$, was calculated using the relation between futures price and the spot prices as

$$\mu(t) = \frac{\ln\left(F(t)/S(t)\right)}{\tau_f(t)}$$

where $F(t)$ denotes a dated future corresponding to the Polymarket contract's settlement time. To model time-varying volatility, an exponentially weighted moving average (EWMA) specification was used, following the RiskMetrics framework described earlier. σ_0 was calculated as a 60-day rolling volatility on hourly log-returns of Bitcoin prices, this was then annualized to match the input required for the Monte-Carlo setup. The conditional variance was updated recursively at each time step, allowing the volatility to adapt dynamically to recent market conditions.

To determine the jump parameters, the price history of Bitcoin perpetual futures between February and December 2025 were examined. After calculating hourly log-returns in Bitcoin price, a rolling standard deviation was calculated using 60 (minimum 30 for early values) recent values. Log-returns with a z-score of 4 or greater were treated as jumps. The z-score threshold of 4 is a modelling assumption and was chosen as a conservative cutoff, meaning that only observations that are extreme relative to recent volatility are classified as jumps. A lower threshold would naturally lead to a higher estimated jump intensity λ, so the resulting parameters should be viewed as dependent on the chosen threshold rather than uniquely determined by the data. The amount of jumps in a year, λ was recorded, as well as the mean, μ_J, and standard deviation, σ_J, of the jumps.

μ_J	-0.0014
σ_J	0.025
λ	50

Parameters for the jump diffusion model

A separate dataframe was constructed for each barrier level, corresponding to a Polymarket contract of the form "Will Bitcoin hit X?". Each row represents a valuation timestamp and contains the associated time-dependent inputs $S_0(t)$, $\sigma(t)$, $\mu(t)$, and $\tau(t)$, used as initial conditions in the Monte Carlo simulation.

The model was then evaluated independently for each row of the dataframe. For a given timestamp, the initial price S_0, volatility σ, drift μ, and time to settlement τ were read and used as inputs. The simulation was carried out using daily time steps, with the number of steps proportional to the remaining time to settlement.

For each timestamp, $10,000$ price paths were generated, and the barrier-hit probability was estimated as the fraction of paths whose maximum (or minimum, depending on barrier type) exceeded the barrier level X at any time before settlement.

In addition, the delta was estimated numerically for each timestamp using a central finite difference scheme, as described above. Both the estimated probability $p_{\text{hit}}(t)$ and the corresponding delta were then appended as new columns to the dataframe, which was exported as a CSV file for comparison with observed Polymarket prices.

Because Method 2 does not depend on Deribit option data, it is additionally evaluated on four contracts (70K, 80K, 110K, 150K) over the full 2025 period, where Method 1 is not feasible due to option illiquidity.

2.2.5 Visualization and Comparisons

Initially, the resulting time series from both methods were plotted against their respective price histories from Polymarket. The Polymarket data is event-driven, with uneven timestamps, while the calculated values populate every minute. This makes formal comparisons a challenge. To mitigate this, the time series were merged using the calculated values' timestamps as reference. For each timestamp, the calculated value was matched with the most recent value in the Polymarket data. This ensures that no look-ahead bias is introduced in the evaluations.

3 Results

3.1 Method 1, Closed Form Approximation

Figures 3 and 4 show how the calculated values from method 1 compares to the prices on the corresponding Polymarket contracts.

Figure 3: Polymarket price and method 1 valuation, $100K strike price.

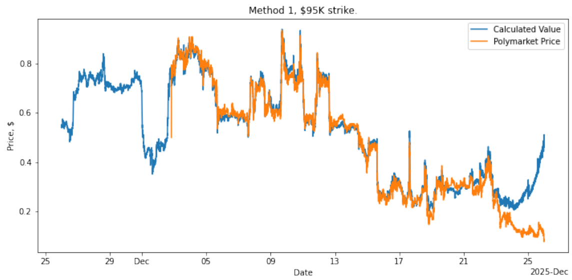

Figure 4: Polymarket price and method 1 valuation, $95K strike price.

In figures 5 and 6 the differences in valuation and price are plotted over time. For the $100K strike price contract, the mean absolute error (MAE) was 0.0405 and the root mean squared error (RMSE) was 0.0685. The corresponding values for the $95K strike price contract were 0.0337 and 0.0633.

