FinTech

Article

Multiscale Stochastic Models for Bitcoin: Fractional Brownian Motion and Duration-Based Approaches

Arthur Rodrigues Pereira de Carvalho [1], Felipe Quintino [1], Helton Saulo [1], Luan C. S. M. Ozelim [2,*], Tiago A. da Fonseca [3] and Pushpa N. Rathie [1]

[1] Department of Statistics, University of Brasília, Brasília 70910-900, DF, Brazil; carvalho.arthur@aluno.unb.br (A.R.P.d.C.); felipe.quintino@unb.br (F.Q.); heltonsaulo@unb.br (H.S.); pushpanrathie@yahoo.com (P.N.R.)
[2] Department of Civil and Environmental Engineering, University of Brasília, Brasília 70910-900, DF, Brazil
[3] Gama Engineering College, University of Brasília, Brasília 72444-240, DF, Brazil; fonsecafga@unb.br
* Correspondence: luanoz@gmail.com

Abstract

This study introduces and evaluates stochastic models to describe Bitcoin price dynamics at different time scales, using daily data from January 2019 to December 2024 and intraday data from 20 January 2025. In the daily analysis, models based on are introduced to capture long memory, paired with both constant-volatility (CONST) and stochastic-volatility specifications via the Cox–Ingersoll–Ross (CIR) process. The novel family of models is based on Generalized Ornstein–Uhlenbeck processes with a fluctuating exponential trend (GOU-FE), which are modified to account for multiplicative fBm noise. Traditional Geometric Brownian Motion processes (GFBM) with either constant or stochastic volatilities are employed as benchmarks for comparative analysis, bringing the total number of evaluated models to four: GFBM-CONST, GFBM-CIR, GOUFE-CONST, and GOUFE-CIR models. Estimation by numerical optimization and evaluation through error metrics, information criteria (AIC, BIC, and EDC), and 95% Expected Shortfall (ES_{95}) indicated better fit for the stochastic-volatility models (GOUFE-CIR and GFBM-CIR) and the lowest tail-risk for GOUFE-CIR, although residual analysis revealed heteroscedasticity and non-normality. For intraday data, Exponential, Weibull, and Generalized Gamma Autoregressive Conditional Duration (ACD) models, with adjustments for intraday patterns, were applied to model the time between transactions. Results showed that the ACD models effectively capture duration clustering, with the Generalized Gamma version exhibiting superior fit according to the Cox–Snell residual-based analysis and other metrics (AIC, BIC, and mean-squared error). Overall, this work advances the modeling of Bitcoin prices by rigorously applying and comparing stochastic frameworks across temporal scales, highlighting the critical roles of long memory, stochastic volatility, and intraday dynamics in understanding the behavior of this digital asset.

Keywords: stochastic differential equations; generalized Langevin equation; fractional Brownian motion; ACD models

JEL Classification: C02; C13; C14; C15; C41; C52

check for updates

Academic Editor: David Roubaud

Received: 13 July 2025
Revised: 5 September 2025
Accepted: 14 September 2025
Published: 19 September 2025

Citation: Carvalho, A.R.P.d.; Quintino, F.; Saulo, H.; Ozelim, L.C.S.M.; Fonseca, T.A.d.; Rathie, P.N. Multiscale Stochastic Models for Bitcoin: Fractional Brownian Motion and Duration-Based Approaches. *FinTech* **2025**, *4*, 51. https://doi.org/10.3390/fintech4030051

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

In recent years, Bitcoin has emerged as one of the most prominent financial assets. This digital currency was introduced in 2009 by Satoshi Nakamoto [1] and is based on

blockchain technology—a digital ledger in which all transactions are securely and transparently recorded. Transactions are carried out in a decentralized manner, meaning they do not rely on intermediaries such as banks or governments. Instead, value is transferred directly between users via digital wallets. These transactions must be verified and recorded on the blockchain through a process known as mining, in which mathematical computations are performed to validate transactions and sequentially add new blocks to the blockchain, thereby ensuring data security and integrity [2]. The literature reveals that there is a growing interest in Bitcoin-related topics, with high quality research [3] being produced on diverse fields, such as safe haven, internet of things (IoT), proof of work (PoW), market efficiency, sentiment analysis, digital currency, and privacy [4].

Unlike traditional currencies, Bitcoin operates continuously, 24 h a day, 7 days a week, and has a maximum supply capped at 21 million units in circulation. This structure results in high liquidity, volatility, and scarcity in the market. In this context, its limited supply and the absence of a regulatory authority contribute to its high volatility, making it an asset that draws the attention of both investors and researchers [5].

Thus, Bitcoin represents a currency that transcends traditional financial markets, while still being influenced by global events such as economic crises and regulatory changes. In some cases, it is even employed as a hedge against inflation in countries with unstable economies. Studies have shown that the inherent unpredictability of Bitcoin prices exhibits a complex, recurring pattern over time, indicating a phenomenon known as long memory, meaning that current prices may be influenced by past market movements [6].

A deep understanding of Bitcoin's volatility is essential across multiple areas of the financial sector, from asset pricing and portfolio risk assessment to the formulation of robust investment strategies. However, forecasting risk measures for cryptocurrencies—such as Value-at-Risk (VaR) and Expected Shortfall (ES)—poses significant challenges due to their inherent characteristics, including high volatility, extreme price movements, and periods of intense market turbulence [7]. To address these challenges, statistical tools, such as heavy-tailed distributions [8,9], have been applied. In particular, special attention has been devoted to models based on fractional Brownian motion (fBm) [6,10], aiming to elucidate the underlying dynamics of Bitcoin.

Moreover, the analysis of high-frequency financial data has proven essential for understanding market dynamics and formulating more effective investment strategies. In this regard, autoregressive conditional duration (ACD) models, as proposed by [11], offer a robust framework for modeling the time between transactions, particularly in irregular datasets such as those found in cryptocurrencies. Although widely used in financial markets to assess transaction intensity and duration dynamics, these models remain underexplored in the context of Bitcoin, despite their relevance for analyzing price risk and market microstructure [12].

The goal of this paper is to propose a new class of processes driven by fBm, capable of modeling the daily prices of Bitcoin and its volatility, and to apply ACD models to analyze the price durations. Aiming to identify temporal and seasonal patterns that contribute to Bitcoin's market behavior, the following specific objectives are addressed:

- Identify and describe the long-memory patterns present in Bitcoin's daily price history.
- Incorporate stochastic volatility into the models proposed by [13], analyzing its impact on forecasting accuracy.
- Conduct a comparative analysis of different stochastic price models.
- Perform an intraday analysis of Bitcoin price durations using the ACD model.

The rest of this paper is organized as follows. In Section 2, we present stochastic differential equations and Generalized Ornstein–Uhlenbeck processes with a fluctuating exponential trend (GOU-FE processes) driven by fBm. In Section 3, we discuss ACD models

and their variants as a tool to allow for intraday price analysis. In Section 4, we discuss a Bitcoin price modeling application, and we present the concluding remarks in Section 5.

2. The GOU-FE Process Driven by fBm

Models based on stochastic differential equations (SDEs) have been widely used to analyze financial data (see, for example [14,15], and the references therein). More recently, considerable interest has arisen in models with memory effects, where non-Markovian drift plays a significant role. An important example of SDE is the generalized Langevin equation (GLE), proposed by [16,17]. In this SDE, the drift component considers the evolution of the process up to a considered time.

Results on the existence and uniqueness of solutions for certain classes of the GLE were obtained by [18–20]. Over the last two decades, numerous authors have further explored theoretical aspects and applications of the GLE, including [21–26].

Issues concerning statistical estimation for the drift term of the GLE are important to better fit the mathematical model to the potential real data aimed to be modeled. They allow us to make predictions about future events within the modeled time series, using accumulated information over time. In this sense, ref. [13] studied drift estimation for the three-parameter Generalized Ornstein–Uhlenbeck process with a fluctuating exponential trend (GOU-FE), which is a solution process of the following class of the GLE:

$$dX(t) = \left(-\theta_1(1-\theta_3)X(t) - \int_0^t X(s)\Gamma_\theta(t-s)ds\right)dt + dN(t), \quad t > 0, \tag{1}$$

where $X(0) = X_0$ is independent of the noise $N(t)$, and, for the drift parameter $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$, $\Gamma_\theta(t)$ is defined by

$$\Gamma_\theta(t) = \begin{cases} e^{-t\theta_1\theta_3/2}(\kappa_1 \cos(\nu t) + \kappa_{2,1} \sin(\nu t)), & \nu > 0, \\ e^{-t\theta_1\theta_3/2}(\kappa_1 \cosh(\nu t) + \kappa_{2,-1} \sinh(\nu t)), & \nu < 0, \\ e^{-t\theta_1\theta_3/2}(\kappa_1 + \kappa_{2,0} t), & \nu = 0, \end{cases} \tag{2}$$

for explicit constants $\nu = \nu(\boldsymbol{\theta}), \kappa_1 = \kappa_1(\boldsymbol{\theta})$, and $\kappa_{2,j} = \kappa_{2,j}(\boldsymbol{\theta})$, $j = -1, 0, 1$, given by the following:

$$\begin{aligned} \nu_0 &= \theta_1\theta_3/2, \\ \nu &= -\nu_0^2 + \theta_2^2(1-\theta_3), \\ \kappa_1 &= \theta_2^2\theta_3 - 2\theta_1(1-\theta_3)\nu_0, \\ \kappa_{2,1} &= \frac{1}{\nu}\left(\theta_1\theta_2^2 - \nu_0\left(\theta_2^2\theta_3 - \theta_1(1-\theta_3)2\nu_0\right) - \theta_1(1-\theta_3)\left(\nu_0^2 + \nu\right)\right), \\ \kappa_{2,-1} &= \frac{1}{\nu}\left(\theta_1\theta_2^2 - \nu_0\left(\theta_2^2\theta_3 - \theta_1(1-\theta_3)2\nu_0\right) - \theta_1(1-\theta_3)\left(\nu_0^2 - \nu\right)\right), \\ \kappa_{2,0} &= \theta_1\theta_2^2 - \nu_0\left(\theta_2^2\theta_3 - \theta_1(1-\theta_3)2\nu_0\right) - \theta_1(1-\theta_3)\nu_0^2. \end{aligned}$$

Here, the parameter space can be $\Theta = (0,\infty)^2 \times [0,1]$. Depending on the desired properties of $\mathbf{X}$ or the estimators, the space may be restricted to some $\Theta_0 \subset \Theta$.

Ref. [13] considered $N(t) = L(t)$ being a Lévy process with a finite second moment (or α-stable, $1 < \alpha \leq 2$).

For the classical Langevin equation, given as

$$dX(t) = -\theta X(t)dt + dL(t), \quad t > 0, \quad X(0) = X_0,$$

it's known that the Ornstein–Uhlenbeck process is its solution and can be obtained if we take $\theta_3 = 0$ in (1). On the other hand, if we take $\theta_3 = 1$, we get the GLE

$$dX(t) = -\theta^2 \int_0^t X(s)dsdt + dL(t), \;\; t > 0, \;\; X(0) = X_0, \;\; \theta \in \mathbb{R},$$

whose Cosine process (proposed by [27]) is the solution. That means θ_3 is a kind of "weight" between a stationary and ergodic Markovian process and a non-Markovian and non-stationary process (the Cosine process). Then, it is natural to think that θ_3 plays an important role in the quality of parameter estimation.

