Evaluation of forecasting ability of risk-neutral density in Bitcoin option

BY

DIAS SAPARBEKOV

THESIS

Submitted in partial fulfillment of the requirements for the

degree of Master of Science in Finance

in the Graduate School of Business

Nazarbayev University, 2024

Astana, Kazakhstan

Advisor: Dr. Thierry Post

Abstract

This study assesses the out-of-sample forecasting capabilities of risk-neutral density models in Bitcoin options market, with a focus on the Normal Inverse Gaussian (NIG) density. Understanding forward-looking price dynamics becomes critical as cryptocurrencies continue to gain a reputation in financial markets. This research examines how the NIG model, with its capability to capture skewness and kurtosis, compares to the benchmark log-normal (LN) distribution. The analysis applies the likelihood ratio test to evaluate the predictive performance of the models. As a result, NIG model improves the accuracy of tail forecasts, outperforming LN in capturing extreme market movements, which holds implications for risk management and market timing in Bitcoin market.

Keywords: Bitcoin · Risk-neutral density · Normal inverse Gaussian · Likelihood ratio test.

1 Introduction

Options are financial derivatives that provide a buyer the right but not the obligation to buy and sell the underlying asset at an established price and period. Because of the increasing volume and turnover in options markets, these contracts have become a vital component of an investor's portfolio for hedging or speculation purposes. The growing popularity of various research that analyses this kind of financial instrument shows that it contains a massive amount of forward-looking information. One of the most interesting bodies of literature is based on extracting market expectations embedded in option prices using the risk-neutral valuation framework of Breeden

and Litzenberger (1978), Cox and Ross (1976) and Rubinstein (1976).

Risk-neutral density is a probability distribution that reflects investors' expectations regarding the future price of the asset, adjusted for risk. The estimation of the risk-neutral density has to be precise for modeling and density forecasting for both European put and call options. Option prices serve as a tool for monitoring current market conditions and they accurately reflect underlying assets, which, as a result, will help investors to be more informed. That is why the correct estimation of risk-neutral density is crucial for further utilizing it for various purposes.

Recently, one of the asset classes that most investors are keen to analyze for future price movements is cryptocurrencies. Although the development of the crypto environment started

more or less recently, the first digital currency was invented in the 1980s (Chaum, 1983). Unfortunately, the project did not gain much popularity, and the particular reason for the circumstance is simply the absence of a digital environment and the problem with double-spending. The successful attempt to introduce cryptocurrencies happened in 2008 when Satoshi Nakamoto (2008) presented the theoretical framework of blockchain technology. His idea of decentralized networks started a new era of the digital world and gave us Bitcoin.

Nowadays, cryptocurrencies and their derivatives are widely used by institutional and retail investors for many purposes, such as investing, hedging, and speculating. Proof of that can be that in mid-January 2024, the US Securities and Exchange Commission (SEC) approved the launch of eleven spot exchange-traded funds (ETFs) investing directly in Bitcoin. This his-

torical event helped to widen the number of institutional investors in crypto assets. Specifically, the SEC allowed funds from companies, including Grayscale, BlackRock and Fidelity Investments, to begin trading shares on the New York Stock Exchange (NYSE), NASDAQ, and the Chicago Board Options Exchange. As of November 2024, the global crypto market capitalization is $ 3.3 trillion, and the monthly volume for bitcoin options is approximately 100 billion.

For November 2024, 85% of the share of open interest across Bitcoin options is retained in the Deribit Exchange, one of the biggest decentralized and leading cryptocurrency futures and options exchanges. I chose Deribit as the data provider for high-quality data and retrieved a historical option chain from it.

I selected Bitcoin as the instrument for analysis using risk-neutral densities due to the significant interest among market participants regarding its future movements, coupled with the scarcity of existing research that could be used by investors. Most of the tools that are used by crypto practitioners are based on backward-looking analysis, which is based on historical data. With rapid development and huge volatility in the crypto environment, a lagging approach can be inappropriate to use. Crypto investors are most of the time wrong in terms of selecting the right time and the right cryptocurrency to invest. The proof for that is, for instance, crypto-tracked futures saw over $ 1.08 billion in liquidations in the past 24 hours in October 5th, 2024 [See CoinDesk, 2024]. For this kind of market, a forward-looking analysis incorporating market expectations for the asset is an essential tool for investors, specifically the

risk-neutral density, which can be very useful in terms of market timing and ranking cryptocurrencies.

The parametric approach was chosen for the estimation of risk-neutral density, precisely the Normal Inverse Gaussian method, which was offered by Eriksson, Ghysels and Wang (2009). It is a continuous stochastic process that is taken as a normal variance-mean mixture where the mixing density is the inverse Gaussian distribution. The normal Inverse Gaussian approach has good behavior in the tails and four parameters that can be calibrated to the empirical moments in the data. This was also demonstrated by Geman et al. (2001) that a superior alignment with the data drives this enhancement, refined option pricing and hedging methods, and theoretical factors.

In this paper, the risk-neutral density estimated using normal inverse Gaussian will be compared to the baseline model,

the Black-Scholes model, which assumes log-normality. The comparison will be made by the likelihood ratio test, which compares the density from the model at time “t” to the subsequent return between t and t+7. I will check whether the general (nesting) specification (NIG) outperforms the specific (nested) specification (LN) by testing whether the forecasting ability improves significantly using a likelihood ratio test.

2 Literature review

Initially, the risk-neutral density was introduced and discussed by Cox and Ross (1976). Brennan (1979) and Rubinstein (1976) analyzed the utility function of a ”representative agent” to achieve risk-neutral pricing. They demonstrated that constant relative risk aversion utility and log-normal returns are specific conditions for option pricing to comply with the Black-Scholes equation. It has been extended by numerous researchers since

then, but the first huge contribution to risk-neutral density literature was made by Breeden and Litzenberger (1978) and Banz and Miller (1978). They showed how to extract whole risk-neutral density from option market prices. To be precise, they expressed it as the second-order derivative of option prices as a function of the strike price; it was extended further by Aït-Sahalia and Lo (1998) and Longstaff (1995). There are many approaches to extracting risk-neutral density, for instance, the method of extracting the shape directly by Rubinstein (1994), Jackwerth and Rubinstein (1996), Melick and Thomas (1997), and others. This approach is similar to the non-parametric approach, which focuses on methods like curve fitting [See Aït-Sahalia and Duarte (2003), Bliss and Panigirtzoglou (2004), Härdle and Hlávka (2009), Figlewski (2008), Birru and Figlewski (2012)]. Besides that, numerous semi-nonparametric approaches

along the lines of functionals are based on the hypergeometric functions of Abadir and Rockinger (2003) and Bu and Hadri (2007); and the density mixture estimation of Bondarenko (2003) and Yuan (2009).