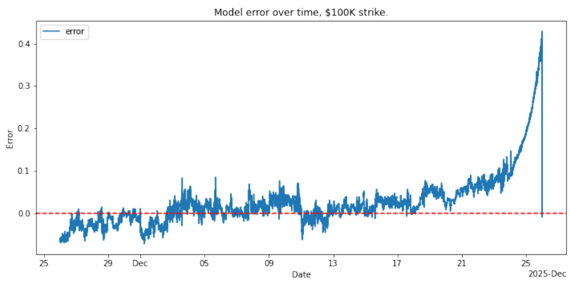

Figure 5: Error over time, $100K strike price.

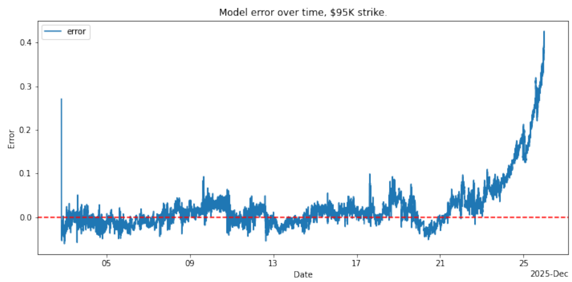

Figure 6: Error over time, $95K strike price.

To test if the errors are mean-reverting, Augmented Dickey Fuller tests were performed on the spread between the calculated values and Polymarket prices. The first test was performed on the full period, for the second test, the last 6 days were removed from the data. The results are presented in table 2. A p-value below 0.05 leads to rejection of the unit-root null, indicating that the time series is stationary.

Barrier	Sample	MAE	RMSE	Bias
100K	Full	0.0405	0.0685	0.0261
100K	Trimmed	0.0225	0.0274	0.0045
95K	Full	0.0337	0.0633	0.0222
95K	Trimmed	0.0156	0.0196	0.0030

Valuation errors

Barrier	Sample	ADF Stat	p-value
100K	Full	-2.72	0.070
100K	Trimmed	-5.51	$1.97 \cdot 10^{-6}$
95K	Full	2.26	0.998
95K	Trimmed	-5.95	$2.15 \cdot 10^{-7}$

ADF test results

Table 2: Model performance and statistical properties for Method 1.

3.2 Method 2, Monte Carlo Simulation

3.2.1 Initial Sample: November–December 2025

Analog to Method 1, the calculated values and Polymarket prices for contracts with barriers $95K and $100K are plotted against each other. The values were calculated using the jump-diffusion model, and using a simpler specification without jumps. Error statistics and ADF test results are presented in table 3.

Figure 7: Method 2 valuation with jumps

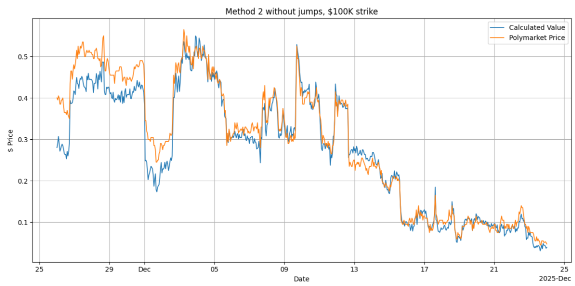

Figure 8: Method 2 valuation without jumps

Figure 9: Method 2 valuation with jumps

Figure 10: Method 2 valuation without jumps

Barrier	Model	MAE	RMSE	Bias
100K	Jump	0.0265	0.0317	0.0153
100K	No Jump	0.0265	0.0370	-0.0183
95K	Jump	0.0232	0.0277	-0.0045
95K	No Jump	0.0311	0.0372	-0.0283

Valuation errors

Barrier	Model	ADF Stat	p-value
100K	Jump	-3.5100	0.0077
100K	No Jump	-3.3998	0.0110
95K	Jump	-2.4170	0.1370
95K	No Jump	-2.3842	0.1463

ADF test results

Table 3: Model performance and statistical properties for 100K and 95K barriers.

3.2.2 Extended Sample: January–December 2025

To assess the robustness of the Monte Carlo approach over a longer horizon and across additional strike levels, the EWMA jump-diffusion model was evaluated on four additional contracts from the Polymarket event “What price will Bitcoin hit in 2025?”: $70,000, $80,000, $110,000, and $150,000. The evaluation covers varying periods within 2025, depending on contract activity, with hourly input data.