In this study, we are interested in investigating the GOU-FE process with multiplicative fBm noise. For the reader's convenience, the definition of fBm is presented below.

Definition 1. *We say that* $\mathbf{B^H} = \{B^H(t); t \geq 0\}$ *is a fractional Brownian motion (fBm) with Hurst parameter* $H \in (0,1)$ *if it is a centered Gaussian process with the covariance function*

$$R_H(t,s) = \frac{1}{2}\Big(s^{2H} + t^{2H} - |t-s|^{2H}\Big).$$

A natural question is whether we can adapt the model proposed by [13] to a GLE with the same drift as (1) but driven by multiplicative fBm $N(t) = B^H(t)$. This led us to propose a more general class of models:

$$dX(t) = \left(-\theta_1(1-\theta_3)X(t) - \int_0^t X(s)\Gamma_\theta(t-s)ds\right)dt + \sigma(X(t))dB^H(t), \;\; t > 0, \;\; X(0) = X_0. \tag{3}$$

Oscillating decays were observed in the autocorrelation functions for both the Cosine process (cf. [27]) and the GOU-FE (cf. [13]), unlike the classical Langevin equation, which has exponential decay. This suggests that the GOU-FE processes, as defined in (1), are capable of modeling an autoregressive time series of an order higher than 1. In summary, the introduction $\sigma(X(t))$ in the noise of the GOU-FE process (1) allows us to enhance the control of the volatility of the process. This impacts the quality of data modeling via (3).

Our approach has potential applications in various fields. One such application is in the modeling of anomalous diffusion based on the GLE. This phenomenon is observed in certain physical systems (cf. [20] and the references therein). Another significant application of our approach is in the modeling of financial assets using the GLE.

As a benchmark, we shall consider the traditional Geometric Brownian Motion process with either constant or stochastic volatilities (defined as $\sigma(t)$):

$$dX(t) = -\mu X(t)dt + \sigma(t)X(t)dW(t), \;\; t > 0, X(0) = X_0, \tag{4}$$

and $W(t)$ is a Brownian motion.

2.1. Stochastic Volatility Model

For this study, we adopt the Cox–Ingersoll–Ross (CIR) process, as originally proposed by [14], for $Y(t) = \sigma(X(t))$. This process is characterized by the following stochastic differential equation (SDE):

$$dY(t) = \kappa\left(\omega - Y(t)\right)dt + \xi\sqrt{Y(t)}\,dB(t), \tag{5}$$

where $\kappa > 0$ is the speed of mean reversion, $\omega > 0$ is the long-term level, $\xi > 0$ is the volatility, and $B(t)$ is a standard Brownian motion.

A fundamental property of the CIR process for discrete-time applications is the existence of an exact conditional distribution for $Y(t + \Delta t)$, given the value $Y(t) = y_0$ over a time interval Δt. As demonstrated by [14], this conditional distribution follows a scaled non-central chi-square distribution (χ'^2):

$$Y(t+\Delta t) \mid Y(t) = y_0 \;\sim\; c_0\, \chi'^2(\mathrm{df}, \lambda), \tag{6}$$

where the scale factor c_0, the degrees of freedom df and the non-centrality parameter λ are defined as functions of the CIR model parameters and the initial value y_0:

$$c_0 = \frac{\xi^2(1-e^{-\kappa\,\Delta t})}{4\,\kappa}, \quad \mathrm{df} = \frac{4\,\kappa\,\omega}{\xi^2}, \quad \lambda = \frac{4\,\kappa\,y_0\,e^{-\kappa\,\Delta t}}{\xi^2(1-e^{-\kappa\,\Delta t})}. \tag{7}$$

The use of this exact distribution is advantageous as it ensures the non-negativity of the process ($Y(t) \geq 0$), which is essential in many financial applications, and avoids biases that could arise from approximate discretization methods. Thus, the probability density function $p(Y(t+\Delta t) \mid Y(t); \kappa, \omega, \xi)$, computable from Equation (6), is essential for parameter estimation of the model using methods such as maximum likelihood or Bayesian inference, as also discussed by [28].

2.2. Parameter Estimation

Let $\{X(t); t \geq 0\}$ be a GOU-FE process (3) with stochastic volatility $Y(t) = \sigma(X(t))$ given by a CIR process (5), where $B^H(t)$ is an fBm with Hurst parameter $H \in (0,1)$.

For the initial estimation of the Hurst parameter, the `hurstexp` function from the `pracma` package in R is employed. This function utilizes the Rescaled Range (R/S) estimator, applied to the log-returns of Bitcoin prices (cf. [29,30]). The obtained value serves as the initial estimate for optimizing the models via the maximum likelihood method (using `optim()`), with distinct lower and upper bounds for each model.

For parameter and latent state estimation in nonlinear and non-Gaussian state-space models, such as the one proposed in this work, we employ the Sequential Monte Carlo (SMC) method, commonly known as the Particle Filter. This method allows for the recursive approximation of the posterior distribution of the latent states $Y(t)$ conditioned on the observations $X(1), X(2), X(3), \ldots, X(t)$, being particularly suitable for handling the nonlinearities and non-Gaussian features inherent in the model [31].

The implementation of the particle filter is based on representing the distribution of interest by a set of weighted samples (particles), $\{Y(t)^{(i)}, w(t)^{(i)}\}_{i=1}^N$, where $Y(t)^{(i)}$ are the particle states and $w_t^{(i)}$ are their respective normalized weights ($\sum_i w(t)^{(i)} = 1$). The iterative process of the filter, as described by [32,33], involves three main steps: propagation, weighting, and resampling.

First, the stochastic dynamic model is formulated in the state-space representation, defining the transition equation for the latent state $X(t)$ and the observation equation for the data $Y(t)$ (prices):

$$\begin{cases} X(t) = f_{\boldsymbol{\theta}}(X(t-1)) + \varepsilon_t, & \text{(State Equation)} \\ Y(t) = g_{\boldsymbol{\theta}}(X(t)) + \eta_t, & \text{(Observation Equation)} \end{cases}$$

where $\boldsymbol{\theta}$ denotes the vector of parameters to be estimated, and ε_t and η_t are the process and observation noise terms, respectively.

The standard particle filter algorithm (Bootstrap Filter) can be summarized as in Algorithm 1. The initialization step consists of generating N particles $\{X(0)^{(i)}\}_{i=1}^N$ from the prior distribution $p(X(0))$ and assigning uniform weights $w(0)^{(i)} = 1/N$.

Subsequently, for each time step $t = 1, \ldots, T$, the propagation and weighting steps are executed. In the propagation step, each particle is advanced according to the model dynamics, $X(t)^{(i)} \sim p_{\boldsymbol{\theta}}(X(t) \mid X(t-1)^{(i)})$. In the case of the CIR model, this transition can be sampled exactly [14].

In the weighting step, the weights are updated proportionally to the likelihood of the observation Y_t given the propagated particle, $w(t)^{(i)} \propto w(t-1)^{(i)} p_{\boldsymbol{\theta}}(Y(t) \mid X(t)^{(i)})$, and then normalized.

Algorithm 1 Particle Filter (Bootstrap Filter).

Initialization ($t = 0$)
for $i = 1, \ldots, N$
  Sample $X(0)^{(i)} \sim p(X(0))$
  Set $w(0)^{(i)} \leftarrow 1/N$
end for
for $t = 1, \ldots, T$
  Propagation and Weighting
  for $i = 1, \ldots, N$
    Propagation: Sample $\tilde{X}(t)^{(i)} \sim p_{\boldsymbol{\theta}}(X(t) \mid X(t-1)^{(i)})$
    Weighting: $\tilde{w}(t)^{(i)} \leftarrow w(t-1)^{(i)} \times p_{\boldsymbol{\theta}}\big(Y(t) \mid \tilde{X}(t)^{(i)}\big)$
  end for
  Normalization
  Evaluate $W(t) \leftarrow \sum_{j=1}^{N} \tilde{w}(t)^{(j)}$
  for $i = 1, \ldots, N$
    Set $w(t)^{(i)} \leftarrow \tilde{w}(t)^{(i)} / W(t)$
  end for
  Resample
  Evaluate ESS $\leftarrow 1 \,/\, \sum_{i=1}^{N} \big(w(t)^{(i)}\big)^2$
  if ESS $< N_{\text{threshold}}$ **then**
    Resample $\{X(t)^{(i)}, w(t)^{(i)}\}_{i=1}^{N}$ from $\{\tilde{X}(t)^{(i)}, w(t)^{(i)}\}_{i=1}^{N}$
    Set $w(t)^{(i)} \leftarrow 1/N \quad \forall\, i$
  else
    $\forall i$: $X(t)^{(i)} \leftarrow \tilde{X}(t)^{(i)}$
  end if
end for

In the context of this algorithm, resampling is employed to mitigate the phenomenon of weight degeneracy. This issue arises when a small number of particles carry most of the probability mass, compromising the representativeness of the distribution. To detect degeneracy, the Effective Sample Size (ESS) is calculated using ESS $= 1/\sum_{i=1}^{N}(w(t)^{(i)})^2$, and if it falls below a predefined threshold, such as $N/2$, resampling is triggered. After resampling, the selected particles are assigned uniform weights, equal to $1/N$.

The particle filter provides an approximation of the marginal likelihood function $p(Y(1), Y(2), Y(3), \ldots, Y(T)) \mid \boldsymbol{\theta})$, which is essential for parameter estimation. The likelihood can be calculated iteratively, and its approximation at time T is given by

$$\hat{L}(\boldsymbol{\theta}) = p(Y_1) \prod_{t=2}^{T} p(Y(t) \mid Y(1), Y(2), \ldots, Y(t-1)) \approx \prod_{t=1}^{T} \left(\frac{1}{N} \sum_{i=1}^{N} p_{\boldsymbol{\theta}}(Y(t) \mid X(t)^{(i)}) \right) \quad (8)$$

A more practical approximation is

$$\hat{L}(\boldsymbol{\theta}) \approx \prod_{t=1}^{T} \left(\sum_{i=1}^{N} w(t-1)^{(i)} p_{\boldsymbol{\theta}}(Y(t) \mid \tilde{X}(t^{(i)})) \right) \quad (9)$$

This likelihood approximation is then used in a numerical optimization procedure to obtain the maximum likelihood estimate (MLE) of the parameters, $\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} \log \hat{L}(\boldsymbol{\theta})$.

In this work, the L-BFGS-B method was used for optimization—a quasi-Newton method efficient for parameter-constrained problems [34]—as well as the Nelder–Mead method [35].

The main advantages of the particle filter include its flexibility in handling complex, nonlinear, and non-Gaussian models and its asymptotic convergence (as $N \to \infty$) to the optimal distribution. The computational complexity is of order $O(N)$ per time step.

