Non-Gaussian Stochastic Volatility Models - Laplace-Variational Bayes Inference

João Pedro Nacinben and Márcio Laurini*
University of São Paulo

Abstract

Stochastic volatility models are fundamental tools in finance for accurately estimating and managing risks, primarily due to their ability to accommodate a dynamic and time-varying volatility structure. However, a notable constraint within these models is the reliance on Gaussian processes to model the latent (log-)variance, which can limit their ability to effectively capture events such as sudden jumps or spikes in the latent volatility. To address this limitation, we employ a non-Gaussian SV model utilizing an inference procedure that combines Laplace and Variational Bayes approximations. Our study showcases the significant advantages of this correction in modeling the conditional variance of Bitcoin's return series.

Keywords: Conditional volatility, financial risks, jumps, Variational inference.

1 Introduction

Stochastic volatility (SV) models, since originally introduced by Taylor (1986), have been reliably employed in the estimation and management of risks in financial markets, due to their flexibility to capture dynamic and time-varying volatility structures. Although SV models incur some degree of estimation complexity, given the unobservable nature of volatility, more recent developments, such as the integrated nested Laplace approximations (INLA) methodology, allow for faster and computationally efficient estimation, even for a considerable increase in numbers of observations or of the dimensionality of the problem itself (Nacinben and Laurini, 2024).

*Correspondence: laurini@fearp.usp.br

The method of Integrated Nested Laplace Approximations (INLA), a Bayesian alternative approach to the usual MCMC (Markov chain Monte Carlo) algorithms that are widely used to approximate posterior distributions, was first proposed by Rue et al. (2009) for the class of latent Gaussian models (LGMs). Inference regarding stochastic volatility models, could then be performed avoiding common MCMC convergence-related issues entirely, as presented in Martino et al. (2011), and extended to multivariate and multifactors models by Nacinben and Laurini (2024).

However, one significant drawback in SV models arises from the conventional reliance on Gaussian processes to model the latent (log-)variance, as their ability to effectively capture abrupt changes or spikes in volatility is limited by the assumption of Gaussianity. In financial series, discontinuous jumps and heavy-tailed distributions for returns are common occurrences, making LGMs not the ideal tools for modeling.

Incorporating heavier-tailed models into the dynamics of latent volatility can better address the inadequacies posed by these assumptions. However, direct INLA-based estimation becomes unfeasible in this context. In order not to give up INLA's advantages in terms of computation efficiency, Cabral et al. (2024) develop a new variational inference (VI) strategy for estimating what they describe as latent non-Gaussian models (LnGMs). The authors have developed an algorithm that decomposes the variational Bayes (VB) inference problem into fitting LGMs. This allows for approximated inference using INLA as an alternative to the significantly slower MCMC sampling method.

With this new literature in mind, our study aims to address the limitations imposed by Gaussian processes in the context of stochastic volatility modeling, proposing a non-Gaussian SV model combining variational Bayes and Laplace approximations. The model is applied in the analysis of the conditional variance of Bitcoin's return series. Evidence suggests that this widely followed cryptocurrency exhibits heavy-tailed patterns in the latent volatility process, indicative of significant discontinuous jumps (Chaim and Laurini, 2018). We focus on the normal-inverse Gaussian (NIG) distribution as a heavy-tailed alternative to the Gaussian distribution, the latter being a special case of the first.

The primary contribution of the methodology presented in this study is a computationally efficient implementation for non-Gaussian SV models, offering a Bayesian alternative to MCMC-based estimation that is both faster and free of convergence issues. Its applications range from portfolio allocation to the formulation of tail risk measures and hedging procedures.

In addition to this introductory section, the paper is structured as follows. Section 2 covers the theoretical aspects regarding INLA and variational inference for the class of latent non-Gaussian models (LnGMs). In Section 3, we discuss the method to be employed.

The results are presented in section 4. Finally, concluding remarks and a summary of the key findings are left for section 5.