This paper focuses on the parametric estimation of risk-neutral density, specifically normal inverse Gaussian, motivated by the comparatively underdeveloped literature on parametric extraction methods. The NIG approach is attractive because it is defined by the first four moments: mean, variance, skewness, and kurtosis. This approach, which is initially proposed by Eriksson, Forsberg, and Ghysels (2004), is similar to the Gram-Charlier series expansion method offered by Madan and Milne (1994), who tested A-type expansion applied in the context of option pricing, and by Rompolis and Tzavalis (2007), who, on the other hand, focused on C-type expansion. Eriks-

son, Ghysels, and Wang (2009) discussed that the usage of the normal inverse Gaussian model has several edges over A-type Gram-Charlier series expansion (A-GCSE) and C-type Gram-Charlier series expansion (C-GCSE) because of the results with negative probabilities of A-GCSE or difficulty of computations of C-GCSE.

Bitcoin options are inverse options, which are option contracts quoted and traded in the units of the underlying cryptocurrency. Lucic and Sepphe (2024) stated that the valuation of inverse options seamlessly integrates with standard no-arbitrage principles by using a proper numéraire, cryptocurrency forward prices. As was discussed earlier, I obtained the option data from Deribit Exchange. All options on Deribit are of the European style, signifying that they can only be exercised at expiration, in contrast to American options, which can

be exercised at any point in time. Options on Deribit are also cash settled, which means when they are exercised, it is only the profits that are paid.

It is pretty intuitive that Bitcoin is a highly volatile and speculative instrument because of the crypto market environment's specifics. Baek and Elbeck (2015), Bouoiyour and Selmi (2015), Dwyer (2015), Pichl and Kaizoji (2017), and Ardia et al. (2019) employed ARCH/GARCH volatility investigation to analyze the time series of Bitcoin. Additionally, Scaillet, Treccani, and Trevisan (2017) found that in the Bitcoin market, volatility jumps are frequent events and that they cluster in time. These explorations make the non-parametric approach for risk-neutral density estimation too implausible. The particular reason for this circumstance is that, even if nonparametric methods are easy to compute and robust, typically, they are very data sen-

sible as they try to take the whole shape of an unknown density [See A it-Sahalia and Lo (1998), Pagan (1999), Ghysels, Patilea, Renault, and Torr‘es (1997)]. According to Malz (2014), his approach requires data of reasonably good quality on the implied volatility smile. All things that have been stated indicate that because of the frequent volatility jumps and limited strike range in the Bitcoin market, the usage of a nonparametric approach is unrealistic for this case.

It is also a fact that risk-neutral density is a biased estimator of physical density, which is also called real-world density, because of the absence of stochastic discount factor (SDF) as a multiplier [See McGee, Post, and Poti (2024)]. The risk-neutral density, physical density, and pricing kernel, which is a stochastic discount factor, are all univariate functions. So, the relationship is such that the risk-neutral density is the product

of the physical density and the pricing kernel. Intuitively, physical density is risk-neutral density discounted by a stochastic discount factor. The way to treat that relationship is that the fundamental equation of asset pricing is that the price of any asset, including options, is the expected payoff under the P distribution discounted at the pricing kernel. However, an identification strategy and high-quality extra data that reflect risk aversion are needed to estimate physical density, which is risky and challenging to complete with inverse options. Instead of focusing on estimating real-world density, we can decrease the dispersion between physical density and risk-neutral density. For that, we have to reduce the impact of the stochastic discount factor on Q-density. One possible way to do that is to use short-dated options to minimize the effect of pricing kernel. The kernel itself doesn't change, but the likelihood of observ-

ing extreme returns and extreme values of the kernel decreases with the shorter time horizon. That is why short-dated and liquid options will not make our risk-neutral density far from the real-world density. Also, the RWD is generally estimated using time series econometrics, which tends to underperform option-implied RND estimates. That's why the RND, despite being a biased estimator of physical density, often outperforms its estimates. Alexander, Deng, Feng, and Wan (2023) demonstrated that the most liquid bitcoin options have less than seven days to maturity, which gives more confidence to proceed with less biased risk-neutral density.

The log-normal model considers that asset returns follow geometric Brownian motion, resulting in a symmetric, unimodal density with two parameters: the mean (which is typically set to zero under risk-neutral assumption) and volatil-

ity. It is a foundation of the Black and Scholes (1973) framework. During the calibration phase, the volatility parameter is modified to reduce the variance between observed and model-derived option prices. Although it is computationally efficient, log-normal models often fail to account for higher-order moments such as skewness and kurtosis, which are common factors in financial markets, especially for volatile assets like Bitcoin.

The Normal Inverse Gaussian (NIG) distribution extends the LN model in the sense of integrating skewness and kurtosis, allowing it to better fit the data. Nevertheless, due to the higher dimensionality and non-convexity of the optimization problem, the random grid search technique was employed, which ensured that estimated parameters aligned the NIG-implied option prices with observed market prices.

Upon constructing densities, their validity is evaluated via tests. The LN and NIG densities, calibrated on a formation date using option prices, are used to forecast the probability distribution of returns over the forecast horizon. Upon expiration, actual returns are juxtaposed with the model predictions. The aggregate of log-likelihoods overall sampling intervals serves as an indicator of the overall forecasting efficacy Likelihood Ratio (LR) test developed by Vuong (1989) to evaluate the predictive effectiveness of rival forecasting models. To address sampling variation, the test is conducted on numerous random pseudo-samples derived from the original sample. The test evaluates the overall predictive capability across all market conditions. However, the Amisano and Giacomini (2007) extended the previous method, where the test is conducted out-of-sample rather than in-sample. It is accurate for both non-

nested and nested forecast models, and it takes time series rather than independent data. The test can evaluate whether the NIG model improves forecasting accuracy over the LN model, ranking out-of-sample performance by the weighted average of logarithmic scores.

3 Data

This study took the Bitcoin futures daily option chain from 2020-01-03 to 2024-11-15 with all available strike prices and maturities from the Deribit exchange. In option chains, there are several maturities, including 1-3 dailies, 1-3 weeklies, 1-3 monthlies, and 3, 6, 9, and 12-month quarterlies of the March, June, September, and December cycles. The raw data contains a timestamp, order type (call or put), option greeks, volume, instrument price (bid and ask), strike price, price of the underlying, implied volatility (bid and ask), open interest, and ex-

piration date. The daily snapshots' timestamps were at 08:00 UTC since all options from Deribit were expiring at that time. New option expiries are often introduced on Thursdays at 08:00 UTC, excluding dailies, resulting in one day when more expiries are available. Because of the liquidity considerations, European out-of-the-money calls and puts were taken into account in this analysis, filtered by delta measure. Only options maturing on Friday, specifically 7-day options, were utilized for further construction of RND. Nonfitting data points and illiquid options characterized by no volume were excluded from analysis at the initial step. Data description is shown in Table 1-2 and Figure 1-3.