3. ACD Models

In this section, we review the essential concepts of the ACD models, first introduced by [11]. This class of models was developed to analyze the time between events, particularly high-frequency financial data such as the duration between trades. By modeling these durations, ACD models provide insights into market liquidity and activity.

The original framework has inspired numerous extensions. Ref. [36] proposed the Generalized Gamma ACD model to offer a more flexible distributional assumption for the durations. The authors in [37] introduced a time-varying ACD model, without stationary assumptions. Others have shifted the modeling focus from the conditional mean to the conditional quantile; for instance, the authors in [38] developed the Birnbaum–Saunders ACD (BS-ACD) to model the conditional median, which was later extended in [39] and approached from a Bayesian perspective in [40]. More recently, the authors in [41,42] proposed conditional quantile frameworks for ACD models, while the authors in [43] highlighted the sensitivity of model estimators to the tail behavior of duration data. For a comprehensive treatment of the topic, we refer the reader to [44].

The main idea of the ACD model [11] is to assume that the duration X_t is the product of the conditional mean duration, ψ_t, and a standardized, independent, and identically distributed (i.i.d.) random error term, ϵ_t. The error term is drawn from a distribution with positive support and a unit mean. The model is specified as

$$X_t = \psi_t \epsilon_t, \quad \text{where} \quad \mathbb{E}[\epsilon_t] = 1. \tag{10}$$

The conditional mean duration ψ_t is defined based on the information set $\mathcal{F}_{t-1}$ available up to time $t-1$. It follows a GARCH-like autoregressive process:

$$\psi_t = \mathbb{E}[X_t | \mathcal{F}_{t-1}] = \omega + \sum_{i=1}^{p} \alpha_i X_{t-i} + \sum_{j=1}^{q} \beta_j \psi_{t-j}. \tag{11}$$

This specification is known as the $ACD(p,q)$ model. To ensure that the conditional duration ψ_t is always positive, the parameters are constrained to be non-negative: $\omega > 0$, $\alpha_i \geq 0$, and $\beta_j \geq 0$. For the process to be weakly stationary, it is required that $\sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{q} \beta_j < 1$. A key distinction between various ACD models lies in the choice of the probability distribution for the error term ϵ_t.

3.1. Exponential ACD Model

The simplest variant is the Exponential ACD (EACD) model, which assumes that the standardized durations ϵ_t follow a standard exponential distribution. In the EACD model, the observed duration X_t at time t (for example, the time elapsed between a price change) is decomposed into the product of a conditional mean duration ψ_t and a standardized innovation ϵ_t. Note that $\psi_t = \mathrm{E}[X_t \mid \mathcal{F}_{t-1}]$ represents the conditional expectation of X_t given the information set $\mathcal{F}_{t-1}$ (information set up to the most recent observation X_{t-1}), while $\epsilon_t = X_t / \psi_t$ has a unit mean and is assumed to follow a standard exponential distribution with a probability density function (PDF):

$$f(\epsilon_t) = e^{-\epsilon_t}, \quad \text{for} \quad \epsilon_t \geq 0.$$

Therefore, given that

$$X_t = \psi_t \epsilon_t, \quad \epsilon_t \sim \text{Exp}(1),$$

then

$$f(x_i \mid \mathcal{F}_{t-1}) = \frac{1}{\psi_t} \exp\Big(-\frac{x_t}{\psi_t}\Big).$$

The primary limitation of the EACD model is its assumption of a constant conditional hazard rate, which may be too restrictive for financial applications where time-varying risk profiles are common.

3.2. Weibull ACD Model

The Weibull ACD (WACD) model, also introduced in [11], provides greater flexibility by allowing for a non-constant hazard function. It assumes that ϵ_t follows a Weibull distribution. To ensure $\mathbb{E}[\epsilon_t] = 1$, the standard Weibull distribution is rescaled.

Let the PDF of a standard Weibull distribution with shape parameter $\gamma > 0$ and scale parameter $\lambda > 0$ be $g(z) = \frac{\gamma}{\lambda}\left(\frac{z}{\lambda}\right)^{\gamma-1} \exp\left(-\left(\frac{z}{\lambda}\right)^{\gamma}\right)$. Its mean is $\lambda\Gamma(1+1/\gamma)$. For ϵ_t to have a unit mean, we set the scale parameter $\lambda = [\Gamma(1+1/\gamma)]^{-1}$. The resulting PDF for ϵ_t is

$$f(\epsilon_t; \gamma) = \gamma(\Gamma(1+1/\gamma))^{\gamma} \epsilon_t^{\gamma-1} \exp(-(\epsilon_t \Gamma(1+1/\gamma))^{\gamma}).$$

The shape parameter γ determines the shape of the hazard function. If $\gamma > 1$, the hazard rate is increasing; if $\gamma < 1$, it is decreasing; and if $\gamma = 1$, the WACD model reduces to the EACD model with a constant hazard rate. Following the main ACD equation ($X_t = \psi_t \epsilon_t$), the conditional PDF for the observed duration X_t given the past information $\mathcal{F}_{t-1}$ is found by applying the change of variables $\epsilon_t = X_t/\psi_t$:

$$f(x_t|\mathcal{F}_{t-1}) = \frac{\gamma}{\psi_t}(\Gamma(1+1/\gamma)))^{\gamma} \left(\frac{x_t}{\psi_t}\right)^{\gamma-1} \exp\left(-\left(\frac{x_t}{\psi_t}\Gamma(1+1/\gamma)\right)^{\gamma}\right).$$

3.3. Generalized Gamma ACD Model

The Generalized Gamma (GG) ACD model, proposed in [36], further enhances flexibility in capturing the distribution of durations. The PDF of the GG distribution for ϵ_t is parameterized by two shape parameters, $\kappa > 0$ and $\gamma > 0$. The conditional density of X_t is given by

$$f(x_t|\mathcal{F}_{t-1}) = \frac{\gamma}{\psi_t \Gamma(\kappa)} \left(\frac{\Gamma(\kappa+1/\gamma)}{\Gamma(\kappa)}\right)^{\kappa\gamma} \left(\frac{x_t}{\psi_t}\right)^{\kappa\gamma-1} \exp\left(-\left(\frac{x_t\Gamma(\kappa+1/\gamma)}{\psi_t\Gamma(\kappa)}\right)^{\gamma}\right).$$

This form arises from rescaling the standard GG distribution to ensure the conditional mean is ψ_t, i.e., $\mathbb{E}[X_t|\mathcal{F}_{t-1}] = \psi_t$. The GG distribution nests several other distributions as special cases, providing a rich framework for duration modeling: Weibull distribution, when $\kappa = 1$; gamma distribution, when $\gamma = 1$; and exponential distribution, when $\kappa = 1$ and $\gamma = 1$.

4. Applications

This section presents the modeling of both daily and intraday data, according to the models detailed in Sections 2 and 3. The results were obtained using R software, version 4.4.1. To ensure reproducibility, the source codes are publicly available in the

GitHub repository, accessible at: https://github.com/Arthur-RPC/BitcoinSDE-ACD-Analysis.git (accessed on 7 September 2025).

For this study, daily price data were retrieved from the CoinMarketCap website using the `crypto2` package in R, which performs automated scraping of the historical records of several financial assets, including Bitcoin, the focus of this research. The analyzed period spans from 1 January 2019 to 31 December 2024. The dataset consists of daily closing prices, which represent the final recorded value of the asset at the end of each trading day.

For the intraday analysis of Bitcoin prices, we computed price durations from high-frequency financial data obtained from the Dukascopy website (available at: https://www.dukascopy.com/trading-tools/widgets/quotes/historical_data_feed), corresponding to 20 January 2025. The removal of seasonality and the model fits within the ACD framework were made using the R package ACDm [45].

4.1. Data Description

It is essential to understand the dynamics of Bitcoin prices over time in order to provide a solid foundation before applying statistical modeling. Thus, Figure 1 shows the trajectory of this cryptocurrency from 1 January 2019 to 31 December 2024.

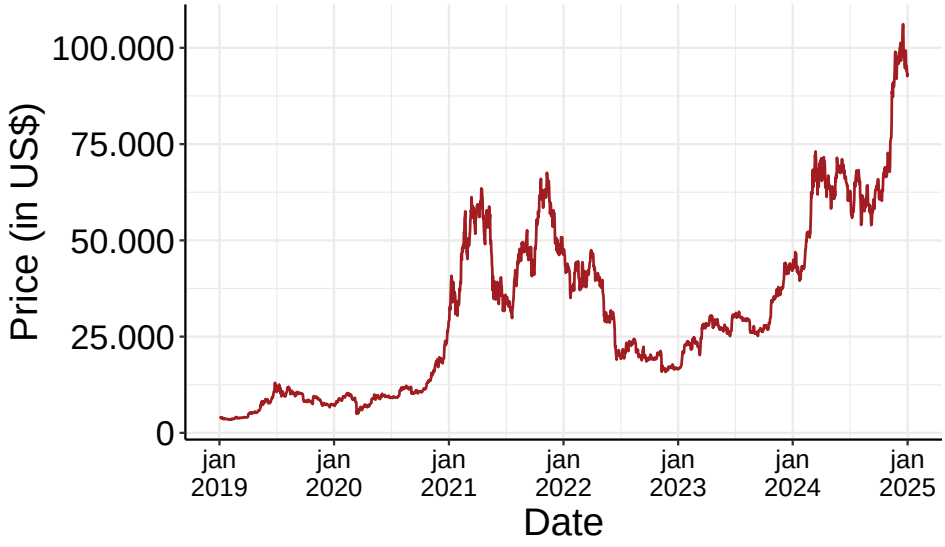

Figure 1. Bitcoin prices (in US$) from 1 January 2019 to 31 December 2024.

As shown in Figure 1, the analyzed period encompasses different market conditions. Between January 2019 and early 2020, Bitcoin prices remained relatively stable, staying below US$15,000. In 2021, there was a strong upward trend in the first half of the year, followed by a sharp correction. This price surge coincided, in part, with the COVID-19 pandemic, during which many investors began to view Bitcoin as a potential hedge against global economic instability [46].

Between 2023 and 2024, Bitcoin experienced another significant increase in value, driven by macroeconomic factors and specific events in the cryptocurrency market. This period culminated in a new price record, surpassing US$98,000 in November 2024, possibly influenced by political events such as the United States elections.

The presence of these different behaviors—from relative calm to high volatility, including sharp uptrends and corrections—makes the 2019–2024 interval ideal for testing the selected stochastic models. This allows for evaluating whether these models adequately adapt to drastic changes in price dispersion. Moreover, the occurrence of persistent trends and abrupt reversals creates a suitable context to test the ability of models based on fractional Brownian motion to capture temporal dependence.

Therefore, the choice of this period is not intended to avoid the influence of external factors or price instability but rather to use them as a stress test for the models. The ability of a stochastic model to describe and predict Bitcoin behavior over such a heterogeneous period, including both calm phases and major turbulence, attests to its robustness and relevance. In this context, the randomness and long memory, intrinsic characteristics of the studied models, will be assessed against a rich and challenging price history.