2 Literature review

Traditionally, stochastic volatility (SV) models rely on Gaussian processes to model the latent variance. Both volatility and returns series, however, often present patterns such as discontinuous jumps or heavy-tailed distributions, which are hardly congruent with the assumption of Gaussianity. Naturally, non-Gaussian SV models have been proposed to deal with such particularities of the data.

As an example, Barndorff-Nielsen (1997) proposes the normal-inverse Gaussian (NIG) distribution as a good prospect for conditional volatility modeling. As pointed out by the author, despite being mathematically simpler, this distribution can approximate the majority of hyperbolic distributions closely, in addition to offering a better fit to stock returns data in some cases. Ando (2006), on the other hand, assumes a Student-t distribution for both the returns and volatility error terms and uses a Markov chain Monte Carlo (MCMC) algorithm to estimate the model parameters. The study shows a good fit to the Nikkei 225 index data is achieved, but estimation is deemed as time consuming.

As a matter of fact, approximate Bayesian inference via MCMC algorithms can face some setbacks related to the simulation-based nature of these methods. Slow chain convergence is a well-known issue, but large sample sizes or increased problem dimensionality can also take a toll on computational efficiency. Rue et al. (2009) introduce, then, the integrated nested Laplace approximations (INLA) methodology, entirely based on deterministic approximations to posterior distributions and parallelizable sparse linear algebra operations.

Standard SV models can be represented as latent Gaussian models (LGMs), as demonstrated by Martino et al. (2011), by constructing the following latent field $\mathbf{x}$:

$$\mathbf{x} = \{h_1, ..., h_n, \mu\} \sim N(0, \mathbf{Q}^{-1}(\boldsymbol{\theta_1})), \tag{1}$$

where μ is a mean parameter with a Gaussian prior, $h_t (t = 1, \ldots, n)$ is the log-volatility series and $\mathbf{Q}$ is a precision matrix. Hyperparameter $\boldsymbol{\theta_1} = \{\tau_h, \phi\}$ contains parameters associated to the volatility. The latent field $|\mathbf{x}|$ undergoes partial observation through a conditionally independent set of returns data $\mathbf{r}$ and likelihood expressed as

$$\pi(\mathbf{r}|\mathbf{x}, \boldsymbol{\theta_2}) = \prod_{t=1}^{n} \pi(r_t|\mathbf{h}, \boldsymbol{\theta_2}), \tag{2}$$

where n is the number of observations and $\boldsymbol{\theta_2}$ represents parameters related to the normally distributed ε_t. Thus, given $\boldsymbol{\theta} = \{\boldsymbol{\theta_1}, \boldsymbol{\theta_2}\}$, SV models estimation can be carried on by computing the following marginals:

$$\pi(\mathbf{x}, \boldsymbol{\theta}|\mathbf{r}) \propto \pi(\boldsymbol{\theta})\pi(\mathbf{x}|\boldsymbol{\theta}) \prod_{t=1}^{n} \pi(r_t|h_t, \boldsymbol{\theta}). \tag{3}$$

The INLA scheme, presented in detail in Rue et al. (2009), can then precisely approximate $\pi(h_t|\mathbf{r})$, $\pi(\mu|\mathbf{r})$ and $\pi(\theta_j|\mathbf{r})$ by initially performing inference on the marginals of $\pi(\mathbf{x}, \boldsymbol{\theta}|\mathbf{r})$.

Authors such as Rue et al. (2017) and Chaim and Laurini (2019), the former tackling specifically long-memory SV models, show that INLA is a considerably faster alternative to MCMC methods, while also being able to estimate parameters of interest just as accurately. These advantages, however, are only available in the case of problems that can be modeled as LGMs, since INLA is only capable of dealing with this class of models.