4 Methodology

For each date t, where the underlying price is S, we observe a string of options op(t): puts and calls with a time to maturity

T. The Breeden-Litzenberg formula gives a way to estimate the probability point associated with the underlying reaching the strike K for each strike. I calibrated a Levy model NIG(κ, θ, δ, μ) so that option prices obtained from this model fit the option prices in the dataset. This enables us to calculate the density $f_{NIG}(x, \kappa, \theta, \delta, \mu)$ of future values of S_{t+T}. The same can be done for the Breeden-Litzenberg model; however, it is essential to mention that this model only yields the density values in the strike points. Then, log return was calculated

$$R = \ln\left(\frac{S_{t+T}}{S_t}\right)$$

After that, we can calculate the log difference in "forecasting" for the two models.

$$1) \ln\left(f_{\text{NIG}}(R, \kappa, \theta, \delta, \mu)\right) - \ln\left(f_{\text{BL}}(R)\right)$$

$$2) \ln\left(f_{\text{NIG}}(R, \kappa, \theta, \delta, \mu)\right) - \ln\left(f_{\text{LN}}(R)\right)$$

Repeated over a number of dates "t", we can then calculate the mean of this log difference and its deviation. We test the null hypothesis: the models are similar in forecasting, and H_1: they differ by comparing the test statistic to a Gaussian quantile.

To fit the NIG model, we first need a pricing function providing option prices according to NIG parameters; then, we will fit those parameters so that the error in option pricing and real option prices is small.

The most delicate issue was in the calibration phase. The process is to use a random grid search for a starting point combined with an L-BFGS algorithm. The random grid is made of a starting point for all parameters (except mu) with the addition of a random noise. Then the L-BFGS method was applied and

provided a convergence point. This is repeated several times (at the hand of the user) in order to maximize the chances of obtaining a minimum.

4.1 Normal Inverse Gaussian (NIG) Model

The Normal Inverse Gaussian (NIG) distribution is part of the generalized hyperbolic distributions and is capable of modeling skewness and excess kurtosis kept in returns. The NIG process is represented as a variance-mean mixture of normal and inverse Gaussian (IG) distributions, making it adaptable for capturing heavy tails and asymmetries.

The foundation of NIG was introduced by Seshadri (1993) in his paper, where he defined Y as a random variable that adheres to an inverse Gaussian probability distribution as follows:

$$\mathcal{L}(Y) = IG\left(\delta, \sqrt{\alpha^2 - \beta^2}\right)$$

By adding normally distributed conditional X with mean and variance, it became NIG, which will be unconditional:

$$\mathcal{L}(X) = NIG(\alpha, \beta, \mu, \delta)$$

So as a result, the probability density function (PDF) of the NIG distribution is:

$$f_{\text{NIG}}(x; \alpha, \beta, \delta, \mu) = \frac{\alpha\delta}{\pi} \exp\left(\delta\sqrt{\alpha^2 - \beta^2} + \beta(x - \mu)\right) \times \frac{K_1\left(\alpha\sqrt{\delta^2 + (x - \mu)^2}\right)}{\sqrt{\delta^2 + (x - \mu)^2}}$$

where steepness or tail thickness parameter is more than zero $(\alpha > 0)$; asymmetry or skewness parameter is represented as follows: $(\beta \in (-\alpha, \alpha))$; scale parameter or width

of the distribution, is also more than zero $(\delta > 0)$; location parameter (mean) is $\mu \in \mathbb{R}$ and $K_1(x)$ is a modified Bessel function of the second kind, also called the third kind, explained by Abramowitz (1974).

The characteristic function simplifies complex computations of risk-neutral density, such as integration over probability densities, into tractable forms. Schmelzle (2010) proposed using the Fourier transformation method for option pricing. In this paper, the NIG characteristic function represents the Fourier transformation of the probability density function of the NIG distribution. I used the Fourier method instead of the Monte Carlo approach because of the speed advantages of Fourier inversion. Computational speed is essential for calibration, as it requires numerous iterations of option pricing, particularly with noisy data from earlier dates for Bitcoin options.

In the NIG approach, the asset's log-returns $\ln\left(\frac{S_T}{S_0}\right)$ are modeled employing the NIG distribution. The characteristic function defines the transformed space of these log returns:

$$X_t \sim NIG(\alpha, \beta, \delta, \mu)$$

For the NIG process, the characteristic function is given by:

$$\phi_{\text{NIG}}(u; \mu, \theta, \sigma, \kappa) = \exp\left(iu\mu + \kappa\left(\sqrt{\theta^2 + \frac{\sigma^2}{\kappa}} - \sqrt{(\theta + iu)^2 + \frac{\sigma^2}{\kappa}}\right)\right)$$

where μ is the location parameter, adjusted under the risk-neutrality, θ is the drift parameter of the Brownian motion, σ is the volatility of the Brownian motion, and κ is the variance parameter of the inverse Gaussian process.

The definitions of characteristic function elements are as follows:

1) The drift term ($iu\mu$) is representing the deterministic trend of the log-returns of the asset. Under risk neutrality, μ is adjusted to ensure the martingale property.

2) The first term in the exponential $\left(\sqrt{\theta^2 + \frac{\sigma^2}{\kappa}}\right)$ describes the baseline behavior, mixing both the Brownian motion (σ) and the variance of the inverse Gaussian process (κ).

3) The second term in the exponential $\left(\sqrt{(\theta + iu)^2 + \frac{\sigma^2}{\kappa}}\right)$ contains the impact of the Fourier variable "u", which translates frequency-domain information into the distribution's moments. It also adjusts the process to mirror the skewness (θ) and excess kurtosis (through κ).