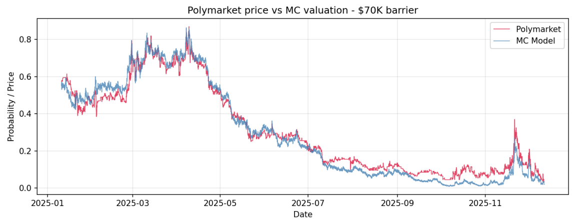

Figure 11: MC valuation vs Polymarket, $70K barrier.

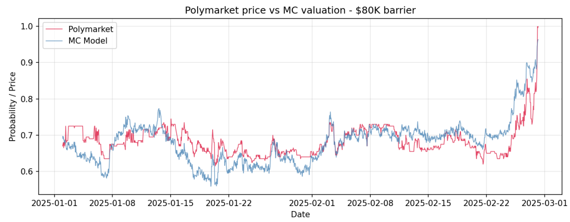

Figure 12: MC valuation vs Polymarket, $80K barrier.

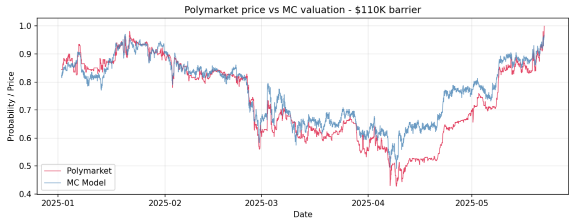

Figure 13: MC valuation vs Polymarket, $110K barrier.

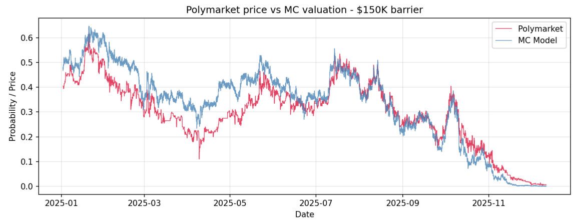

Figure 14: MC valuation vs Polymarket, $150K barrier.

Barrier	Model	MAE	RMSE	Bias
150K	Jump	0.0552	0.0684	0.0349
150K	No Jump	0.1161	0.1364	-0.1156
110K	Jump	0.0428	0.0563	0.0334
110K	No Jump	0.0682	0.0829	-0.0616
80K	Jump	0.0351	0.0424	-0.0032
80K	No Jump	0.2062	0.2293	-0.2057
70K	Jump	0.0367	0.0420	-0.0132
70K	No Jump	0.1239	0.1363	-0.1239

Valuation errors

Barrier	Model	ADF Stat	p-value
150K	Jump	-2.2451	0.1903
150K	No Jump	-1.9512	0.3084
110K	Jump	-2.9121	0.0440
110K	No Jump	-3.5922	0.0059
80K	Jump	-2.7409	0.0672
80K	No Jump	-1.3150	0.6223
70K	Jump	-3.5321	0.0072
70K	No Jump	-3.4704	0.0088

ADF test results

Table 4: Model performance and statistical properties across barrier levels.

3.3 Delta-hedging

3.3.1 Initial Sample: Closed-Form Delta

The following Δ-values were calculated using the closed form delta-approximation in equation (8) for the one-touch option with a barrier of $100000. These values are visualized in figure 15 and summarized in table 5.

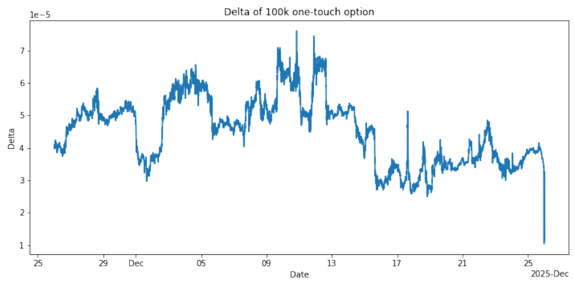

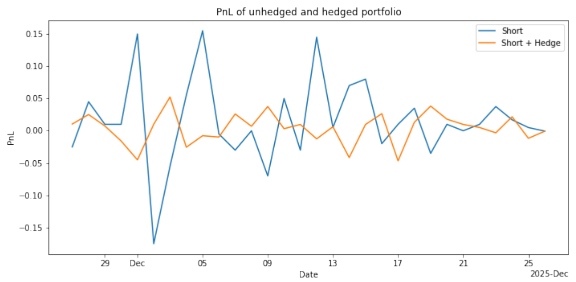

Figure 15: Delta of Polymarket contract over time. Figure 16: PnL of hedged and unhedged portfolio.