Cabral et al. (2024) propose a new methodology for computationally efficient inference for the class of latent non-Gaussian models (LnGMs), using variational Bayes algorithms and INLA. The authors construct these LnGMs by substituting the Gaussian process in the latent field $\mathbf{x}$ of the LGM structure with a non-Gaussian model, using a mixture representation of $\mathbf{x}$ for an assumed NIG-distributed driving noise. In general terms, if the latent field $\mathbf{x}^G$ follows an n-dimensional Gaussian distribution, with $\mathbf{0}$ mean and precision matrix $\mathbf{Q} = \mathbf{D}^T\mathbf{D}$, we have that:

$$\mathbf{D}\mathbf{x}^G \stackrel{d}{=} \mathbf{Z}, \tag{4}$$

in which $\mathbf{Z} = [Z_1, \ldots, Z_n]^\top$ is a vector of Gaussian variables. So, according to the authors, an extension of $\mathbf{x}^G$ to non-Gaussianity could be expressed by:

$$\mathbf{D}\mathbf{x} \stackrel{d}{=} \boldsymbol{\Lambda}, \tag{5}$$

where $\boldsymbol{\Lambda}$ is a vector generalized hyperbolic (GH) random variables depending on a particular parameter η that controls non-Gaussianity through changes in the kurtosis of the distribution

In their study, the authors assume Λ_i follows a symmetric NIG distribution, which is demonstrated by Cabral et al. (2022) to be a flexible extension to Gaussian models. In fact, as the authors show, when $\eta = 0$ the model is identical to the Gaussian case, while increasingly higher values for η result in the distribution converging to the Cauchy distribution. And, since the NIG can be represented as a normal variance-mean mixture with an inverse Gaussian distribution as the mixing density, it is also demonstrated by the authors that the latent field x has the following mixture representation:

$$\mathbf{x}|\mathbf{V} \sim \mathrm{N}\left[\mathbf{0}, \quad \mathbf{D}^{-1}\mathrm{diag}(\boldsymbol{V})\mathbf{D}^{-T}\right], \quad V_i|\eta \overset{ind.}{\sim} \mathrm{IG}(h_i, \eta^{-1}h_i^2), \tag{6}$$

where h_i are predetermined constants[1] and $V = [V_1, \ldots, V_n]^\top$ is a mixing vector in which $V_i \sim IG(1, \eta^{-1})$. When the Gaussian assumption on the latent field in a LGM is replaced by this mixture representation, the resulting hierarchical structure can be defined as a latent non-Gaussian Model.

2.1 Variational inference for Latent Non-Gaussian models

In this section we present the basic elements of the variational approach proposed by Cabral et al. (2024) to estimate non-Gaussian latent models. Variational inference (VI) techniques (Jordan et al., 1999), depart from conventional Markov Chain Monte Carlo sampling methods for approximating a posterior distribution $\pi(\mathbf{z}|\mathbf{y})$. Instead, VI seeks a surrogate density $q(\mathbf{z})$ by solving an optimization problem:

$$q(\mathbf{z}) = \underset{q\in\mathcal{Q}}{\arg\min}\left\{\mathrm{KLD}(q(\mathbf{z})|\pi(\mathbf{z}|\mathbf{y}))\right\} = \underset{q\in\mathcal{Q}}{\arg\max}\, E_{q(\mathbf{z})}\left(\log\left(\frac{\pi(\mathbf{y},\mathbf{z})}{q(\mathbf{z})}\right)\right), \tag{7}$$

where KLD denotes the Kullback-Leibler divergence. The expectation on the right side of Equation (7) represents the evidence lower bound (ELBO), and maximizing the ELBO is equivalent to minimizing the KLD due to their relationship, expressed as $\mathrm{ELBO}(q(\mathbf{z})) = -\mathrm{KLD}(q(\mathbf{z})|\pi(\mathbf{z}|\mathbf{y})) + \log \pi(\mathbf{y})$. In this context, $\pi(\mathbf{y})$ denotes the evidence. As the Kullback-Leibler divergence is always non-negative, the evidence lower bound (ELBO) furnishes a lower limit on the log-evidence. Wang and Blei (2019) and related works provide theoretical insights into variational inference (VI) techniques.