4.2 NIG option pricing

The Fourier inversion technique is used to calculate call (C) or put (P) option prices:

$$C = S_0 \cdot Q_1(k, \phi_{NIG}) - Ke^{-rT} \cdot Q_2(k, \phi_{NIG}),$$

$$P = -S_0 \cdot (1 - Q_1(k, \phi_{NIG})) + Ke^{-rT} \cdot (1 - Q_2(k, \phi_{NIG}))$$

where Q_1 and Q_2 are integrals derived from the characteristic function:

$$Q_1(k) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re} \left(\frac{e^{-iuk}}{iu} \frac{\phi(u-i)}{\phi(-i)} \right) du$$

$$Q_2(k) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re} \left(\frac{e^{-iuk}}{iu} \phi(u) \right) du$$

where:

$$\phi(u) = e^{t\left(i\mu u + \frac{1-\sqrt{1-2i\kappa\theta u + \kappa\sigma^2 u^2}}{\kappa}\right)}$$

and k = $\ln\left(\frac{K}{S_0}\right)$ defined as log-moneyness.

The integrals are solved numerically using **Fourier transform**, which speeds up the computation.

4.3 Martingale correction

To implement the no-arbitrage condition under risk-neutrality Q, the drift parameter is modified such that:

$$\mathbb{E}^Q[S_T] = S_0 e^{rT}$$

The martingale correction changes the drift to validate that the expected discounted price is aligned with the risk-free rate.

$$w = \frac{1 - \sqrt{1 - 2\beta\delta - \delta\sigma^2}}{\delta}$$

The standard deviation of the NIG process was formalized as:

$$\text{dev}_X = \sqrt{\sigma^2 + \theta^2\kappa}$$

4.4 Calibration of NIG Parameters

The calibration phase aligns the NIG model with market option prices by minimizing pricing errors; to be precise, it min-

imizes the mean squared error between observed and model-implied option prices.

$$\text{Cost Function} = \frac{1}{N}\sum_{i=1}^{N}\left(\frac{C_{\text{model}}(K_i,T_i) - C_{\text{market}}(K_i,T_i)}{C_{\text{market}}(K_i,T_i)}\right)^2$$

This optimization was completed over the NIG parameters (κ, θ, σ) utilizing L-BFGS-B, which is a limited-memory algorithm for decoding extensive nonlinear optimization problems subject to straightforward bounds, ensuring:

$$\kappa > 0,\ \sigma > 0,\ -10 \leq \theta \leq 10$$

4.5 Risk-neutral density under the Black-Scholes (BS) with different volatility for every strike

Risk-neutral density was defined as "strike gamma" since the second derivative of the identified option price for the strike prices was taken.

I derive the Breeden-Litzenberg RND by numerically estimating the second derivative of the call price with respect to the strike price. The first derivative with respect to strike price:

$$\frac{\partial C(K,T)}{\partial K} = -e^{-rT}\mathbb{P}^Q(S_T > K)$$

where $P^Q(S_T > K)$ is the risk-neutral probability that the asset price exceeds the strike.

The second derivative:

$$\frac{\partial^2 C(K,T)}{\partial K^2} = e^{-rT} f_Q(K,T)$$

connects the option price's curvature to the RND.

$\frac{\partial^2 C}{\partial K^2} > 0$ - implies a positive density for strikes within the market range, which means it assumes the no-arbitrage condition. The density $f_Q(K,T)$ shows the chance of the asset price reaching K under the risk-neutral measure.

4.6 Risk-neutral density under lognormal assumption

In the Black-Scholes (BS) model, which is the benchmark model for option pricing , including Bitcoin options, the price of an asset follows a geometric Brownian motion:

$$S_T = S_0 \exp\left(\left(r - \frac{\sigma^2}{2}\right)T + \sigma W_T\right)$$

where: S_T - Price of the asset at maturity; S_0 - Current price of the asset; r - Risk-free rate; σ - Volatility of the asset's returns; T - Time to maturity; W_T - A standard Brownian motion. This implies that S_T is **lognormally distributed**:

$$\ln\left(\frac{S_T}{S_0}\right) \sim \mathcal{N}\left(\left(r - \frac{\sigma^2}{2}\right)T, \sigma^2 T\right)$$

Black-Scholes Formula for a European Call Option is shown below:

$$C = S_0\Phi(d_1) - Ke^{-rT}\Phi(d_2)$$

where:

$$d_1 = \frac{\ln(S_0/K) + \left(r + \frac{\sigma^2}{2}\right) T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

Then, finding the constant volatility (σ) that minimizes the relative squared error between market and model price was done. The same calibration method as with NIG approach was compiled. My approach directly calibrates volatility using Black-Scholes prices and uses the lognormal RND formula. Using Haug's (2007) method of differentiating, we get:

$$\frac{n(d_2)e^{-rT}}{X\sigma\sqrt{T}}$$

where X: strike; sigma: implied volatility; T: years to maturity; r: interest rate; S: underlying futures price. And d_2 is the

probability that the option will expire in the money. And $n(d_2)$ gives us the expected value of paying at the strike price.

4.7 Risk-neutral density under Normal Inverse Gaussian

The NIG RND is derived using the subordinated Lévy process framework proposed by Cont and Tankov (2003). The NIG RND $f_{\text{NIG}}(x)$ (PDF) is obtained by applying the inverse Fourier transform to the CF:

$$f_{\text{NIG}}(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iux} CF_{\text{NIG}}(u)\, du$$

where $CF_{\text{NIG}}(u)$ - Characteristic function evaluated at u.

Further, risk-neutral density, denoted f(x), represents the probability density of prices under the risk-neutrality measure. In the NIG space, f(x) is formulated as:

$$f(x) = C \cdot \exp(A \cdot x) \cdot \frac{\text{Bessel}_1(B \cdot D)}{D}$$

where the elements C, A, B, and D were derived from the NIG parameters (κ, θ, σ, μ) and the time to maturity (t). The scaling constant "C" assures normalization and includes time decay and subordination, defined as:

$$C = \frac{t}{\pi} \exp\left(\frac{t}{\kappa}\right) \sqrt{\frac{\theta^2}{\kappa\sigma^2} + \frac{1}{\kappa}}$$

The exponential term (Ax), where A is: $A = \frac{\theta}{\sigma^2}$ adapts for skewness, which is written as drift parameter θ, while the modified Bessel function of the second kind, $Bessel_1$, gives flexibility to the density's shape, adjusting both skewness and kurtosis. The arguments B and D further refine the density's shape, with $B = \frac{\sqrt{\theta^2 + \frac{\sigma^2}{\kappa}}}{\sigma^2}$ meaning scaling and $D = \sqrt{x^2 + \frac{t^2\sigma^2}{\kappa}}$ operating as a distance metric.

5 Results

:

The mean calibration error for LN was found to be around 2.1 basis points (bps), which is more than double the errors marked in the other two different models. LN's pricing inaccuracies were mainly related to its strict assumption of constant volatility, leading to high probabilities assigned to average returns and notable underestimation of tail probabilities. This point was noticeable in extreme market conditions, where LN failed to capture the asymmetries in Bitcoin's volatility smile.