4.2. Daily Analysis

In this section, we use the abbreviations GOUFE-CIR and GOUFE-CONST to refer to the GOU-FE process driven by fBm (3), with the CIR model as the volatility function $Y(t) = \sigma(X(t))$, and the GOU-FE process with constant volatility ($\sigma \equiv 1$), respectively. Similarly, we use GFBM-CIR and GFBM-CONST to denote the corresponding GBM models obtained from (4).

To asses the performance of all four models (GOUFE-CIR, GOUFE-CONST, GFBM-CIR, and GFBM-CONST), their parameters were estimated. The results are presented in Tables 1 and 2, where dashes (–) indicate parameters not applicable to the constant volatility specification.

Table 1. GOUFE models' parameter estimates.

Model	Log-Likelihood	θ_1	θ_2	θ_3	κ	ω	ξ	H
GOUFE-CIR	20,888.66	0.0149634	0.005039290	0.9051001	2.790833	0.1454547	0.4370782	0.5511263
GOUFE-CONST	31,871.26	0.0100000	0.001028255	0.8999774	–	–	0.3624386	0.5461502

Table 2. GFBM models' parameter estimates.

Model	Log-Likelihood	μ	κ	ω	ξ	H
GFBM-CIR	21,170.86	0.0009889554	2.785846	0.1000000	0.3624460	0.5461402
GFBM-CONST	25,599.41	0.0035963469	–	–	0.7000000	0.5413182

From Tables 1 and 2, it can be observed that the models with constant volatility (GOUFE-CONST and GFBM-CONST) yield higher log-likelihood values compared to their respective counterparts with CIR-type stochastic volatility. This initial analysis suggests that, for the daily Bitcoin prices, the additional complexity introduced by modeling volatility through the CIR process may not be justified, and assuming constant volatility provides a more efficient fit to the data. Among all analyzed models, GOUFE-CONST stands out with the highest log-likelihood (31,871.26), indicating the best overall fit. Furthermore, the estimated Hurst exponent (H) remains consistently in the range of 0.54 to 0.55 across all models, highlighting the presence of long memory in the daily Bitcoin prices.

After evaluating the estimated parameters, predictive performance was assessed using error metrics that quantify the deviation between observed and predicted values. Accordingly, Figure 2 presents a comparison between the observed daily Bitcoin prices and the forecasts generated by the four models throughout the analysis period, while Table 3 summarizes this comparison based on the prediction errors.

Table 3. RMSE, MAE, MAPE, R^2, and bias for the prediction models (best metrics are in bold format).

Model	RMSE	MAE	MAPE	R^2	Bias
GFBM-CIR	**1201.802**	716.3009	**0.0225726**	0.9969002	**9.7625**
GOUFE-CONST	1213.196	**716.2269**	0.0225984	**0.9969822**	34.0668
GFBM-CONST	1217.385	722.6033	0.0227552	0.9967674	−72.3004
GOUFE-CIR	1235.990	742.0912	0.0236312	0.9968740	215.2351

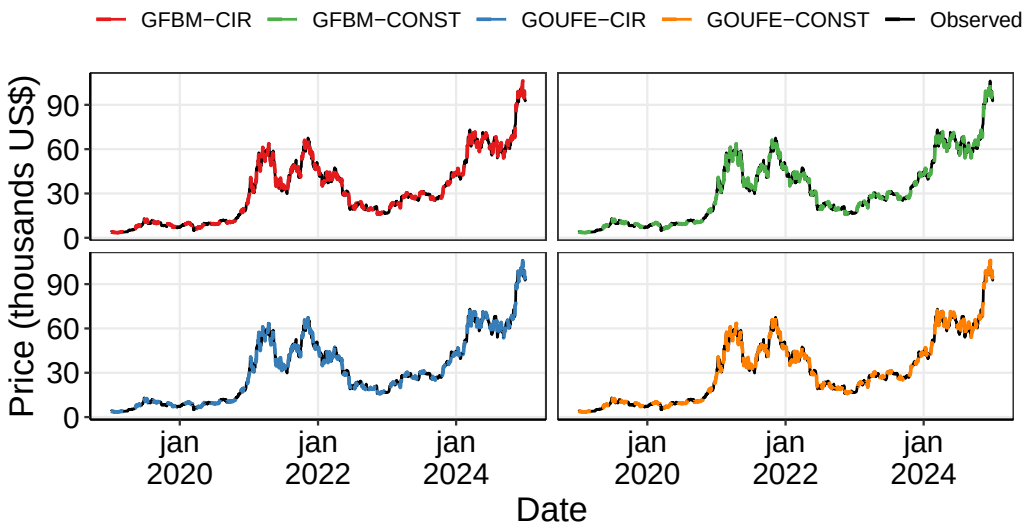

Figure 2. Predicted Bitcoin prices (in US$1000) vs. observed time series from 1 January 2019 to 31 December 2024.

Analyzing Figure 2, it can be seen that all models are able to follow the general trajectory of the observed prices, capturing the main market trends and movements. However, Table 3 shows that, from the perspective of prediction errors and the coefficient of determination, the GFBM-CIR model had the best performance, presenting the lowest RMSE (US$1201.802) and the lowest MAPE (2.25%). On the other hand, GOUFE-CONST achieved the lowest MAE (US$716.2269).

Additionally, the GOUFE-CONST model also obtained the highest R^2 value (0.9969822), indicating that it explains approximately 99.70% of the variability in observed prices, although the other models have very similar values. Regarding bias, the GFBM-CIR model showed the smallest absolute bias (US$9.7625), suggesting a very small tendency toward overestimation, whereas GOUFE-CIR presented the highest bias (US$215.2351).

Overall, the error metrics suggest that in terms of predictive accuracy, the GFBM-CIR and GOUFE-CONST models performed the best, although the other models are not far behind. It is worth noting that, for the GOUFE model, the adoption of stochastic volatility (CIR) did not bring a significant gain in accuracy despite the additional parameters. In contrast, for the GFBM model, introducing CIR volatility resulted in a noticeable improvement, with the CIR version having a bias of only US$9.7625, compared to US$−72.3004 for the constant volatility specification.

Nevertheless, to more confidently determine which model performs best, it is necessary to consider additional factors, such as the Information Criteria (AIC, BIC, and EDC), where lower values indicate a better fit by penalizing model complexity. Table 4 presents these results.

Table 4. Information criteria (AIC, BIC, and EDC) evaluation for the models (best metrics are in bold format).

Model	AIC	BIC	EDC
GOUFE-CIR	**46,161.33**	**58,638.44**	**46,249.44**
GFBM-CIR	46,725.71	59,202.82	46,813.82
GFBM-CONST	55,582.82	68,059.94	55,670.93
GOUFE-CONST	68,126.51	80,603.62	68,214.62

According to the results in Table 4, the GOUFE-CIR model achieved the lowest values for all three information criteria (AIC, BIC, and EDC), indicating the best overall fit by

balancing likelihood and model complexity. The GFBM-CIR model ranked second in this evaluation.

These results suggest that introducing stochastic volatility via the CIR (Cox–Ingersoll–Ross) process provided a significant gain in model fit that justifies the increased number of parameters, particularly when compared to the constant volatility versions (GOUFE-CONST and GFBM-CONST), which yielded consistently higher AIC, BIC, and EDC values.

However, it is important to note that information criteria are based on in-sample fit. When comparing these findings with the out-of-sample predictive performance reported in Table 3, a discrepancy in model rankings emerges. While GFBM-CIR and GOUFE-CONST performed better in terms of predictive accuracy, GOUFE-CIR remained the best model based on in-sample fit.

4.2.1. Residual Analysis

To evaluate model adequacy, a residual analysis was performed. This allows us to examine the distribution of residuals, including median, quartiles, and outliers, as shown in Figure 3 and Table 5.

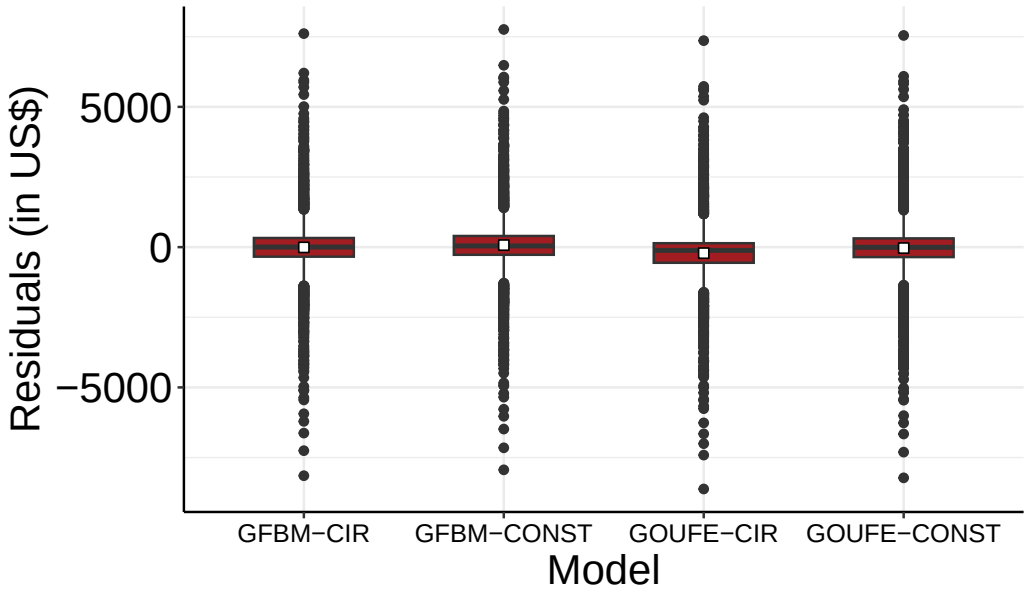

Figure 3. Residuals boxplot for GFBM-CIR, GFBM-CONST, GOUFE-CIR, and GOUFE-CONST models. GFBM-CIR and GOUFE-CONST provide the lowest bias in forecasting the central trend.

Table 5. Descriptive statistics of residuals for each model.

Model	Min.	Q1	Med.	Q3	Max.	SD
GOUFE-CONST	−8223.34	−355.14	−5.84	311.58	7544.64	1212.72
GFBM-CIR	−8147.72	−341.43	2.81	325.76	7610.12	1212.76
GFBM-CONST	−7937.89	−273.95	42.36	397.36	7757.97	1215.24
GOUFE-CIR	−8620.95	−558.77	−116.96	139.57	7359.55	1271.11

Based on Figure 3 and Table 5, we conclude that the GFBM-CIR and GOUFE-CONST models exhibit low bias in forecasting the central trend of Bitcoin prices, as their medians are very close to zero (2.81 and −5.84, respectively).

Regarding the dispersion of the residuals, the interquartile range (IQR = Q3 − Q1) is similar across models, ranging from 666.72 (GOUFE-CONST) to 698.34 (GOUFE-CIR). In this aspect, GOUFE-CIR shows the highest IQR and also the highest standard deviation (SD = 1271.11), indicating greater variability in its residuals. On the other hand, GFBM-CIR

(SD = 1212.76) and GOUFE-CONST (SD = 1212.72) have the lowest standard deviations, indicating less dispersion.