In the structural variational inference (VI) framework, the objective is to discover the optimal surrogate density $q(\mathbf{x}, \boldsymbol{\theta}, \mathbf{V}, \eta) = q(\mathbf{x}, \boldsymbol{\theta})q(\mathbf{V})q(\eta)$. Here, the sole constraint within the space $\mathcal{Q}$ is the posterior independence among the blocks $(\mathbf{x}, \boldsymbol{\theta})$, $\mathbf{V}$, and η. The main result is presented in Theorem 1 of Cabral et al. (2024), where $q(\mathbf{x}, \boldsymbol{\theta})$ represents the posterior distribution of a Latent Gaussian Model (LGM), and the mixing variables V_i and parameter η adhere to a generalized inverse Gaussian distribution (GIG).

The result is given in Theorem 1 of Cabral et al. (2024), with $q(\mathbf{x}, \boldsymbol{\theta})$ being the posterior distribution of an LGM, mixing variables V_i and parameter η following a generalized inverse

[1] Constants are equal to 1 for models defined in discrete space. See Cabral et al. (2022) for more details.

Gaussian distribution (GIG). The GIG distribution is characterized by the density:

$$\pi_{\text{GIG}}(x; p, a, b) = \frac{(a/b)^{p/2}}{2K_p(\sqrt{ab})} x^{(p-1)} e^{-(ax+b/x)/2}, \quad x > 0, \tag{8}$$

where $K_\lambda(x)$ is the modified Bessel function of the second kind of order λ.

We present the Theorem 1 of Cabral et al. (2024) below:

Theorem 1. *Cabral et al. (2024). The surrogate density* $q(\mathbf{x}, \boldsymbol{\theta}, \mathbf{V}, \eta) = q(\mathbf{x}, \boldsymbol{\theta}) q(\mathbf{V}) q(\eta)$ *that minimises* $\text{KLD}(q(\mathbf{x}, \boldsymbol{\theta}, \mathbf{V}, \eta) | \pi(\mathbf{x}, \boldsymbol{\theta}, \mathbf{V}, \eta | \mathbf{y}))$ *is a solution of the system:*

$$\begin{aligned}
q(\mathbf{x}, \boldsymbol{\theta}) &\sim \text{pLGM}\{\pi(\mathbf{y}|\mathbf{x}, \boldsymbol{\theta}_1),\ \mathbf{m} = \mathbf{0},\ \mathbf{Q} = \mathbf{D}(\boldsymbol{\theta}_2)^\top \text{diag}(\mathbf{V}^{(-)}) \mathbf{D}(\boldsymbol{\theta}_2),\ \pi(\boldsymbol{\theta})\}, \\
q(V_i) &\sim \text{GIG}\left(-1,\ E_{q(\eta)}(\eta^{-1}),\ E_{q(\mathbf{x},\theta)}([\mathbf{D}\mathbf{x}]_i^2]) + h_i^2 E_{q(\eta)}(\eta^{-1})\right), \quad i = 1, \ldots, N, \\
q(\eta) &\sim \text{GIG}\left(-N/2 + 1,\ 2\alpha_\eta,\ \sum_{i=1}^{N} E_{q(V_i)}(V_i) - 2h_i + h_i^2 E_{q(V_i)}(V_i^{-1})\right),
\end{aligned}$$

where $V_i^{(-)} = E_{q(V_i)}(V_i^{-1})$.