To test the effectiveness of the delta-hedge, the PnL variances of two portfolio strategies were compared: I) Being short one Polymarket contract. II) Being short one Polymarket contract and hedging by owning Δ Bitcoin as well. The rebalancing was done daily.

Table 5: Hedge statistics

Barrier	Live Days	Var Unhedged	Var Hedged	Reduction	Hit
$100K	30	0.00433	0.00055	87.2%	No

As portrayed in table 5, the variance of the PnL of the hedged portfolio, was 87.2% lower compared to the unhedged portfolio.

3.3.2 Extended Sample: Monte Carlo Delta

The delta-hedging evaluation was repeated for the extended sample, now using Δ values computed from the EWMA jump-diffusion Monte Carlo model via central differencing. Four barrier levels were evaluated: $70,000, $80,000, $110,000, and $150,000. The hedging methodology is identical to the initial sample: the PnL variances of two portfolio strategies were compared: I) Being short one Polymarket contract. II) Being short one Polymarket contract and hedging by owning Δ Bitcoin as well. The rebalancing was done daily. For contracts where the barrier was reached during the sample period ($80K on 28 February 2025, $110K on 21 May 2025), only the live trading period before settlement is included in the hedge statistics.

The Δ-values for the $70,000 barrier are visualized in figure 17 and the PnL for the hedged and unhedged strategies are shown in figure 18.

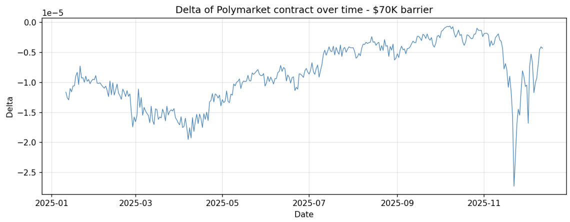

Figure 17: Delta of Polymarket contract over time, $70K barrier.

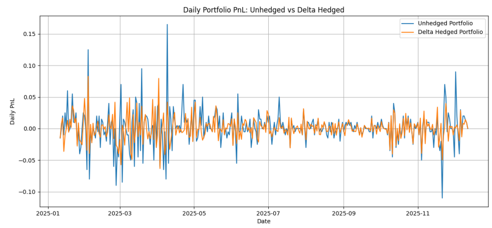

Figure 18: PnL of hedged and unhedged portfolio, $70K barrier.

The corresponding results for the $150,000 barrier are shown in figures 19 and 20:

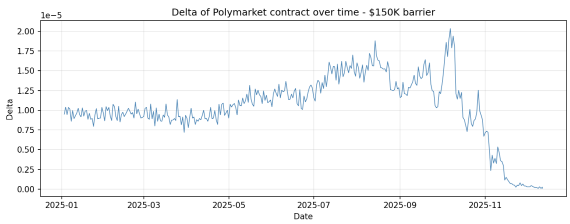

Figure 19: Delta of Polymarket contract over time, $150K barrier.

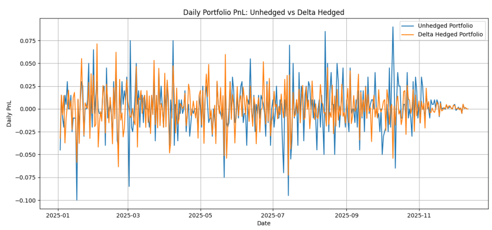

Figure 20: PnL of hedged and unhedged portfolio, $150K barrier.

The hedge statistics across all four barriers are summarized in table 6. The variance reduction ranges from 27.7% to 56.9%, with the strongest reduction observed for the $70K barrier.

