Examining Historical Returns and Volatility in Bitcoin

Determining some best practices for forecasting models

Introduction

A couple of weeks ago, we went on a journey of exploration into the world of real estate markets. Our mission? To not only understand how things have performed in the past, but also, maybe even more importantly, to figure out a way to predict future prices. This led us down an unexpected rabbit hole — forecasting Bitcoin's returns. But it was worth it. We came up with some really fascinating conclusions, which kickstarted a deep-dive investigation into Bitcoin's historical returns and the best ways to predict future price trends.

The silver lining is that we've reached some intriguing conclusions. Thus, the research before the research was born - a comprehensive examination of Bitcoin's historical returns, their distribution, and the potentially optimal methods for forecasting price paths.

Grasping the distribution of past returns is a pivotal step in the creation of any asset simulation model. This understanding enables us to perform more precise simulations of future asset price behavior. For assets like Bitcoin and other financial instruments, which don't seem to abide by the normal distribution, extra scrutiny is necessary to pinpoint the most fitting statistical distribution that accurately mirrors its historical returns.

Bitcoin Historical Return Distribution

Start Date	2010-07-17
End Date	2023-05-29
Rolling Periods	1 day(s)
Mean (n = 1)	0.28%
Mean Annualized	179.51%
Geometric Mean (n = 1)	0.28%
Geometric Mean Annualized	179.92%
Volatility (n = 1)	5.99%
Volatility Annualized	114.19%
Skewness (Symmetry) (close to 0 = symmetric)	2.8101 Highly Positively Skewed
Kurtosis (Tails) (close to 3 = bell-shaped)	104.4469 Heavy Tails (Leptokurtic)

In this research piece, we delve into this subject with an objective to establish a robust methodology for executing Monte Carlo simulations for Bitcoin, and potentially, for other similar assets. We will scrutinize historical price data and assess various statistical distributions to identify the one that offers the best fit.

As always, open source code and a functioning web app will be available.

Historical Bitcoin Returns

We kicked off our analysis by looking at Bitcoin prices going back to July, 2010. That's over ten years of data and includes almost three Bitcoin halvings. It's important to cast a wide net because even though recent data might be more relevant as Bitcoin becomes more mainstream, the impact of halving cycles is still significant. Whether halvings are factored into prices is a topic of ongoing debate (maybe we'll tackle that in another post soon). We'll set that aside for now, but it's clear that Bitcoin has a history of moving in roughly four-year cycles, either by design or by coincidence. A shorter data set might have overlooked this pattern.

It doesn't take long to figure out that Bitcoin's returns don't follow a normal distribution, a bell curve where most of the data huddles around the average, with less probable outcomes trailing off on either side.