Collapsed variational inference (CVI) capitalizes on the concept of analytically integrating specific model parameters (Zhang et al. (2018)). By reducing the number of parameters requiring estimation and eliminating hierarchical correlations, inference procedures are generally expedited. For the LnGM model we can integrate out η from $\pi(\mathbf{V}|\eta)$ obtaining:

$$\pi(\mathbf{V}) = \int_0^\infty \left(\prod_{i=1}^{N} \pi(V_i|\eta) \right) \pi(\eta) d\eta, \tag{9}$$

We use the structured and collapsed variational inference (SVCI) method proposed by Cabral et al. (2024) to estimate the non-Gaussian SV model. The SVCI method combines the use of INLA and CVI to estimate the parameters of the latent Gaussian model and the mixing parameter η. Details of the algorithm and the implementation can be found in Cabral et al. (2024).

3 Model

As a starting point, we consider the traditional formulation for the stochastic volatility (SV) model, which is represented by the following equations:

$$r_t = \exp\{h_t/2\}\varepsilon_t, \quad \varepsilon_t \sim N(0, 1) \tag{10}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \xi_t, \quad \xi_t \sim N(0, [(1-\phi^2)\tau]^{-1}) \tag{11}$$

where r_t denotes the log-returns, h_t is the log-volatility, μ is a mean parameter and ϕ is the persistence for the autoregressive process given by h_t. Both ε_t and ξ_t are normally distributed error terms.

This standard SV representation can be thought of as a latent Gaussian model (LGM), in which the latent field $\mathbf{x}$ encompasses the volatility series and its mean parameter, that is, $\mathbf{x} = \{h_1, \ldots, h_n, \mu\} \sim N(0, \mathbf{Q}^{-1}(\boldsymbol{\theta_1}))$, where $\boldsymbol{\theta_1} = \{\tau, \phi\}$. In line with the methodology introduced by Cabral et al. (2024), we propose a non-Gaussian SV model with a symmetric NIG distributed noise for the unobserved log-volatility equation, therefore extending the latent field $\mathbf{x}$ to non-Gaussianity.

Thus, keeping the LGM formulation as a benchmark, our non-Gaussian SV model can be represented in LnGM form as follows:

$$\mathbf{r}|\mathbf{x}, \boldsymbol{\theta_2} \sim \prod_{i \in I} \pi(r_i|x_i, \boldsymbol{\theta_2}) \tag{12}$$

$$\mathbf{x}|\mathbf{V}, \boldsymbol{\theta_1} \sim \mathrm{N}\left[\mathbf{0}, \mathbf{D}(\boldsymbol{\theta_1})^{-1}\mathrm{diag}(\boldsymbol{V})\mathbf{D}(\boldsymbol{\theta_1})^{-T}\right], \tag{13}$$

$$V_i|\eta \overset{ind.}{\sim} \mathrm{IG}(1, \eta^{-1}), \tag{14}$$

$$\boldsymbol{\theta} \sim \pi(\boldsymbol{\theta}), \eta \sim \mathrm{Exp}(\alpha_\eta) \tag{15}$$

where $\boldsymbol{\theta_2}$ are parameters related to the process ε_t and η, the non-Gaussianity control parameter, is assigned a exponential prior. Estimation can then be conducted according to the variational Bayes scheme proposed by Cabral et al. (2024), using a structured and collapsed variational inference (SCVI) algorithm.

Regarding the implementation of the model in a practical sense, we first fit an univariate SV model (LGM) using the **R-INLA** package, which then generates an 'inla' object for the model that can be used for the estimation of the non-Gaussian SV model we propose. The 'ngvb' function, included in the package of same name[2], takes the 'inla' object as an argument and fits the LnGM we're interested in, extending to non-Gaussianity a selection of the components of the LGM previously fitted.

[2]The **ngvb** package was developed by Cabral et al. (2024) and can be found in https://github.com/rafaelcabral96/ngvb.

4 Results

We present in this section the estimation results for the proposed non-Gaussian SV model, which is to be compared to its traditional Gaussian counterpart. First, however, we briefly discuss the adopted dataset.

4.1 Data

The empirical segment of this study utilizes daily returns data for Bitcoin (BTC) over the specified time frame, spanning from January 2nd, 2019, to December 12th, 2023, resulting in a total of 1822 observations. Figure 1 visually represents the log-returns of Bitcoin, offering a graphical insight into its price dynamics over the specified date range.