In terms of symmetry in the residual distribution, the GFBM-CIR model appears to be the most symmetric around its median, with similar distances between the median and quartiles (Med $-$ $Q1$ $\approx$ 344.24; $Q3$ $-$ Med $\approx$ 322.95). The GFBM-CONST model shows a slight right-skew, while the GOUFE-CIR model exhibits a more pronounced left-skew, consistent with its negative median.

A common feature in all four models is the presence of numerous outliers, which contribute to large prediction errors. The total range of residuals (Max–Min) is quite similar across models, varying from approximately 15,694 (GFBM-CONST) to 15,979 (GOUFE-CIR).

In Figure 4, we present the scatter plots of residuals against the observed Bitcoin prices for the four models fitted to the daily data: GFBM-CIR, GFBM-CONST, GOUFE-CIR, and GOUFE-CONST.

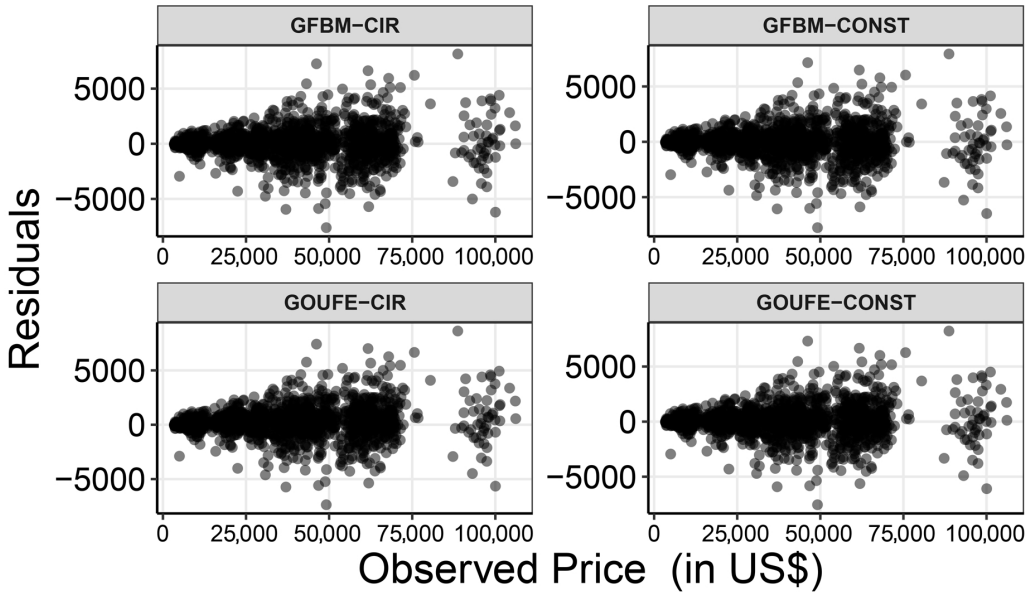

Figure 4. Residual dispersion plots for each model, which show that dispersion of residuals is not constant for the models.

According to Figure 4, in all four models, the dispersion of residuals is not constant across the range of observed prices. For lower prices (approximately below US$25,000), the residuals are more concentrated around zero, indicating smaller errors. However, as the price increases, the variability of the residuals also increases considerably, forming a pattern resembling a cone or fan. This behavior indicates the presence of heteroskedasticity, meaning that the error variance is not constant and tends to be larger at higher price levels. Thus, the models' prediction errors tend to increase as Bitcoin prices rise.

To gain further insight into the distribution of residuals, QQ-plots are provided below, comparing the quantiles of the residuals with the quantiles of a standard normal distribution $N(0,1)$. This is presented in Figure 5.

According to Figure 5, all four models show significant deviations from the reference line, with a large number of outliers and heavy tails, suggesting that the residuals do not follow a standard normal distribution.

This heavy-tailed behavior implies that prediction errors can occasionally be much larger than what is suggested by average-based measures such as RMSE or MAE. To better assess the risk associated with such extreme events, we use the Expected Shortfall (ES), also known as Conditional Value at Risk. The $ES_{1-\alpha}$ measures the average expected loss in the worst α% of scenarios.

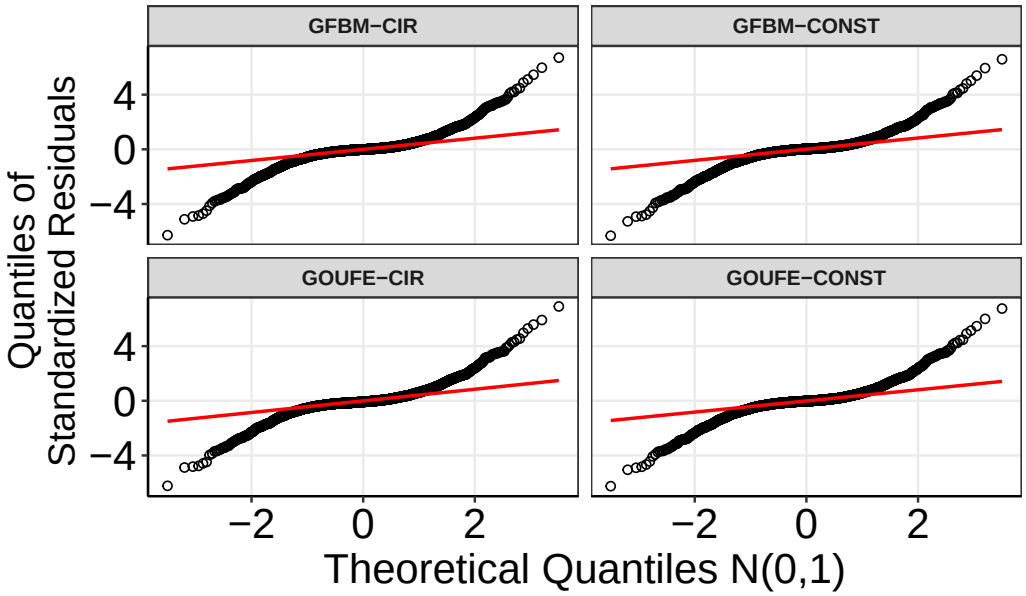

Figure 5. QQ-plots of standardized residuals for each model, which show significant deviations from the reference line representing the standard normal distribution.

Table 6 presents the Expected Shortfall values calculated at a 95% confidence level (ES_{95}) for each fitted model.

Table 6. Expected Shortfall at the 95% confidence level (ES_{95}) for each model (best metrics are in bold format).

Model	ES_{95}
GOUFE-CIR	−**2764.197**
GOUFE-CONST	−3020.207
GFBM-CIR	−3070.517
GFBM-CONST	−3212.101

Analyzing the results in Table 6, we conclude that the GOUFE-CIR model presented the lowest expected shortfall (US$−2764.197), considering the average losses in the worst 5% of scenarios. Next in ranking are the GOUFE-CONST (US$−3020.207) and GFBM-CIR (US$−3070.517) models, with very similar values.

In contrast, the GFBM-CONST model recorded the highest expected shortfall (US$−3212.101), indicating that it is the least conservative model.

4.2.2. Normality and Autocorrelation Tests on Residuals

To formally verify the normality hypothesis suggested in Figure 5, the Shapiro–Wilk test was applied and its results presented in Table 7. The null hypothesis (H_0) of this test is that the data follow a normal distribution.

As shown in Table 7, the p-values for all four models are extremely small, leading to rejection of the null hypothesis of normality in all cases.

To assess whether autocorrelation is present in the residuals of the fitted models, the Autocorrelation Function (ACF) plots are presented in Figure 6. These plots display the estimated autocorrelations of the residuals at different lags for each model. The blue dashed lines represent approximate significance limits (usually $\pm 1.96/\sqrt{N}$, where N is the sample size). Bars exceeding these limits indicate statistically significant autocorrelation at that specific lag.

Table 7. Shapiro–Wilk test for the standardized residuals.

Model	W	*p*-Value
GFBM-CONST	0.8615	2.86×10^{-40}
GOUFE-CIR	0.8605	2.24×10^{-40}
GFBM-CIR	0.8597	1.87×10^{-40}
GOUFE-CONST	0.8593	1.70×10^{-40}

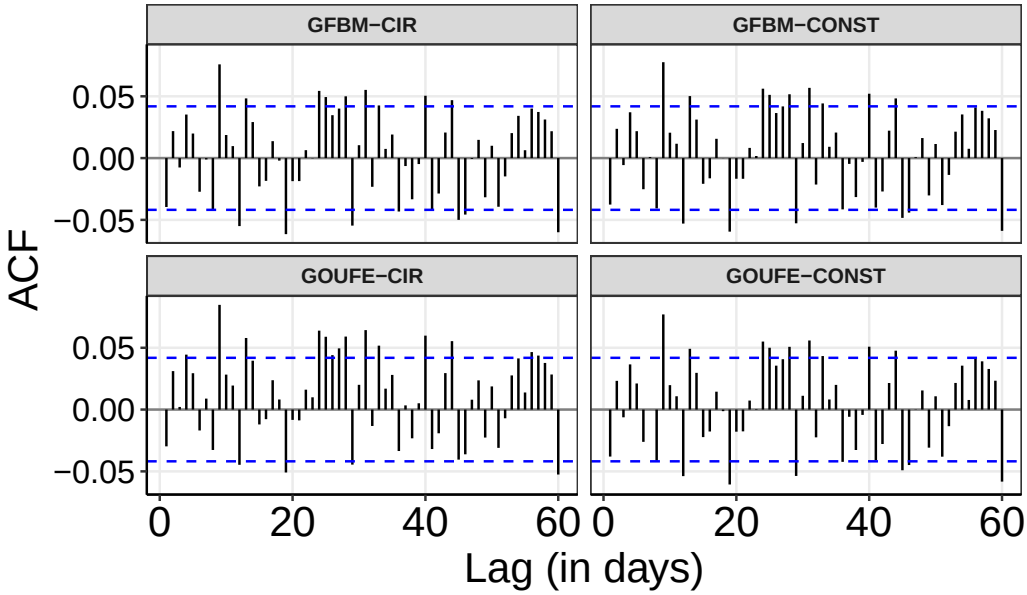

Figure 6. Autocorrelation function (ACF) plots of residuals for each model. For all models, we can see several bars exceeding the significance limits (plotted as blue dashed lines), which points to autocorrelation still present in residuals.

As seen in Figure 6, all four models show several bars exceeding the significance limits, suggesting that some autocorrelation structures are not fully captured by the models.

To formally and globally test for the presence of autocorrelation, the Ljung–Box test was applied, and its results are presented in Table 8. Its null hypothesis (H_0) is that autocorrelations up to a given lag are jointly equal to zero, i.e., the residuals are independent. According to Burns [47], the chosen lag should not exceed 5% of the sample size. For the current study, a lag of 60 days was selected, which falls well below the upper limit (109).

Table 8. Ljung–Box test for standardized residuals (lag = 60 days).