Even if we cannot call BS as a model, it was added as comparable for calibration and pricing results. The BS achieved almost perfect calibration error, providing robust alignment with market prices. It did well in catching central returns and outperformed NIG in the at-the-money region, offering a precise

fit due to its non-parametric structure. While BS showed moderate precision in analysis, its dependence on market-implied volatilities limits its benefit for extrapolating beyond observed data.

The Normal Inverse Gaussian (NIG) model demonstrated excellent flexibility and precision, reaching a low mean calibration error of 0.7 bps. Its parametric construction qualified it to catch the skewness and kurtosis in Bitcoin's market dynamics,makingit effective in modeling extreme events like crises or rapid downfalls. NIG's capability to align with market-implied probabilities in the tails is noticeable in the density plots. However, NIG slightly underperformed in central regions, allocating less density to average returns compared to BS. This tradeoff recalls its priority on tail modeling, which is favorable for

risk management applications but slightly decreases its fit in non-extreme scenarios.

The comparison of RND densities delivered further understanding of the models' attributes. The densities of BS and NIG were similar across most cases, with deviations of less than 0.02 density units in the tails and even more minor discrepancies near the center. Similarity shows the robustness of NIG model, which is, for a moment, parametric, in capturing market-implied distributions. However, the LN model's density deviated significantly, showing some overestimating probabilities around average returns and underestimating tail probabilities compared to NIG and BS. This difference shows LN's inability to reflect the asymmetries and fat tails typical of Bitcoin's price movements. Last statistical testing reinforced these findings. Likelihood ratio tests (LRTs) conducted with

weights applied to different regions of returns demonstrated that NIG dominated in tail regions, capturing extreme probabilities. The LN model was better than NIG in average return ranges because it assumed constant volatility. After scaling densities to ensure total probability equals one, the results revealed minimal changes, further validating their robustness. LN, however, continued to display significant misalignments, particularly in the tails. The tail-weighted LRT tests highlighted NIG's slight outperformance in modeling extreme price movements, making it a good choice for tail risk analysis in Bitcoin markets. The statistical similarity between BS and NIG, despite their methodological differences, is noteworthy. While BS directly interpolates market prices, NIG's parametric design allows for extrapolation, suggesting theoretical advantages for scenarios concerning unobserved market dy-

namics. It is crucial to mention that LRT tests were using a 5% level. All results can be explored in Figure 4-9 and Table 3.

6 Conclusion

This research indicates the effectiveness of the Normal Inverse Gaussian (NIG) distribution in forecasting risk-neutral densities in Bitcoin options, delivering a robust alternative to the traditional log-normal model. The NIG model's ability to contain skewness and kurtosis demonstrates pronounced asymmetries and fat tails inherent in Bitcoin's price dynamics. And through investigated calibration and out-of-sample testing, my findings show that the NIG model delivers prime tail risk predictions, which makes it a solid tool for risk managers and institutional investors steering volatile markets. Even if the log-normal model remains a computationally efficient benchmark, its limitations in capturing higher-order moments highlight the

requirement for more parametric approaches like NIG. My results underscore the growing relevance of forward-looking analytics in cryptocurrency markets and allow further exploration into advanced modeling techniques to enhance predictive accuracy and decision-making in digital asset trading.

References:

1. Abadir, K. M. & Rockinger, M. (2003). Density functionals, with an option-pricing application. Econometric Theory, 19(5), 778–811.

2. Abramowitz, M. (1974). Handbook of mathematical functions: With formulas, graphs, and mathematical tables. Dover Publications.

3. Aït-Sahalia, Y., & Duarte, J. (2003). Nonparametric option pricing under shape restrictions. Journal of Econometrics, 116(1–2), 9–47.

4. Aït-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. The Journal of Finance, 53(2), 499–547.

5. Alexander, C., Deng, J., Feng, J., & Wan, H. (2023). Net buying pressure and the information in bitcoin option trades. Journal of Financial Markets, 63(C).

6. Amisano, G., & Giacomini, R. (2007). Comparing density forecasts via weighted likelihood ratio tests. Journal of Business & Economic Statistics, 25(2), 177–190.

7. Ardia, D., Bluteau, K., & Rüede, M. (2019). Regime changes in Bitcoin GARCH volatility dynamics. Finance Research Letters, 29(C), 266–271.

8. Baek, C., & Elbeck, M. (2015). Bitcoins as an investment or speculative vehicle? A first look. Applied Economics Letters, 22, 30–34.

9. Banz, R. W., & Miller, M. H. (1978). Prices for state-contingent claims: Some estimates and applications. The Journal of Business, 51(4), 653–672.

10. Birru, J., & Figlewski, S. (2012). Anatomy of a meltdown: The risk-neutral density for the S&P 500 in the fall of 2008. Journal of Financial Markets, 15(2), 151–180.

11. Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3), 637–654.

12. Bliss, R., & Panigirtzoglou, N. (2004). Option implied risk aversion estimates. The Journal of Finance, 59(2), 407–446.

13. Bondarenko, O. (2003). Statistical arbitrage and securities prices. The Review of Financial Studies, 16(3), 875–919.

14. Bouoiyour, J., & Selmi, R. (2015). Bitcoin price: Is it really that new round of volatility can be on way? MPRA Paper 65580. University Library of Munich, Germany.

15. Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. The Journal of Business, 51(4), 621–651.

16. Brennan, M. J. (1979). The pricing of contingent claims in discrete time models. The Journal of Finance, 34(1), 53–68.

17. Bu, R., & Hadri, K. (2007). Estimating option-implied risk-neutral densities using spline and hypergeometric functions. The Econometrics Journal, 10(1), 216–244.

18. Chaum, D. (1983). Blind signatures for untraceable payments. In D. Chaum, R. Rivest, & A. Sherman (Eds.), Advances in Cryptology Proceedings of Crypto 82 (pp. 199–203).

19. Cont, R., & Tankov, P. (2003). Financial modeling with jump processes. Chapman and Hall/CRC.

20. Cox, J. C., & Ross, S. A. (1976). The valuation of options for alternative stochastic processes. Journal of Financial Economics, 3(1–2), 145–166.

21. Dwyer, G. P. (2015). The economics of Bitcoin and similar private digital currencies. Journal of Financial Stability, 17(C), 81–91.

22. Eriksson, A., Forsberg, L., & Ghysels, E. (2004). Approximating the probability distribution of functions of random variables: A new approach. Discussion Paper CIRANO.