Figure 1: Daily returns for Bitcoin.

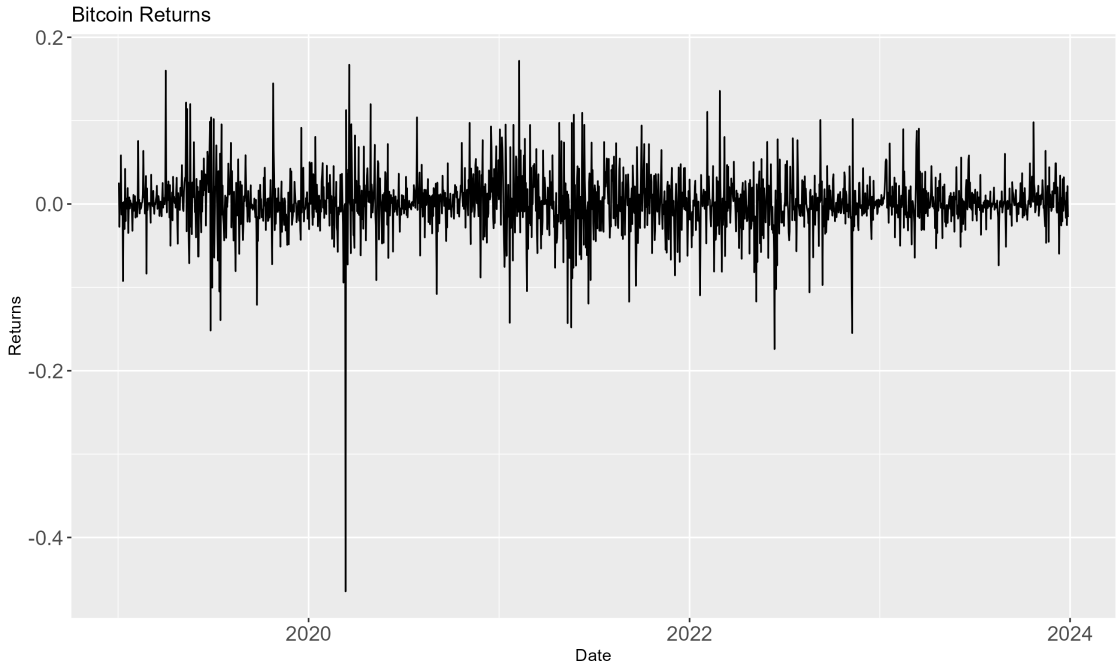

Cryptocurrency returns data are believed to display some features of particular interest to the analysis, as discussed by Phillip et al. (2018). The authors identify long memory patterns, leverage effect and heavy-tailed return distributions, noting varying asymmetry in the relationship between returns and volatility across crypto-assets. Specifically for evidence of non-Gaussianity in returns series for Bitcoin, the most traded cryptocurrency in the world, Chaim and Laurini (2018) highlight, in their study, the significance of jumps in explaining

large negative price variations, while jumps in the permanent volatility component enhanced their stochastic volatility model's fit during periods of heightened volatility.

Finally, Table 1 offers some descriptive statistics concerning the log-returns of Bitcoin, encompassing key metrics such as mean, median, standard deviation, skewness, kurtosis and the total number of observations.

Table 1: Descriptive statistics for the Bitcoin return series

	mean	median	sd	skewness	kurtosis	n
BTC	0.0013	0.0007	0.0355	−1.227	21.961	1822

4.2 Empirical analysis

We present in Tables 2 and 3 the posterior distribution for the latent Gaussian and non-Gaussian SV model specifications. We can observe that the means μ_g and μ_{ng} for the latent log-volatility are similar in both estimations. The persistence parameters ϕ_g and ϕ_{ng} are substantially different between the two specifications, with a posterior mean of 0.653 for the Gaussian specification and 0.587 for the non-Gaussian specification, indicating that the adoption of a non-Gaussian dynamic implies less persistence for the latent volatility dynamics. Note that this lower persistence is compensated by a lower precision. τ_{ng} is estimated with a posterior mean of 0.282 for the non-Gaussian process, while the precision τ_{ng} of the Gaussian specification is estimated with a posterior mean of 0.653.