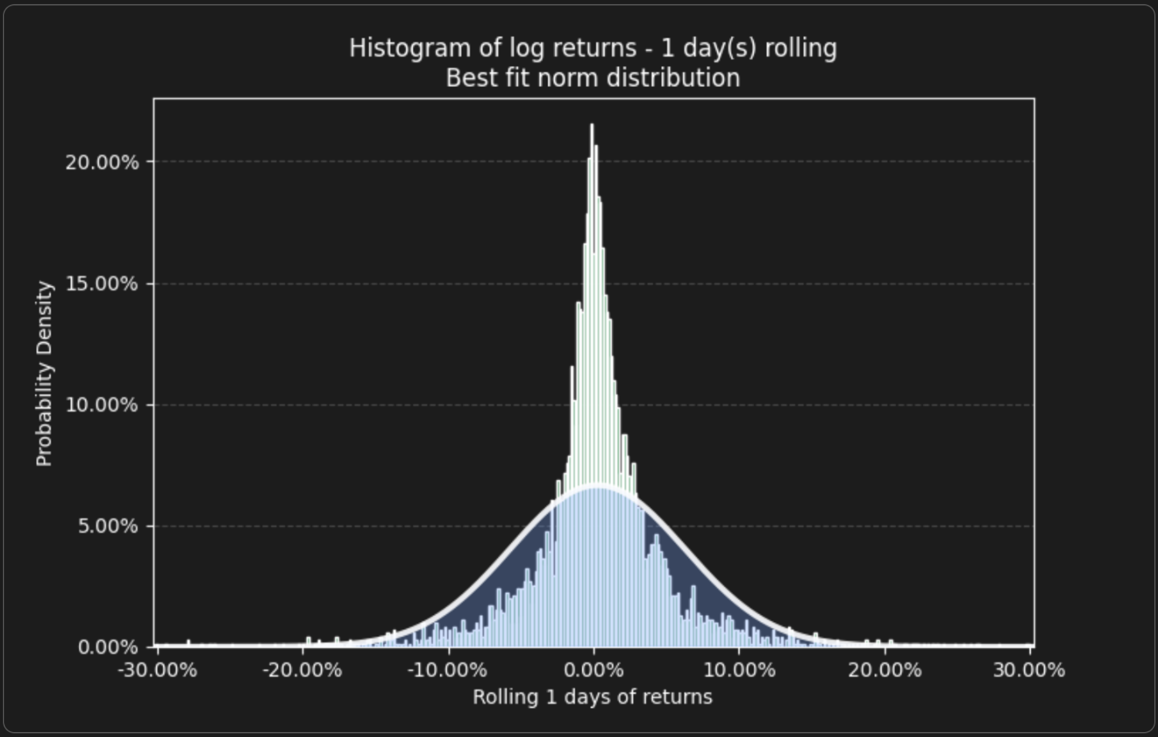

To get a better handle on Bitcoin's return patterns, we looked at two measures called skewness and kurtosis. Skewness tells us how symmetrical a distribution is, while kurtosis measures how heavy the tails are.

In the case of Bitcoin's returns, the skewness value of +2.8, which deviates significantly from zero, indicates that the distribution is not symmetrical. The kurtosis was around 105, way higher than the 3 you'd expect in a normal distribution. This means that Bitcoin's price changes can have extreme values and heavy tails compared to a normal distribution.

In comparison, the S&P 500 has a long term distribution of returns with skewness of ~-0.3 and kurtosis of ~10 - in other words, a distribution that is more symmetrical and with less extreme outcomes compared to Bitcoin.

Let's explore other distribution models.

Distribution Comparison and Goodness of Fit Tests

Next up, we compared Bitcoin's *1-day* returns against several theoretical distributions. We started with the Kolmogorov-Smirnov test, a nonparametric test that checks how well a sample distribution fits a reference probability distribution. Following a suggestion from Nick (@btconometrics), we also checked the distribution fit using Akaike information criteria or AICs.

KS Tests results **(Lower KS-Stat Numbers = Better)**

Distribution		KS Statistic	P-Value	Status
Johnson SU Distribution	(johnsonsu)	0.0138	0.3354	Cannot Reject
Generalized Hyperbolic Distribution	(genhyperbolic)	0.0154	0.2151	Cannot Reject
Normal Inverse Gaussian Distribution	(norminvgauss)	0.0160	0.1830	Cannot Reject
Generalized Normal Distribution	(gennorm)	0.0200	0.0479	Reject
Student's t-distribution	(t)	0.0238	0.0101	Reject
Double Gamma Distribution	(dgamma)	0.0411	0.0000	Reject
Double Weibull Distribution	(dweibull)	0.0490	0.0000	Reject
Log-Laplace Distribution	(loglaplace)	0.0555	0.0000	Reject
Laplace Distribution	(laplace)	0.0560	0.0000	Reject

AIC Tests results **(Lower AIC Numbers = Better)**

Distribution		AIC
Generalized Hyperbolic Distribution	(genhyperbolic)	-16,830
Generalized Normal Distribution	(gennorm)	-16,790
Student's t-distribution	(t)	-16,775
Double Weibull Distribution	(dweibull)	-16,654
Double Gamma Distribution	(dgamma)	-16,532
Log-Laplace Distribution	(loglaplace)	-16,370
Laplace Distribution	(laplace)	-16,366

The Kolmogorov-Smirnov test tells us how well our data fits a specific theoretical distribution by measuring the biggest distance between the cumulative distribution function of the sample and the reference distribution.

The Akaike Information Criterion (AIC) is a popular way of measuring statistical models. It basically combines the fit and simplicity of the model into one statistic. When comparing two models, the one with the lower AIC is usually considered "better".

After examining the results and running the data for different periods, we found that a generalized hyperbolic distribution looks like a good fit.

In simple terms, a generalized hyperbolic distribution is a flexible model that can handle skewness and kurtosis, characteristics commonly observed in financial returns like Bitcoin. When compared with the normal distribution, it's clear that the generalized hyperbolic curve fits much better.

Comparing the chart below with the normal distribution makes it clear that the generalized hyperbolic curve has significantly better fit.