23. Eriksson, A., Ghysels, E., & Wang, F. (2009). The Normal Inverse Gaussian distribution and the pricing of derivatives. The Journal of Derivatives, 16, 23–37.

24. Figlewski, S. (2008). Estimating the implied risk-neutral density for the U.S. market portfolio. In T. Bollerslev, J. R. Russell, & M. Watson (Eds.), Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle. Oxford University Press.

25. Geman, H., Madan, D. B., & Yor, M. (2001). Asset prices are Brownian motion: Only in business time. In Quantitative Analysis in Financial Markets: Collected Papers of the New York University Mathematical Finance Seminar (Volume II) (pp. 103–146).

26. Haug, E. G. (2007). The complete guide to option pricing formulas. McGraw-Hill.

27. Härdle, W., & Hlávka, Z. (2009). Dynamics of state price densities. Journal of Econometrics, 150(1), 1–15.

28. Jackwerth, J. C., & Rubinstein, M. (1996). Recovering probability distributions from option prices. The Journal of Finance, 51(5), 1611–1631.

29. Longstaff, F. A. (1995). Option pricing and the martingale restriction. Review of Financial Studies, 8(4), 1091–1124.

30. Lucic, V., & Sepp, A. (2024). Valuation and hedging of cryptocurrency inverse options. Quantitative Finance, 24(7), 851–869.

31. Madan, D. B., & Milne, F. (1994). Contingent claims valued and hedged by pricing and investing in a basis. Mathematical Finance, 4(3), 223–245.

32. Malwa, S. (2024, August 5). Crypto futures record $1B$ in liquidations as Bitcoin nosedives, Ether slumps most since 2021. CoinDesk.

33. Malz, A. (2014). A simple and reliable way to compute option-based risk-neutral distributions. Technical Report 677, FRB of New York Staff Report.

34. McGee, R., Post, T., & Potì, V. (2024). Option-implied physical distributions. Michael J. Brennan Irish Finance Working Paper Series Research Paper No. 24-6.

35. Melick, W. R., & Thomas, C. P. (1997). Recovering an asset's implied PDF from option prices: An application to crude oil during the Gulf Crisis. The Journal of Financial and Quantitative Analysis, 32(1), 91–115.

36. Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Decentralized Business Review, 21260.

37. Pagan, A., & Ullah, A. (1999). Nonparametric econometrics. Cambridge University Press.

38. Pichl, L., & Kaizoji, T. (2017). Volatility analysis of Bitcoin. Quantitative Finance and Economics, 1(4), 474–485.

39. Rubinstein, M. (1976). The valuation of uncertain income streams and the pricing of options. Bell Journal of Economics, 7(2), 407–425.

40. Rubinstein, M. (1994). Implied binomial trees. The Journal of Finance, 49(3), 771–818.

41. Scaillet, O., Treccani, A., & Trevisan, C. (2017). High-frequency jump analysis of the Bitcoin market. International Finance eJournal.

42. Schmelzle, M. (2010). Option pricing formulae using Fourier transform: Theory and application.

43. Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. Econometrica, 57(2), 307–333.

44. Yuan, M. (2009). State price density estimation via non-parametric mixtures. Annals of Applied Statistics, 3(3), 963–984.

7 Appendix

Statistic	Value
Mean	0.01463
Standard deviation	0.09202
Skewness	0.00071
Kurtosis	3.06508

Table 1: Descriptive statistics on weekly returns.