Thus, we can observe that the non-Gaussian specification indicates a process with less persistence, but greater variability, which is consistent with the variance structure observed in the Bitcoin series, which presents large variations in returns, compatible with jump processes with low predictability. The parameter η_{ng} is estimated with a posterior distribution that does not contain the zero value in the 95% credibility interval, evidencing the deviation from Gaussianity for this process.

Figure 2 shows the temporal dynamics of the posterior mean of the log-variance fitted by the two models. We can observe that the log-variance of Gaussian and non-Gaussian specifications are similar. In Figure 3 a comparison is presented between the posterior average of volatilities estimated by Gaussian and non-Gaussian specifications, superimposed on the absolute returns of the Bitcoin series. We can observe that the non-Gaussian specification better adjusts the most extreme values in absolute returns, which is consistent with the assumption of a distribution with heavy tails for the latent volatility process.

Table 2: Posterior Distribution of Estimated Parameters - Latent Gaussian Model (LGM)

		mean	**sd**	**0.025q**	**0.5q**	**0.975q**
BTC	μ_g	-7.41	0.084	-7.573	-7.411	-7.244
	τ_g	0.653	0.063	0.541	0.648	0.789
	ϕ_g	0.725	0.042	0.639	0.725	0.805

Note: μ_g denotes the mean parameter for the volatility h_t, while τ_g denotes the marginal precision and ϕ_g the persistence parameter.

Table 3: Posterior Distribution of Estimated Parameters - Latent Non-Gaussian Model (LnGM)

		mean	**sd**	**0.025q**	**0.5q**	**0.975q**
BTC	μ_{ng}	-7.44	0.069	-7.578	-7.444	-7.307
	τ_{ng}	0.282	0.025	0.238	0.280	0.336
	ϕ_{ng}	0.587	0.049	0.491	0.586	0.682
	η_{ng}	1.130	0.108	0.919	1.127	1.360

Note: μ_{ng} denotes the mean parameter for the volatility h_t, while τ_{ng} denotes the marginal precision and ϕ_{ng} the persistence parameter. Lastly, η_{ng} is the non-Gaussianity parameter (=0 if Gaussian).

Figure 2: Log-variance fitted for both models.