Table 6: Hedge summary across all barriers (live period only)

Barrier	Live Days	Var Unhedged	Var Hedged	Reduction	Hit
$70K	337	0.000773	0.000333	56.9%	No
$80K	57	0.000733	0.000384	47.6%	Feb 28
$110K	140	0.000692	0.000500	27.7%	May 21
$150K	345	0.000648	0.000444	31.5%	No

4 Analysis & Conclusion

4.1 Analysis

4.1.1 Method 1, Closed Form Approximation

Figures 3 and 4 illustrate that the theoretical valuations from Method 1 track the Polymarket prices reasonably closely over most of the sample period. The error plots in Figures 5 and 6 show that deviations between market prices and model values fluctuate around zero rather than drifting persistently in one direction.

The magnitude of the pricing errors is moderate. For the $100K barrier contract, the mean absolute error (MAE) is 0.0405 and the root mean squared error (RMSE) is 0.0685. For the $95K barrier, the corresponding values are slightly lower (MAE = 0.0337, RMSE = 0.0633). The RMSE exceeding the MAE in both cases indicates the presence of occasional larger deviations, consistent with short-lived spikes rather than systematic bias. Overall, the error statistics suggest that the model captures a substantial portion of the variation in market prices, though short-term discrepancies occur.

To assess whether these discrepancies are temporary or persistent, Augmented Dickey–Fuller (ADF) tests were conducted on the error $e_t = V_t - P_t$, where V_t denotes the implied value from the model and P_t the observed Polymarket price. Stationarity of e_t implies deviations from the model value are mean-reverting rather than permanent. Over the full sample, the ADF results differ across barriers. For the $100K contract, the test statistic of -2.72 (p = 0.070) is close to conventional significance levels but does not allow rejection of the unit root null at the 5% level. For the $95K contract, the positive test statistic (2.26, p = 0.998) strongly indicates non-stationarity.

Inspection of the data reveals that the final six days of the sample are affected by a mismatch in contract expiry between Deribit and Polymarket contracts, which mechanically inflates implied volatility and induces large, artificial movements in the spread (see Appendix I). These movements are unrelated to the economic performance of the pricing model and instead reflect a data construction issue. When this mechanically distorted period is excluded, the ADF statistics change substantially. For both barriers, the unit root null is strongly rejected (-5.51 and -5.95, with p-values around 10^{-6}). These results indicate that, conditional on consistent contract specifications, the spread between market prices and model-implied values is stationary.

Taken together, the evidence suggests that deviations from the theoretical valuation are temporary when the underlying data are economically consistent. In such periods, market prices appear to revert toward the model-implied value rather than diverging persistently. The failure to reject a unit root in the full sample is attributable to a structural distortion in the final portion of the data rather than to systematic model misspecification.

4.1.2 Method 2, Monte Carlo Simulation

The Monte Carlo model tracks Polymarket prices closely across all tested barrier levels, with the deviations between model values and observed prices generally fluctuating within a relatively narrow range. The error statistics show that the model achieves low MAE and RMSE across most barriers, particularly for barriers closer to the prevailing Bitcoin price.

However, it is important to emphasize that proximity to Polymarket prices should not be interpreted as a measure of correctness in itself. The model is derived under a risk-neutral framework, whereas Polymarket prices reflect a combination of risk preferences, sentiment, and market microstructure effects. As a result, a model that consistently matches Polymarket prices may simply be replicating these effects rather than providing a theoretically grounded valuation. The objective is therefore not to minimize the distance to Polymarket prices, but to assess whether deviations are systematic or temporary.

The jump parameters were estimated from hourly log-returns using a rolling z-score threshold, making the identification regime-aware rather than applying a fixed cutoff. The resulting mean μ_J = -0.0014 is close to zero, reflecting approximately symmetric jump direction, which is economically reasonable.

The effect of the jump component is also noteworthy. The inclusion of jumps reduces both MAE and RMSE for barriers both above and below the prevailing spot level. This suggests that the improvement is not driven by directional bias, but rather by a more accurate representation of tail risk. If the jump component were introducing a systematic upward or downward drift, the improvement would be expected to be asymmetric across barriers, which is not observed.

At the same time, the specification has clear limitations. Jump sizes are drawn from a normal distribution centered near zero, meaning that most simulated jumps are small and do not resemble the discrete shocks observed in practice. As a result, the model captures increased variance rather than distinct jump events. A more realistic specification would allow for larger, more discrete moves, for example through a bimodal or heavy-tailed distribution. This limitation is likely most relevant for barriers where tail events dominate the valuation.