Model	Statistic	DF	*p*-Value
GFBM-CIR	160.6510	60	4.65×10^{-11}
GFBM-CONST	160.8942	60	3.67×10^{-11}
GOUFE-CONST	161.2322	60	3.29×10^{-11}
GOUFE-CIR	175.6447	60	2.80×10^{-13}

As evident from Table 8, the null hypothesis (H_0) is strongly rejected for all four models, as the *p*-values are significantly lower than any conventional significance level. Thus, residuals exhibit dependency at a 60-day lag, confirming the initial suspicion that these models retain some long-range temporal dependence.

4.3. Intraday Analysis

The choice of the price difference threshold is a crucial step in computing price durations, as it determines which price changes should be considered significant for generating durations (intervals between meaningful price changes). If the threshold is too small, the analysis may be dominated by market noise, leading to an excessive number of short durations. Conversely, if the threshold is too large, important market movements might be overlooked.

To find an appropriate balance, tests were conducted by varying the price threshold between 0.05%, 0.075%, 0.10%, and 0.15% of the average Bitcoin price on 20 January 2025. The impact of this choice on the resulting distribution of durations is presented in Figure 7 and Table 9.

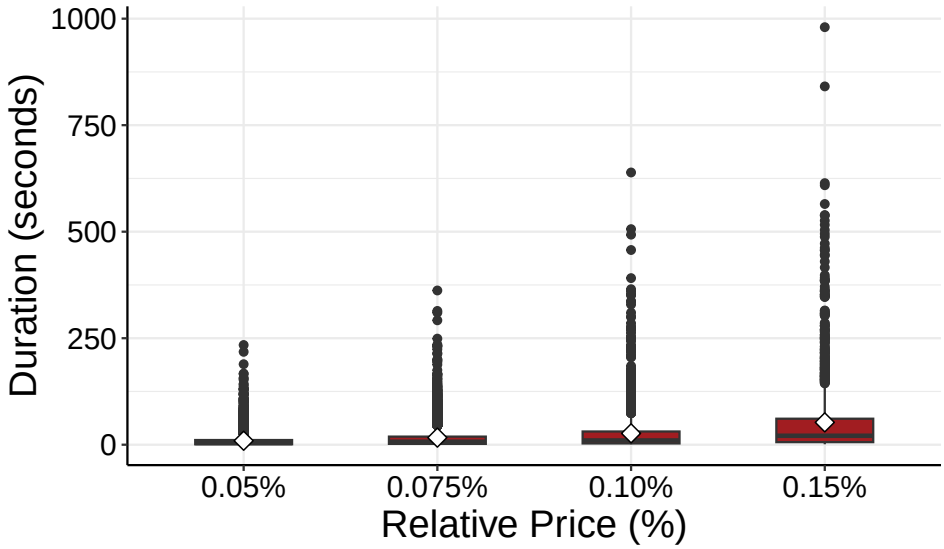

Figure 7. Boxplot of durations for different relative price thresholds. The 0.10%-price-band threshold offers a balanced trade-off by capturing important market movements without being overly influenced by noise.

As shown in Figure 7 and Table 9, a 0.05% threshold is highly sensitive to minor price fluctuations, generating 8893 durations with a median of just 4 s. In contrast, a 0.15% threshold drastically reduces the event count to 1579 and increases the median duration to 21 s. The 0.10% threshold offers a balanced trade-off, producing a substantial number of durations (3114) with a median of 10 s, thereby capturing significant market movements without being overly influenced by noise. Based on these findings, a threshold of 0.10%, corresponding to price changes of approximately US$104.74, was adopted for the remainder of this study.

Table 9. Descriptive statistics of durations by relative price threshold.

Statistic	**0.05%**	**0.075%**	**0.10%**	**0.15%**
# of durations	8893	4983	3114	1579
Min.	1	1	1	1
Q1	1	2	3	6
Median	4	7	10	21
Mean	9.31	16.61	26.58	52.44
Q3	11	19	31	61
Max.	234	362	639	980
SD	14.52	26.54	45.04	86.94

4.3.1. Duration Analysis Throughout the Day

To understand how market activity varies over a 24 h period, the generated durations were analyzed. In ACD models, shorter durations indicate periods of high market activity, whereas longer durations suggest lower liquidity or reduced volatility. Figure 8 displays the evolution of durations over the course of the day, highlighting periods of higher and lower price variation intensity.

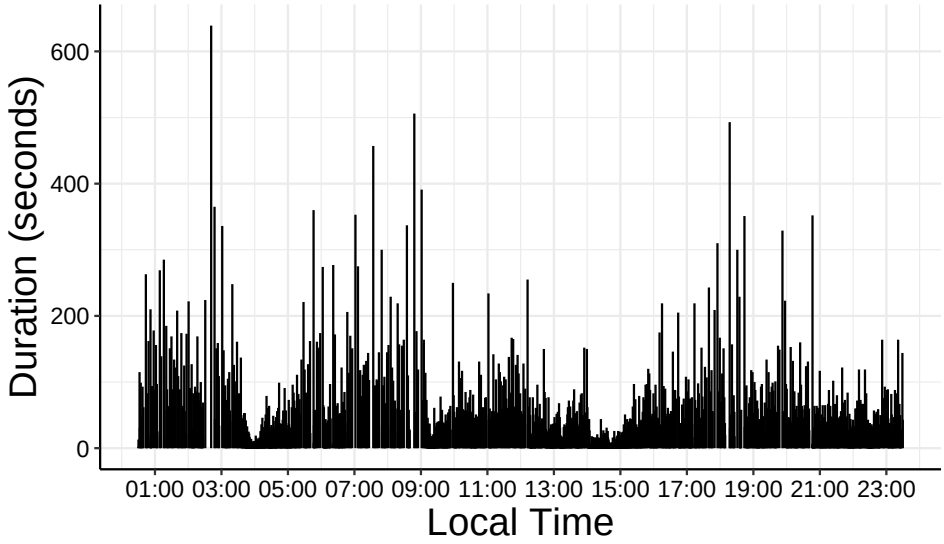

Figure 8. Intraday behavior of durations (GMT-3) for a relative price threshold of 0.10%.

From Figure 8, we note that significant variation in durations is observable throughout the day. The longest durations, where prices remained stable for over 6 minutes, occurred around 3:00 a.m., 7:00 a.m., 9:00 a.m., and 7:00 p.m. (GMT-3 timezone).

To remove the underlying diurnal pattern from durations, a smoothing technique known as the "Super Smoother", proposed by Friedman [48], was applied. This method identifies long-term trends in how durations vary over the day. The resulting pattern is shown in Figure 9.

Figure 9 reveals a cyclical intraday pattern in event durations. Durations are initially high, then drop sharply around 5:00 a.m., signaling an increase in trading activity. Between 5:00 a.m. and 3:00 p.m., successive peaks and troughs suggest fluctuating volatility, likely influenced by overlapping market sessions and liquidity windows. Towards the end of the day (after 3:00 p.m.), durations begin to rise again as market activity subsides.

After adjusting the durations to remove the time-of-day effect, the distribution of the adjusted durations is presented in Figure 10. The histogram reveals a highly right-skewed distribution, with the vast majority of durations concentrated near zero, indicating that most transactions occur within extremely short time intervals.

This sharp concentration around zero is consistent with the microstructure of financial markets, where bursts of activity tend to cluster in high-liquidity periods. Despite the adjustment, the presence of a long right tail suggests that some prolonged gaps between transactions still occur, potentially driven by episodic drops in trading intensity or structural breaks not fully captured by the diurnal correction.

The resulting shape underscores the need for flexible parametric models, such as the Generalized Gamma distribution, that can accommodate both heavy tails and asymmetry. This distributional feature will directly inform the choice of conditional duration models in subsequent sections and highlights the limitations of more restrictive specifications like the exponential distribution, which assumes memorylessness and constant hazard rate.

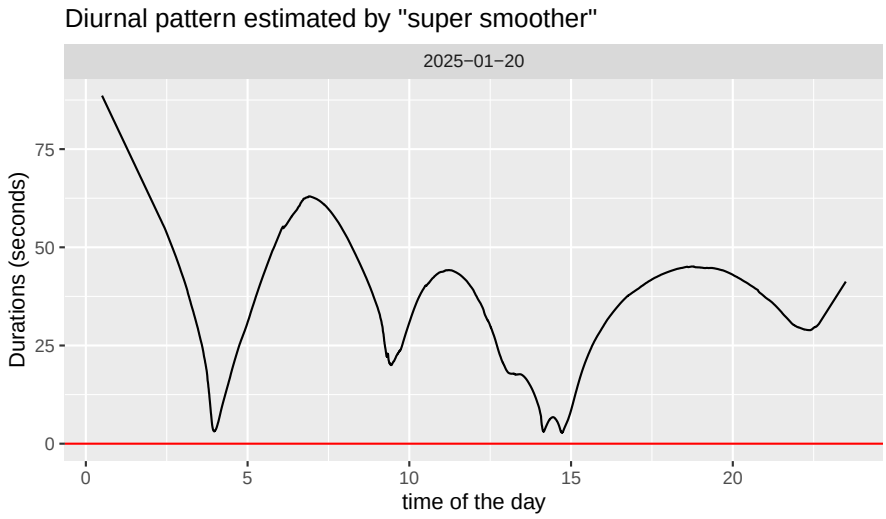

Figure 9. Estimated intraday (GMT-3) duration pattern using the Super Smoother method.

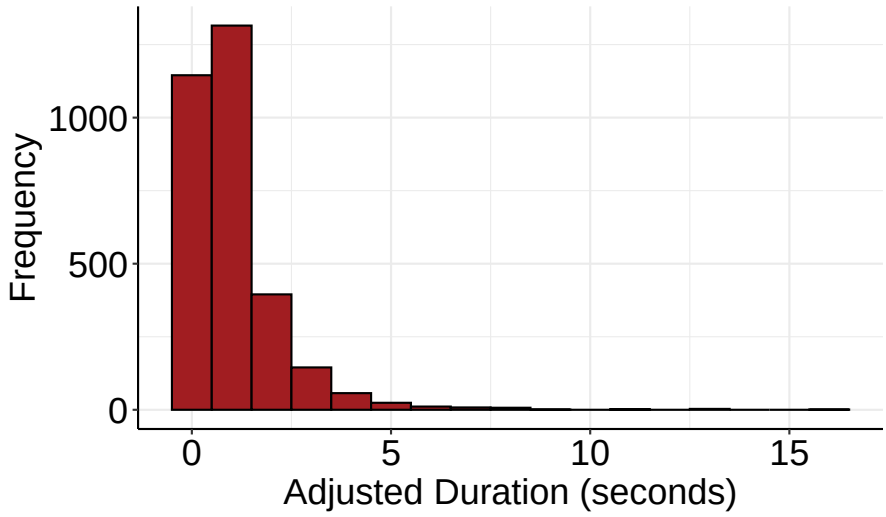

Figure 10. Adjusted duration (in seconds) histogram, showing that the majority of transactions is occurring in very short intervals, and also a long right tail.