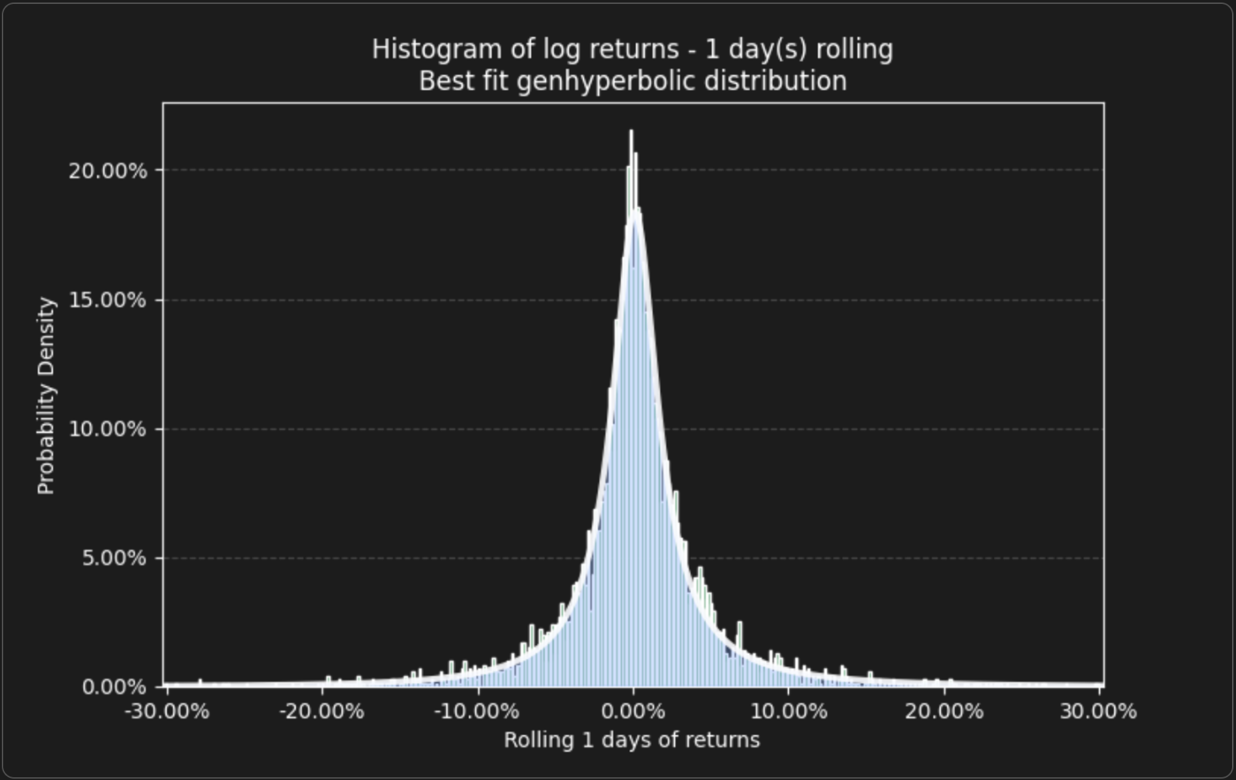

So, why are we looking for a good distribution? These distributions are continuous, so we can use them to find any specific probability (or return). When we run Monte Carlo simulations, we pick a random number and find its corresponding return in the distribution. This brings us to another important aspect of modeling results: significant price changes are often followed by even more significant changes, a phenomenon known as "volatility clustering".

GARCH Model

Another topic to explore is whether Bitcoin's returns are autocorrelated. Autocorrelation means that Bitcoin's returns are not independent and that past returns can affect future returns. Unfortunately this is ignored by many models out there but that's also a topic for another research piece.

Considering the shortcomings of basic statistical distributions in capturing the characteristics of Bitcoin returns, we turned to the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model. GARCH models are useful because they assume that the variance of the residuals follows an autoregressive process, which means they can capture volatility clustering.

The results showed strong evidence of autocorrelation in Bitcoin returns, which isn't unusual in financial time series data and confirms that GARCH models are suitable for modeling volatility.

We won't go into the full details of the GARCH model results here, but you can easily modify and access them in the open-source code repository linked to this document. In short, we used an Ljung-box test to check for autocorrelation at a given number of lags. The results are pretty convincing that returns are autocorrelated. You can check out the open-source code for the full calculations.

In summary, these results strongly support the idea that Bitcoin's returns are autocorrelated. The statistically significant ARCH and GARCH terms show that past returns and their volatility significantly affect current volatility, which clearly indicates autocorrelation.

Backtesting and Monte Carlo Simulations

Now for the fun part.

Backtesting and Monte Carlo Simulations are important steps in validating the robustness and predictive power of any statistical model. The idea is to test how the model performs against actual historical data (backtesting) and to see how it behaves under a wide variety of random scenarios (Monte Carlo Simulations).

We tested two forecasting methods: a generalized hyperbolic distribution and a historical simulation approach. The historical simulation approach is just a fancy way of saying that we're assuming future returns will follow a similar pattern to past returns.