The ADF results show that the persistence of the pricing errors varies across barriers and model specifications. For some contracts, the residuals are stationary, indicating that deviations from the model are short-lived, while in other cases the unit root null cannot be rejected, suggesting more persistent discrepancies.

This indicates that the model does not consistently capture the time dynamics of pricing errors. In particular, some contracts exhibit low MAE and RMSE while still showing non-stationary residuals, implying that small mispricings can persist over time rather than being quickly corrected.

4.1.3 Delta-Hedge Performance

The hedging results suggest that both valuation frameworks are useful not only for pricing, but also for practical risk management. In the initial sample, where the closed-form approximation was used to calculate the delta for the $100,000 one-touch contract, the estimated Δ remained small and relatively stable throughout most of the sample, with a mean of 0.000045. The pronounced drop around December 26 is best understood as a consequence of the previously discussed expiry mismatch rather than as an economically meaningful feature of the contract itself.

To evaluate hedge effectiveness in this initial setting, the variance of PnL was compared across two portfolio strategies: (I) a short position in one Polymarket contract, and (II) a delta-hedged portfolio consisting of a short Polymarket contract and a long position of Δ Bitcoin, rebalanced daily. The variance falls from 0.00433 in the unhedged portfolio to 0.00055 in the hedged portfolio, corresponding to an 87.2% reduction. Figure 16 also shows that the hedged portfolio follows a substantially smoother PnL path. This indicates that, in the short initial sample, a large share of price variation in the contract can be attributed to first-order exposure to the underlying Bitcoin price, and that this exposure is captured well by the closed-form delta.

The extended-sample results provide a broader and more realistic test of hedge performance. Here, delta values were computed numerically from the EWMA jump-diffusion Monte Carlo model, and hedge effectiveness was evaluated across four barrier levels: $70,000, $80,000, $110,000, and $150,000. The same daily rebalancing strategy was applied, with hedge statistics for the $80,000 and $110,000 contracts calculated only over their live trading periods prior to settlement. Across all four barriers, the delta-hedged portfolios exhibit lower PnL variance than the corresponding unhedged portfolios. As summarized in Table 6, the variance reduction ranges from 27.7% to 56.9%, with the strongest hedge performance observed for the $70,000 barrier and the weakest for the $110,000 barrier.

Taken together, these findings show that the model-implied deltas remain economically useful even when moving from the stylized closed-form setting to the more flexible Monte Carlo framework. At the same time, the variance reductions in the extended sample are materially smaller than in the initial sample, suggesting that hedge performance becomes less precise once longer horizons, jump risk, and barrier-specific market dynamics are incorporated. This is consistent with the fact that daily rebalancing cannot eliminate all risk. The remaining variation in hedged PnL likely reflects higher-order exposures such as gamma risk, discrete rebalancing error, jump sensitivity, and model misspecification.

Overall, the results indicate that the pricing models capture an important share of the contracts' first-order sensitivity to Bitcoin and therefore provide hedge ratios that are useful in practice. However, the evidence also shows that one-touch prediction market contracts are not driven solely by contemporaneous movements in the underlying asset. Residual volatility remains meaningful, especially in the extended sample, implying that contract prices also reflect features that are not fully neutralized by delta hedging alone.

4.1.4 Theoretical Implications & Limitations

The finding that both models produce stationary error series suggests that Polymarket prices are not arbitrary or purely sentiment-driven, but track theoretically motivated valuations derived from traditional derivatives markets. This lends empirical support to the view that prediction markets aggregate information efficiently. The results also highlight an underappreciated structural connection between prediction markets and exotic options. Most participants likely think of these contracts as simple bets, but the pricing dynamics are governed by the same probabilities that professional options desks compute daily. Sentiment remains a complicating factor however, as the models are risk-neutral while Polymarket prices reflect subjective beliefs influenced by narrative shifts and momentum.

The most significant data quality constraint is the illiquidity of certain Deribit option contracts, which traded as infrequently as once every thirty minutes. This means the implied volatility input is a stale estimate that may not reflect current conditions, particularly problematic during rapid price movements. Wide bid-ask spreads introduce additional noise into the IV calculation. This illiquidity, combined with the expiry mismatch between Deribit and Polymarket contracts, likely accounts for a meaningful portion of residual error in both models. Beyond data quality, the short one-month sample period limits the generality of the conclusions.