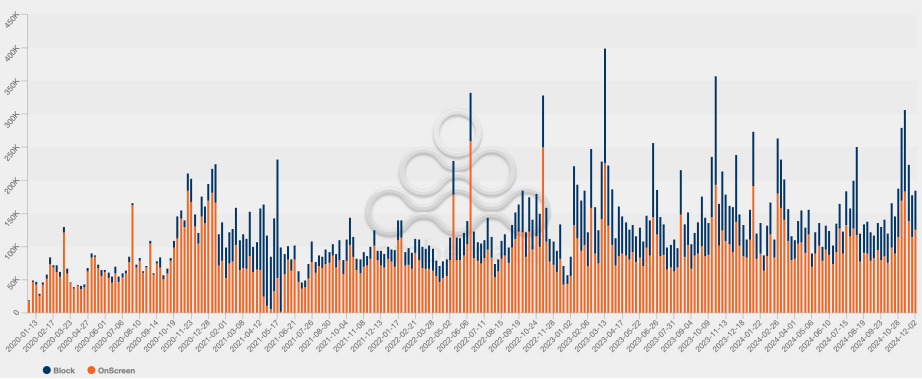

Fig. 1: Weekly Volume of Bitcoin Option Contracts

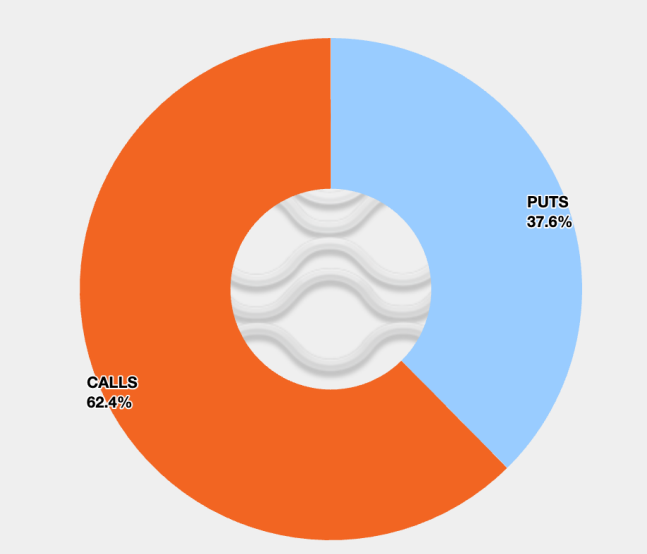

Fig. 2: Puts vs Calls Volume

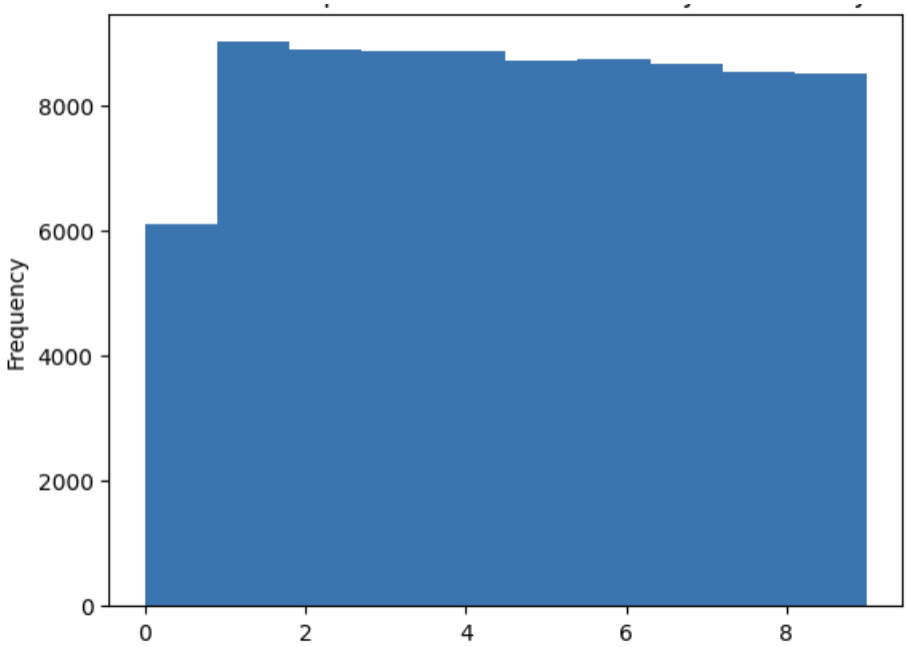

Fig. 3: Data with options with less than 10 days to maturity

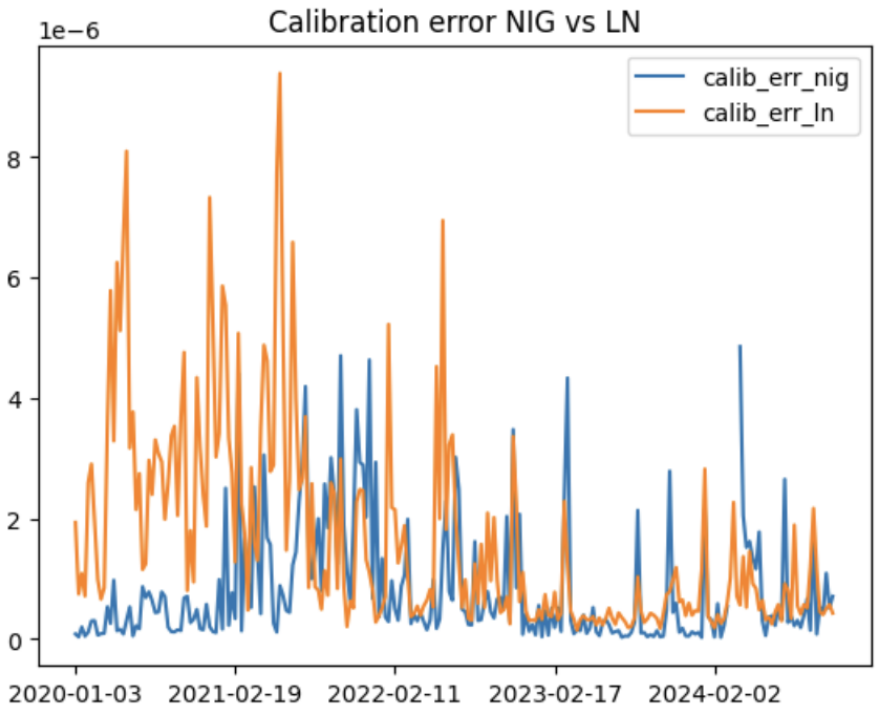

Fig. 4: Callibration error between NIG and LN

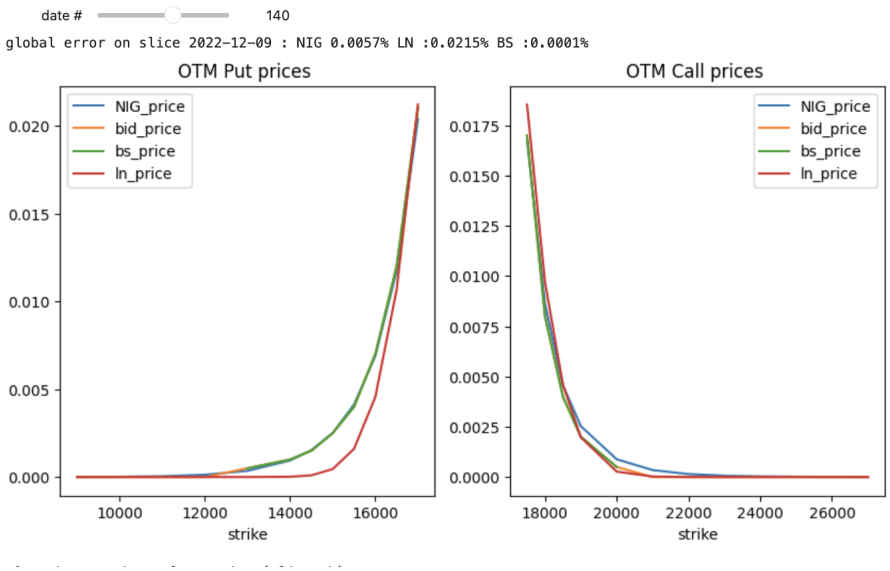

Fig. 5: Example of calibration errors showing on OTM calls and puts

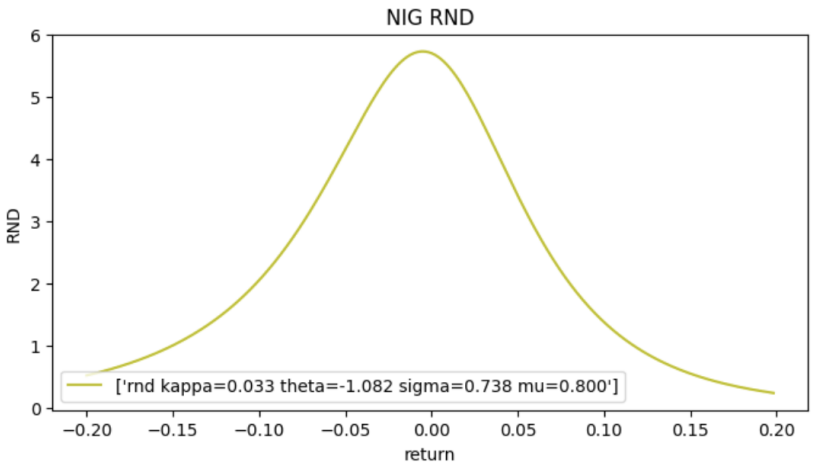

Fig. 6: One normal inverse Gaussian risk-neutral density

Days to maturity I	Start date I	Maturity date
0	114	114
1	235	235
2	236	236
3	236	236
4	237	237
5	237	237
6	237	237
7	237	237
8	237	237
9	237	237
10	236	236
11	236	236

Table 2: Counting of start dates according to days to maturity

Test Stat	Result	Conclusion
NIG_vs_LN	0.416814	NIG and LN are performing equally at 0.05% level
NIG_vs_LN_Skewright	-0.549505	NIG and LN are performing equally at 0.05% level
NIG_vs_LN_Skewleft	0.696314	NIG and LN are performing equally at 0.05% level
NIG_vs_LN_Center	-5.221990	LN is significantly better than NIG at 0.05% level
NIG_vs_LN_Tails	1.506117	NIG and LN are performing equally at 0.05% level

Table 3: LR Test results comparing NIG and LN models.