Table 10 provides summary statistics for the adjusted durations, which are essential for selecting and parameterizing an appropriate ACD model. The statistics confirm right skewness (3.877) and high kurtosis (26.698). The mean (1.055 s) is greater than the median (0.704 s), which is characteristic of a right-skewed distribution. The high coefficient of variation (110.62%) underscores the significant dispersion of the data.

Table 10. Descriptive statistics of adjusted durations.

Statistic	Value
Min.	0.012
Median	0.704
Mean	1.055
Max.	15.506
Standard Deviation (SD)	1.168
Coefficient of Variation (%)	110.618
Skewness	3.877
Kurtosis	26.698
Range	15.494

4.3.2. Model Fit Analysis for Durations

To model the adjusted durations, three $ACD(1,1)$ models were fitted and compared: Exponential, Weibull, and Generalized Gamma. The parameter estimates for each model are presented in Tables 11–13.

Table 11. Exponential $ACD(1,1)$ model parameter estimates.

Parameter	Estimate	SD	*p*-Value
ω	0.0405	0.00519	0.000
α_1	0.1063	0.01184	0.000
β_1	0.9626	0.01004	0.000

Table 12. Weibull $ACD(1,1)$ model parameter estimates.

Parameter	Estimate	SD	*p*-Value
ω	0.0402	0.00440	0.000
α_1	0.1035	0.00989	0.000
β_1	0.9655	0.00821	0.000
γ	1.1722	0.01528	0.000

Table 13. Generalized Gamma $ACD(1,1)$ model parameter estimates.

Parameter	Estimate	SD	*p*-Value
ω	0.0469	0.00557	0.000
α_1	0.1226	0.01194	0.000
β_1	0.9473	0.01107	0.000
κ	20.7112	8.20236	0.016
γ	0.2464	0.04950	0.000

In all three models, the parameters α_1 and β_1 are highly statistically significant, indicating a strong temporal dependence where past durations heavily influence the current duration.

- Exponential Model: The parameter estimates show high persistence, with the sum $\alpha_1 + \beta_1 \approx 1.0689$.
- Weibull Model: This model also indicates strong temporal dependence. The shape parameter $\gamma = 1.1722$ is significant and greater than 1, suggesting that the hazard rate of durations is increasing and that the distribution is more peaked with heavier tails than the exponential distribution.
- Generalized Gamma Model: The estimates for α_1 and β_1 again confirm temporal persistence. The additional shape parameters, $\kappa = 20.7112$ and $\gamma = 0.2464$, are both significant, affording the model greater flexibility to capture the complex skewness and kurtosis of the duration data.

The evaluation of model fit for the ACD(1,1) specification was conducted using a combination of log-likelihood values and standard model selection criteria: AIC (Akaike Information Criterion), BIC (Bayesian Information Criterion), and MSE (Mean Squared Error). As presented in Table 14, the Generalized Gamma distribution produced the best overall performance, attaining the highest log-likelihood and the lowest values for AIC and BIC. Although the MSE differences across models are subtle, they remain consistent with the ranking obtained from the likelihood-based criteria.

These results indicate that the Generalized Gamma model provides a more flexible and accurate representation of the conditional distribution of durations, particularly when compared to the more restrictive Exponential and Weibull alternatives.

Table 14. $ACD(1,1)$ models' fit assessment (best metrics are in bold format).

Model	Log-Likelihood	AIC	BIC	MSE
Exponential	−3102.44	6210.88	6229.01	**1.2376**
Weibull	−3034.89	6077.78	6101.95	1.2377
Generalized Gamma	**−2885.06**	**5780.13**	**5810.35**	1.2381

In addition to the numerical criteria, the adequacy of the fitted models was also examined through Cox–Snell residual analysis. QQ-plots of the residuals (Figure 11) illustrate the extent to which the transformed residuals conform to the theoretical $\text{Exp}(1)$ distribution. The Exponential and Weibull models display clear deviations from the reference line, especially in the upper quantiles, indicating poor tail modeling. Conversely, the Generalized Gamma model yields residuals that closely align with the diagonal, suggesting a more accurate capture of the underlying duration dynamics.

Overall, the combination of information criteria and residual diagnostics supports the Generalized Gamma distribution as the most appropriate choice for modeling intraday durations under the ACD(1,1) structure.

To quantify the fit in the tail of the distribution, the Expected Shortfall (ES) was computed at a 95% confidence level (ES_{95}), representing the average of the most extreme 5% values; see Table 15. From this table, we note that the Generalized Gamma model achieved an $\text{ES}_{95\%}$ of 3.9753, which is lower than the $\text{ES}_{95\%}$ of the Exponential model (4.0759) and the Weibull model (4.882). This result indicates that the Generalized Gamma model provides a better fit to the expected behavior of an $\text{Exp}(1)$ distribution.

Finally, to ensure that the models successfully captured all temporal dependence, the Cox–Snell residuals were tested for autocorrelation using ACF plots (Figure 12) and the Ljung–Box test (Table 16). The Ljung–Box test results show p-values well above 0.05 for all models, so the null hypothesis of no autocorrelation is not rejected. The ACF plots in Figure 12 visually support this, with most correlations falling within the confidence bands.

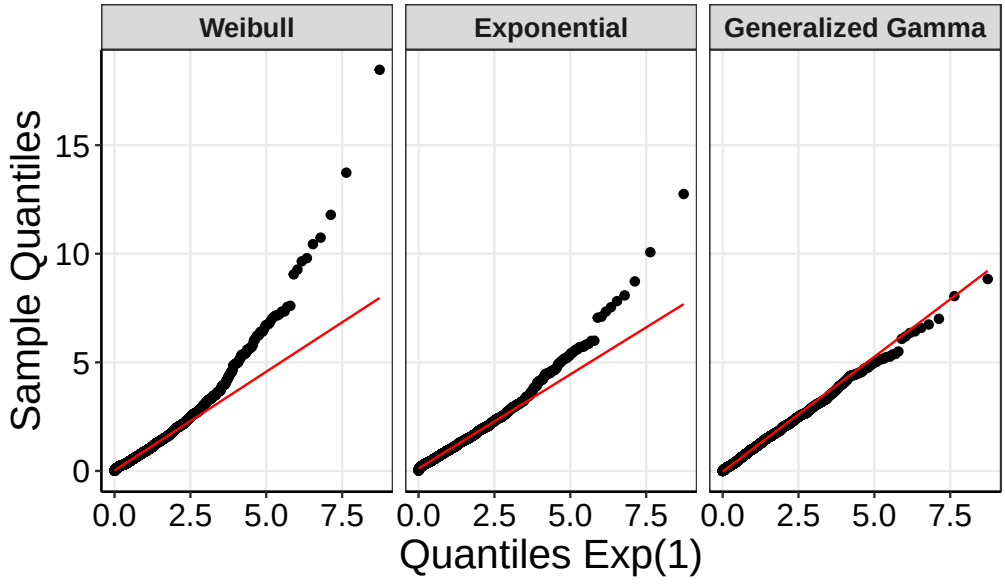

Figure 11. QQ-plots of Cox–Snell residuals for each model. The Generalized Gamma closely follows the reference line confirming its good fit.

Table 15. Expected Shortfall at the 95% confidence level (ES_{95}) for each model (best metric shown in bold format).

Model	ES_{95}
Generalized Gamma	**3.975315**
Exponential	4.075854
Weibull	4.881840

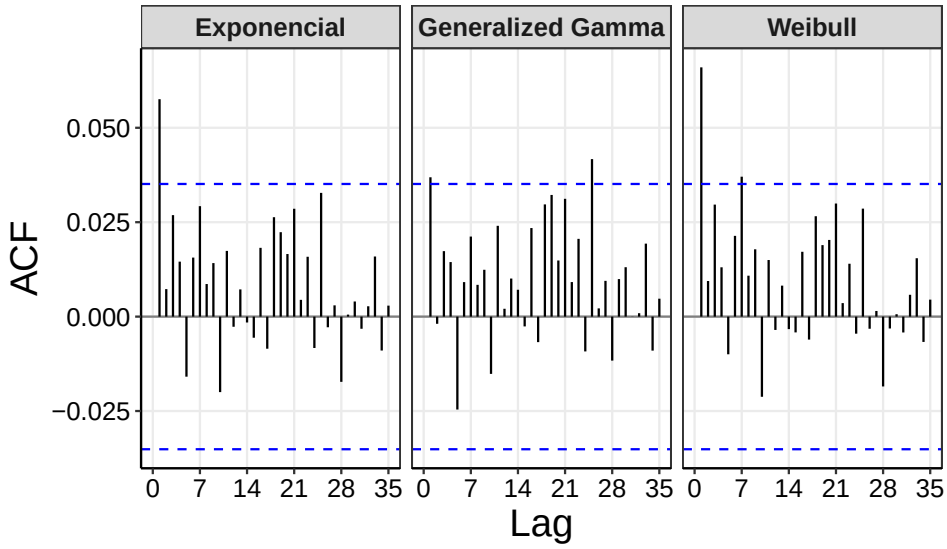

Figure 12. Cox–Snell residuals ACF plots for each model (lag = 35). Exceeding bars (that stretch outside of the region delimited by blue dashed lines) point to the presence of autocorrelation in residuals for all models.

Table 16. Ljung–Box test for Cox–Snell residuals (lag = 35).

Model	Statistic	DF	*p*-Value
Generalized Gamma	35.16990	33	0.3657253
Exponential	36.02138	33	0.3289920
Weibull	41.00386	33	0.1596591

5. Conclusions

This study focuses on the stochastic modeling of Bitcoin prices, addressing both daily and intraday dynamics, to capture complex features such as long memory, stochastic volatility, and patterns in transaction durations. The objective was to propose and evaluate models based on SDEs driven by fBm for daily data and ACD models for intraday analysis.

In the daily analysis, GFBM and GOUFE models were compared, incorporating both constant and stochastic volatility specifications via the CIR model. The evaluation—based on error metrics, information criteria, and risk analysis—indicated that models with stochastic volatility provided a better fit, suggesting that incorporating volatility dynamics is indeed relevant. Residual analysis, however, revealed the presence of heteroskedasticity (i.e., increasing variance with price) across all models. Although autocorrelation was mitigated, these findings indicate that despite progress, the models still fall short of fully capturing the complexity of the distribution of Bitcoin's daily returns.

For the intraday analysis, Exponential, Weibull, and Generalized Gamma ACD models were applied to price durations. The results are seen to be quite favorable to the Generalized Gamma model in terms of model fitting.

The improved model identification discussed in this paper has practical implications for real-world dynamic trading and risk management, as it enables a more accurate un-

derstanding of price dynamics and time durations. Importantly, statistical comparisons using information criteria and error measures demonstrate that our proposed methods consistently outperform conventional specifications, thereby confirming the robustness of the modeling approach.

However, there are some challenges in practical applications. One limitation is the computational costs associated with making near-instantaneous decisions in fast-moving extreme markets. The algorithms showed good convergence when using Nelder–Mead optimization, whereas L-BFGS-B did not perform as well. Additionally, the choice of initial guesses plays an important role in the drift estimation of the GOUFE-CIR model. Another limitation is the unobservability of the stochastic volatility process, which complicates real-time estimation and may require the use of filtering techniques or proxy variables to infer volatility dynamics.