4.2 Conclusion

This report shows that Polymarket contracts on whether Bitcoin reaches a given price level can be meaningfully analyzed as one-touch options. While presented as prediction-market tokens, their payoff structure closely resembles that of exotic derivatives in traditional finance. The results suggest that this is not only a theoretical analogy, but a practically useful framework. Across both the closed-form approximation and the simulation-based framework, the calculated valuations track observed Polymarket prices reasonably well over large parts of the sample. This suggests that prices in these markets are shaped by economically structured probabilities and not solely by narrative or speculative sentiment. At the same time, the results do not imply that Polymarket prices coincide exactly with risk-neutral option values. Rather, the findings support the view that one-touch option theory provides a strong benchmark for understanding and interpreting these contracts.

Between the two valuation approaches considered, the Monte Carlo framework appears to be the more robust method. While the closed-form approximation performs reasonably well in the short initial sample, its accuracy is sensitive to the assumptions required for the digital-to-one-touch relationship and to distortions in the underlying data. In contrast, the Monte Carlo model remains effective across a broader range of barrier levels and over longer sample periods. Its flexibility allows it to incorporate time-varying volatility, path dependence, and jump risk, all of which are especially relevant when pricing Bitcoin-related barrier events. The extended-sample results show that this richer specification produces valuation errors that are generally lower than those from the simpler alternatives, while also providing economically meaningful hedge ratios. The comparison between jump and no-jump specifications further suggests that model choice should depend on the barrier being studied, since distant barriers appear more sensitive to tail dynamics than contracts closer to the money. Taken together, these results indicate that the simulation-based framework is better suited for the practical valuation of prediction-market contracts with option-like payoffs.

At the same time, the analysis makes clear that the main limitations of the framework do not arise from the core idea of treating these contracts as derivatives, but from the quality of the available inputs and the frictions present in the market environment. The expiry mismatch between Deribit and Polymarket contracts, the illiquidity of certain option series, stale implied-volatility estimates, and the simplified jump specification all affect the precision of the valuations. These issues are particularly visible in the final part of the initial sample, where mechanically distorted inputs lead to artificially large errors. The hedging results point in the same direction. Although model-implied deltas substantially reduce PnL variance, especially in the initial sample, the remaining residual risk shows that practical hedge performance is constrained by discrete rebalancing, higher-order exposures, and imperfect model inputs. The broader conclusion is therefore twofold: first, exotic option theory offers a useful and credible framework for analyzing crypto prediction markets; second, further improvements in data quality and model specification are likely to matter at least as much as further theoretical refinement. In that sense, the report suggests that prediction markets such as Polymarket can be studied with many of the same tools used in traditional derivatives markets, but that their full potential for pricing, forecasting, and trading applications depends on overcoming substantial empirical and microstructural limitations.

5 References

References

[1] Polymarket. Polymarket 101, n.d. URL `https://docs.polymarket.com/polymarket-101`. Polymarket Documentation.

[2] Intercontinental Exchange. ICE announces strategic investment in Polymarket, October 2025. URL `https://ir.theice.com/press/news-details/2025/ICE-Announces-Strategic-Investment-in-Polymarket/default.aspx`. Press Release, October 2025.

[3] Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973. doi: 10.1086/260062. URL `https://www.sfu.ca/~kkasa/BlackScholes_73.pdf`.

[4] Noshaba Zulfiqar and Saqib Gulzar. Implied volatility estimation of Bitcoin options and the stylized facts of option pricing. *Financial Innovation*, 7(67), 2021. doi: 10.1186/s40854-021-00280-y. URL `https://link.springer.com/article/10.1186/s40854-021-00280-y`.

[5] Akhilesh Ganti. Understanding one-touch options and their key benefits, November 2025. URL `https://www.investopedia.com/terms/o/onetouchoption.asp`.

[6] RiskMetrics Group. The riskmetrics 2006 methodology, 2006. URL `https://www.msci.com/documents/10199/d0905614-2771-46dc-b000-1a033146586a`. RiskMetrics Methodology.

[7] Deribit. public/get_tradingview_chart_data, n.d. URL `https://docs.deribit.com/api-reference/upcoming/market-data/public-get_tradingview_chart_data`. Deribit API Documentation.