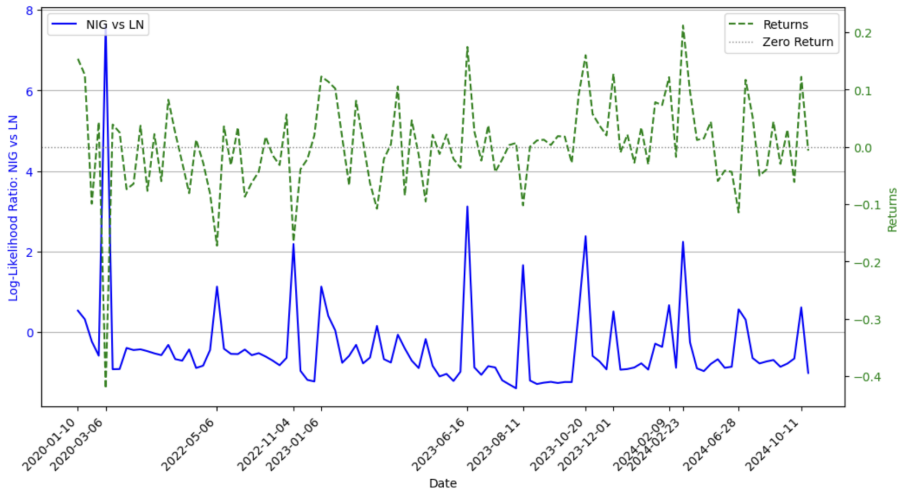

Fig. 7: Log-Likelihood Ratio (NIG vs LN) and Returns Over Time

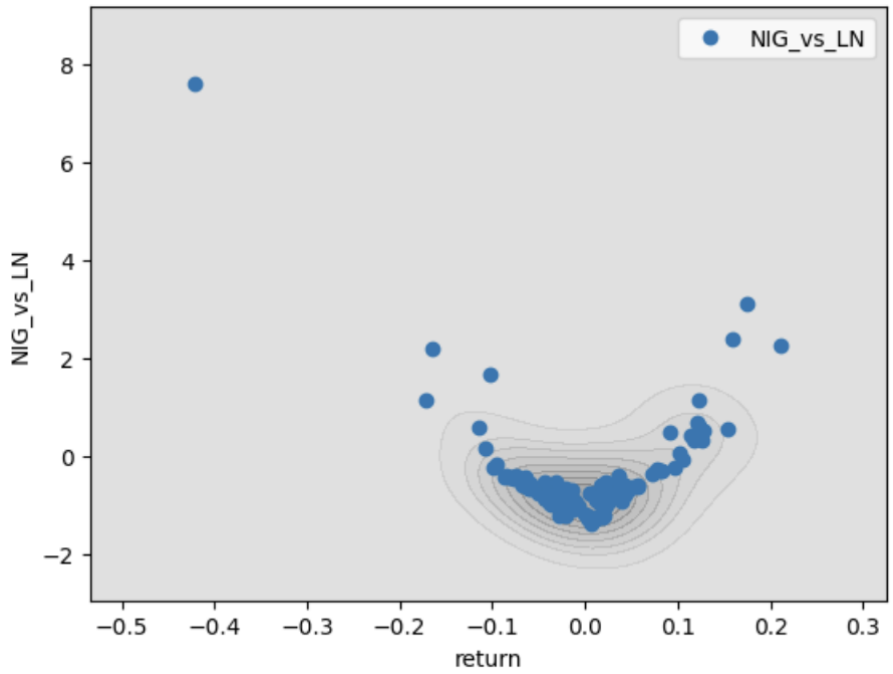

Fig. 8: Comparison charts of log-likelihood score between NIG and LN

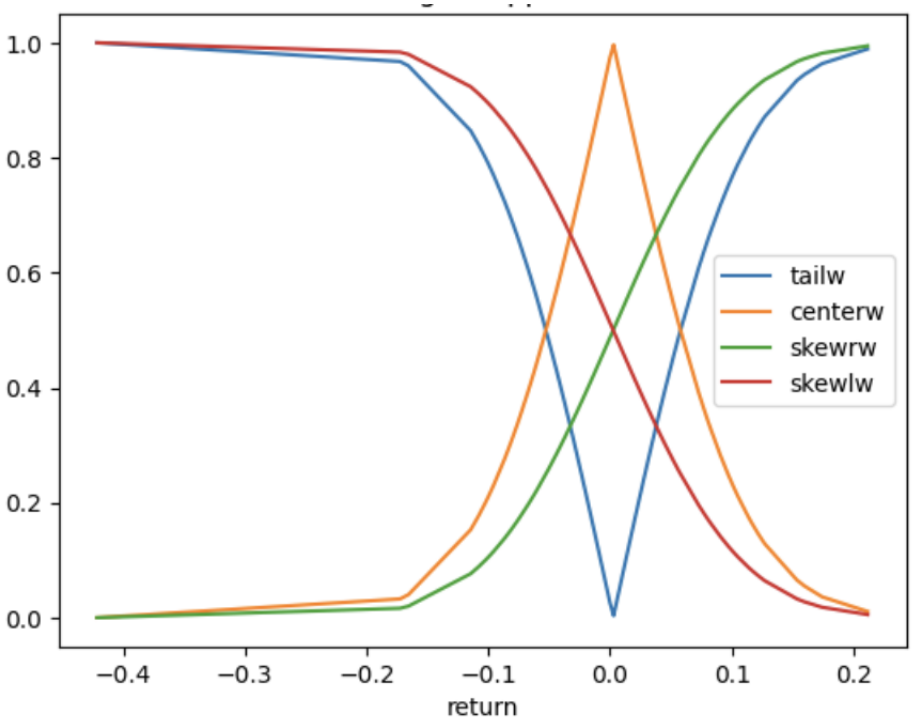

Fig. 9: Applied weights in likelihood ratio test