Overall, the framework we explored in this work can be a starting point to study the GOUFE-CIR with high-frequency data observations and how to estimate the parameters in this case.

Author Contributions: Conceptualization, A.R.P.d.C., F.Q. and H.S.; methodology, A.R.P.d.C., F.Q. and H.S.; software, A.R.P.d.C.; validation, F.Q. and H.S.; formal analysis, F.Q., H.S., T.A.d.F., L.C.S.M.O. and P.N.R.; investigation, A.R.P.d.C.; writing—original draft preparation, A.R.P.d.C., F.Q. and H.S.; writing—review and editing, T.A.d.F., L.C.S.M.O. and P.N.R. All authors have read and agreed to the published version of the manuscript.

Funding: Helton Saulo gratefully acknowledges financial support from the University of Brasília and the Brazilian National Council for Scientific and Technological Development (CNPq) through grant number 304716/2023-5. The APC was kindly waived by the Editorial Office of FinTech.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data can be accessed via the links in the article.

Acknowledgments: The authors acknowledge the support provided by the University of Brasilia (UnB). The authors would like to thank the editor and three anonymous referees for their useful comments, which helped us improve the manuscript.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Nakamoto, S. Bitcoin: A Peer-to-Peer Electronic Cash System. *SSRN* **2008**. [CrossRef]
2. Antonopoulos, A.M. *Mastering Bitcoin: Programming the Open Blockchain*, 2nd ed.; O'Reilly Media: Sebastopol, CA, USA , 2017; pp. 115–130.
3. Merediz-Solà, I.; Bariviera, A.F. A bibliometric analysis of bitcoin scientific production. *Res. Int. Bus. Financ.* **2019**, *50*, 294–305. [CrossRef]
4. Rejeb, A.; Rejeb, K.; Alnabulsi, K.; Zailani, S. Tracing Knowledge Diffusion Trajectories in Scholarly Bitcoin Research: Co-Word and Main Path Analyses. *J. Risk Financ. Manag.* **2023**, *16*, 355. [CrossRef]
5. Antonopoulos, A.M. *Mastering Bitcoin: Unlocking Digital Cryptocurrencies*, 1st ed.; O'Reilly Media: Sebastopol, CA, USA, 2014; pp. 89–101.
6. Garnier, J.; Solna, K. Chaos and order in the bitcoin market. *Phys. A Stat. Mech. Its Appl.* **2019**, *524*, 708–721. [CrossRef]
7. Hotta, L.K.; Trucíos, C.; Valls Pereira, P.L.; Zevallos, M. Forecasting Bitcoin and Ethereum risk measures through MSGARCH models: Does the specification matter? *Braz. Rev. Financ.* **2025**, *23*, e202503. [CrossRef]
8. Rathie, P.; Ozelim, L. Exact and approximate expressions for the reliability of stable Lévy random variables with applications to stock market modelling. *J. Comput. Appl. Math.* **2017**, *321*, 314–322. [CrossRef]
9. Rathie, P.; Ozelim, L.; Otiniano, C. Exact distribution of the product and the quotient of two stable Lévy random variables. *Commun. Nonlinear Sci. Numer. Simul.* **2016**, *36*, 204–218. [CrossRef]

10. Alhagyan, M.; Yassen, M.F. Incorporating stochastic volatility and long memory into geometric Brownian motion model to forecast performance of Standard and Poor's 500 index. *AIMS Math.* **2023**, *8*, 18581–18595. [CrossRef]
11. Engle, R.; Russell, J. Autoregressive Conditional Duration: A New Method for Irregularly Spaced Transaction Data. *Econometrica* **1998**, *66*, 1127–1162. [CrossRef]
12. Dimpfl, T.; Odelli, S. Bitcoin Price Risk—A Durations Perspective. *J. Risk Financ. Manag.* **2020**, *13*, 157. [CrossRef]
13. Quintino, F.S.; Medino, A.V.; Dorea, C.C. Drift estimation for a class of generalized Ornstein-Uhlenbeck process with fluctuating exponential trend. *Commun. Stat.-Simul. Comput.* **2025**, *54*, 1161–1174. [CrossRef]
14. Cox, J.C.; Ingersoll, J.E.; Ross, S.A. A theory of the term structure of interest rates. *Econometrica* **1985**, *53*, 385–407. . . [CrossRef]
15. Barndorff-Nielsen, O.; Shephard, N. Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *J. R. Stat. Soc. Ser. B* **2001**, *63*, 167–241. [CrossRef]
16. Mori, H. Transport, Collective Motion, and Brownian Motion. *Prog. Theor. Phys.* **1965**, *33*, 423–455. [CrossRef]
17. Kubo, R. The fluctuation-dissipation theorem. *Rep. Prog. Phys.* **1966**, *29*, 255–284. [CrossRef]
18. Kannan, D. On the Generalized Langevin Equation. *J. Math. Phys. Sci.* **1977**, *11*, 1–24.
19. Kannman, D.; Bharucha-Reid, A. Random integral equation formulation of a generalized Langevin equation. *J. Stat. Phys.* **1972**, *5*, 209–233. [CrossRef]
20. Medino, A.; Lopes, S.; Morgado, R.; Dorea, C. Generalized Langevin equation driven by Lévy processes: A probabilistic, numerical and time series based approach. *Phys. A Stat. Mech. Its Appl.* **2012**, *391*, 572–581. [CrossRef]
21. Zwanzig, R. *Nonequilibrium Statistical Mechanics*; Oxford University Press: Oxford, UK, 2001.
22. Ottobre, M.; Pavliotis, G.A. Asymptotic analysis for the generalized Langevin equation. *Nonlinearity* **2011**, *24*, 1629. [CrossRef]
23. McKinley, S.A.; Nguyen, H.D. Anomalous diffusion and the generalized Langevin equation. *SIAM J. Math. Anal.* **2018**, *50*, 5119–5160. [CrossRef]
24. DiTerlizzi, I.; Ritort, F.; Baiesi, M. Explicit solution of the generalised Langevin equation. *J. Stat. Phys.* **2020**, *181*, 1609–1635. [CrossRef]
25. Pavliotis, G.A.; Stoltz, G.; Vaes, U. Scaling limits for the generalized Langevin equation. *J. Nonlinear Sci.* **2021**, *31*, 8 . [CrossRef]
26. Zhu, Y.; Venturi, D. Generalized Langevin equations for systems with local interactions. *J. Stat. Phys.* **2020**, *178*, 1217–1247 . [CrossRef]
27. Stein, J.; Lopes, S.; Medino, A. Continuous processes derived from the solution of generalized Langevin equation: Theoretical properties and estimation. *J. Stat. Comput. Simul.* **2016**, *86*, 2819–2845. [CrossRef]
28. Glasserman, P. *Monte Carlo Methods in Financial Engineering*; Springer: New York, NY, USA, 2004. [CrossRef]
29. Hurst, H.E. Long-term storage capacity of reservoirs. *Trans. Am. Soc. Civ. Eng.* **1951**, *116*, 770–808. [CrossRef]
30. Mandelbrot, B.B.; Wallis, J.R. Robustness of the rescaled range R/S in the measurement of noncyclic long run statistical dependence. *Water Resour. Res.* **1969**, *5*, 967–988. [CrossRef]
31. Doucet, A.; de Freitas, N.; Gordon, N.J. (Eds.) *Sequential Monte Carlo Methods in Practice*; Springer: New York, NY, USA, 2001. [CrossRef]
32. Gordon, N.J.; Salmond, D.J.; Smith, A.F.M. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proc. F (Radar Signal Process.)* **1993**, *140*, 107–113. [CrossRef]
33. Arulampalam, M.S.; Maskell, S.; Gordon, N.; Clapp, T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Trans. Signal Process.* **2002**, *50*, 174–188. [CrossRef]
34. Liu, J.S.; West, M. Combined parameter and state estimation in simulation-based filtering. In *Sequential Monte Carlo Methods in Practice* ; Springer: New York, NY, USA, 2001; pp. 197–223.
35. Nelder, J.A.; Mead, R. A Simplex Method for Function Minimization. *Comput. J.* **1965**, *7*, 308–313. [CrossRef]
36. Lunde, A. A Generalized Gamma Autoregressive Conditional Duration Model. Technical Report, 1999. Available online: https://www.researchgate.net/publication/228464216_A_generalized_gamma_autoregressive_conditional_duration_model (accessed on 7 September 2025).
37. Bortoluzzo, A.B.; Morettin, P.A.; Toloi, C.M. Time-varying autoregressive conditional duration model. *J. Appl. Stat.* **2010**, *37*, 847–864. [CrossRef]
38. Bhatti, C.R. The Birnbaum–Saunders autoregressive conditional duration model. *Math. Comput. Simul.* **2010**, *80*, 2062–2078. [CrossRef]
39. Cunha, D.R.; Vila, R.; Saulo, H.; Fernandez, R.N. A general family of autoregressive conditional duration models applied to high-frequency financial data. *J. Risk Financ. Manag.* **2020**, *13*, 45. [CrossRef]
40. Fernando, N.; Jeremias, L.; Saulo, H. Bayesian inference for the Birnbaum–Saunders autoregressive conditional duration model with application to high-frequency financial data. *Commun. Stat. Case Stud. Data Anal. Appl.* **2021**, *7*, 215–228. [CrossRef]
41. Saulo, H.; Balakrishnan, N.; Vila, R. On a quantile autoregressive conditional duration model. *Math. Comput. Simul.* **2023**, *203*, 425–448. [CrossRef]

42. Saulo, H.; Pal, S.; Souza, R.; Vila, R.; Dasilva, A. Parametric Quantile Autoregressive Conditional Duration Models With Application to Intraday Value-at-Risk Forecasting. *J. Forecast.* **2025**, *44*, 589–605. [CrossRef]
43. Cavaliere, G.; Mikosch, T.; Rahbek, A.; Vilandt, F. Tail behavior of ACD models and consequences for likelihood-based estimation. *J. Econom.* **2024**, *238*, 105613. [CrossRef]
44. Hautsch, N. *Modelling Irregularly Spaced Financial Data: Theory and Practice of Dynamic Duration Models*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2004.
45. Belfrage, M. *ACDm: Tools for Autoregressive Conditional Duration Models*, R package version 1.0.4.3. 2024. Available online: https://cran.r-project.org/web/packages/ACDm/ACDm.pdf (accessed on 7 September 2025).
46. Goodell, J.W.; Goutte, S. Co-movement of COVID-19 and Bitcoin: Evidence from wavelet coherence analysis. *Financ. Res. Lett.* **2021**, *38*, 101625. [CrossRef]
47. Burns, P. Robustness of the Ljung-Box Test and its Rank Equivalent. *SSRN Electron. J.* **2002** . [CrossRef]
48. Friedman, J.H. *A Variable Span Smoother*; Technical Report LCS Technical Report 5; Department of Statistics, Stanford University: Stanford, CA, USA, 1984.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.