Appendix I - Input Variables

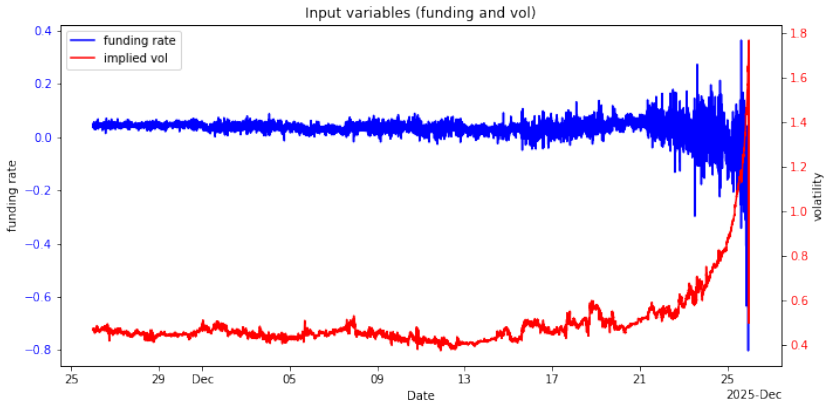

Figure 21: Input values for option valuations. Computed from option and spot data from Deribit. The spike in implied volatility near 26th of December is due to the Deribit contract expiring that date.

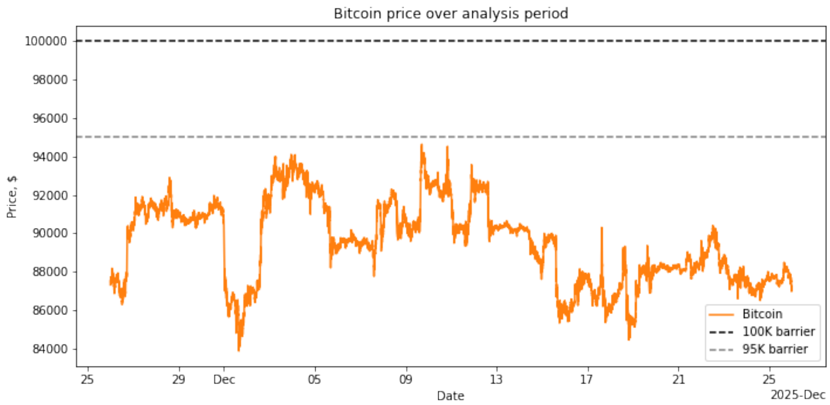

Figure 22: Bitcoin's price history during period of analysis (initial sample).

Disclaimer

Disclaimer

These analyses, documents, and any other information originating from LINC Research & Analysis (henceforth "LINC R&A") are created for information purposes only, for general dissemination, and are not intended to be advisory. The information in the analysis is based on sources, data, and persons which LINC R&A believes to be reliable. LINC R&A can never guarantee the accuracy of the information. The forward-looking information found in this analysis is based on assumptions about the future, and is therefore uncertain by nature, and using information found in the analysis should therefore be done with care. Furthermore, LINC R&A can never guarantee that the projections and forward-looking statements will be fulfilled to any extent. This means that any investment decisions based on information from LINC R&A, any employee or person related to LINC R&A, are to be regarded to be made independently by the investor. These analyses, documents, and any other information derived from LINC R&A are intended to be one of several tools involved in investment decisions regarding all forms of investments regardless of the type of investment involved. Investors are urged to supplement with additional relevant data and information, as well as consult a financial adviser before any investment decision. LINC R&A disclaims all liability for any loss or damage of any kind that may be based on the use of analyses, documents, and any other information derived from LINC R&A.

Conflicts of Interest and Impartiality

To ensure LINC R&A's independence, LINC R&A has established compliance rules for analysts. In addition, all analysts have signed an agreement in which they are required to report any conflicts of interest. These terms have been designed to ensure that COMMISSION DELEGATED REGULATION (EU) 2016/958 of 9 March 2016, supplementing Regulation (EU) No 596/2014 of the European Parliament and of the Council concerning regulatory technical standards for the technical arrangements for objective presentation of investment recommendations or other information recommending or suggesting an investment strategy and for disclosure of particular interests or indications of conflicts of interest.

Other