To measure the accuracy of these methods, we used the Root Mean Square Error (RMSE) method. RMSE measures the difference between the values predicted by a model and the actual values. You can think of it as the "average mistake" made by the model. The smaller the RMSE, the better the model is at predicting.

In our backtesting, the generalized hyperbolic distribution method had a lower RMSE than the historical simulation method, which means it made smaller "mistakes" on average and is therefore better at predicting future returns.

So, the generalized hyperbolic distribution method seems promising for simulating Bitcoin returns, providing a superior fit and predictive accuracy. Along with this, the GARCH model proved useful for modeling volatility, especially given Bitcoin's tendency for volatility clustering.

To wrap up, for a comprehensive simulation of Bitcoin prices, we recommend using both the generalized hyperbolic distribution to model returns and the GARCH model to simulate volatility. This combo should give a more realistic picture of Bitcoin's price movements. The example below is a Monte Carlo simulation of a forward-looking forecast of Bitcoin prices assuming a generalized hyperbolic distribution. The orange line is the average price across all simulations.

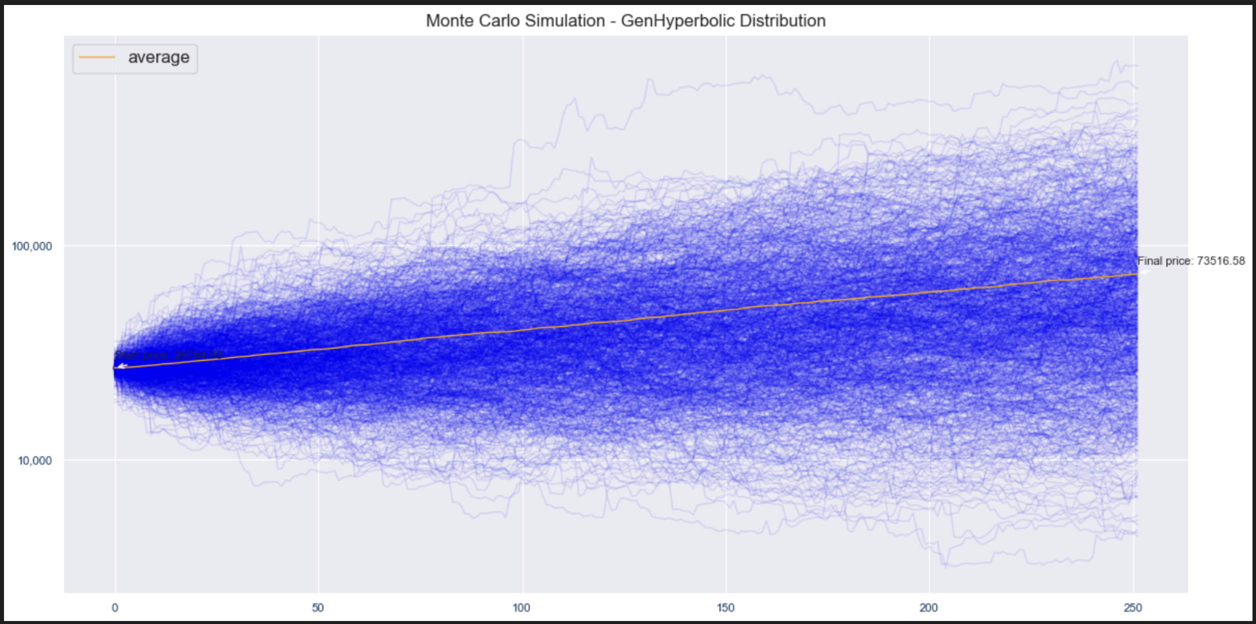

In our work at Nakamoto Portfolio, we're comfortable using this distribution for our forecasts and it will be our go-to choice moving forward. We will have a few apps coming soon that will let the user explore with different Monte Carlo forecasts.

It's important to remember that no statistical model can guarantee future outcomes. Bitcoin is complex and unpredictable, so any predictions should be taken with a few grains of salt. We can say with certainty that our forecasts will be wrong. Instead of trying to pinpoint exact future prices, we suggest focusing on forecasting to understand potential price movements and their potential impacts. For example, what happens to your portfolio when Bitcoin experiences wild volatility? Can you stomach the changes? Forecasting and looking at extreme scenarios and understanding potential drawdowns can be beneficial. Understanding drawdowns can guide risk management strategies.

Special thanks to Nick, Breno Brito, Ed Gotham and Alex Krüger for the suggestions and review of this piece. I never cease to be impressed by the quality of bitcoiners all around the world!

Bitcoiners Run the Numbers
The code used for this article is open-source and available here.
Clone, Fork, modify, verify.