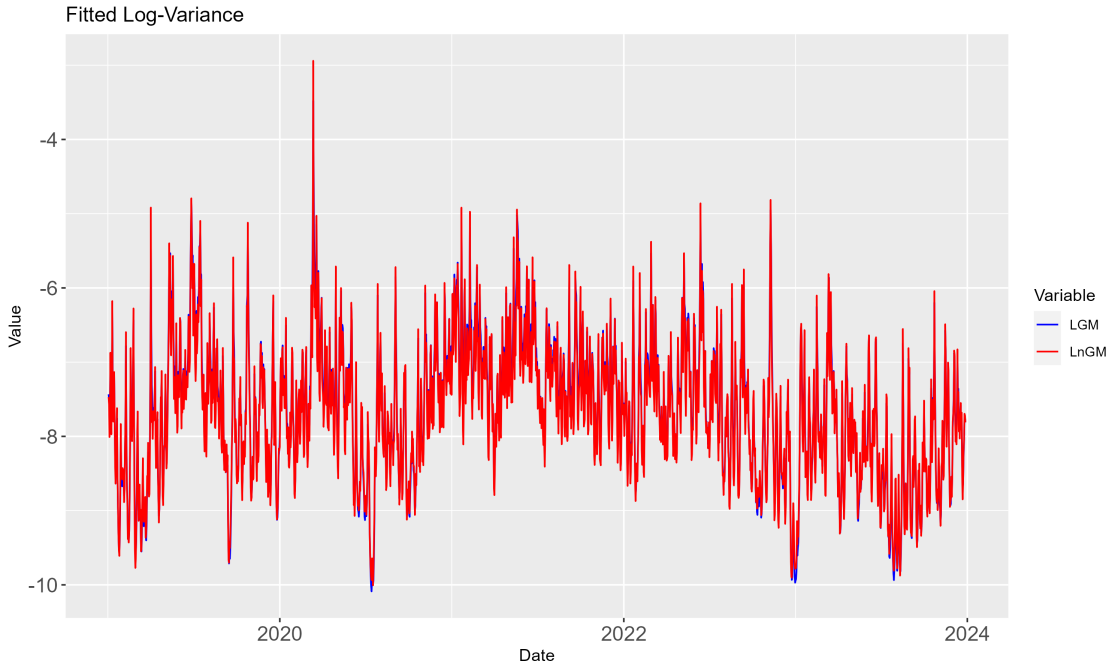

Figure 3: Estimated volatilities vs. absolute returns.

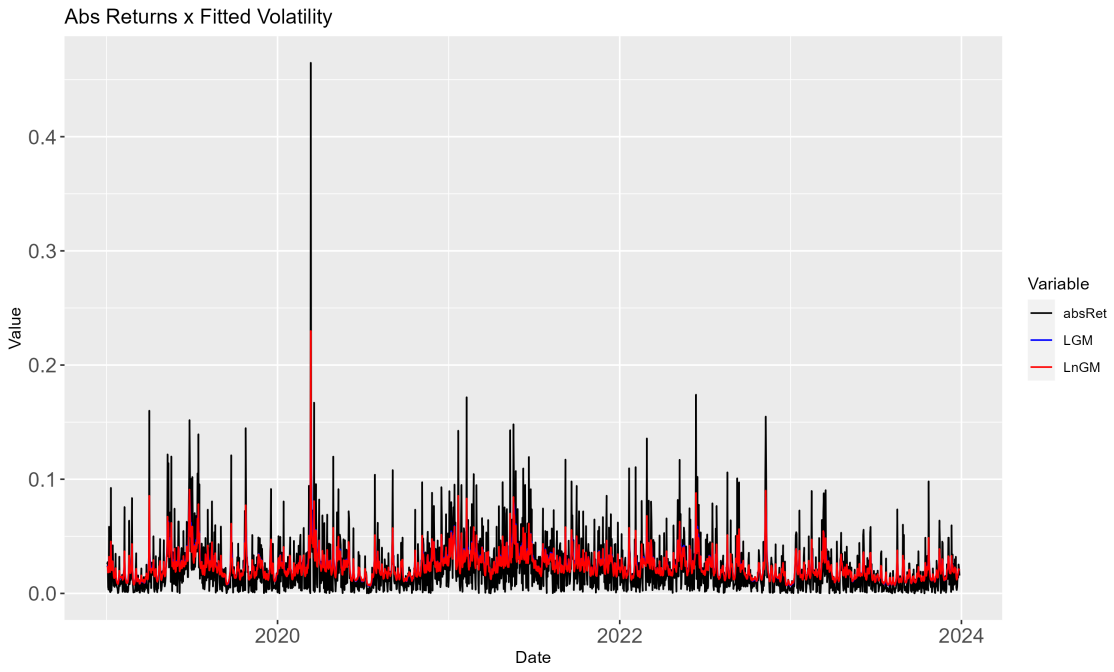

To verify the predictive gain of using a non-Gaussian specification for latent volatility, we present in-sample and out-of-sample error measures comparing the two specifications. Table 4 shows a comparison between the mean error (ME), Root Mean Square Error (RMSE) and Mean Absolute (MAE) error between the volatility measured by the two specifications compared to the absolute returns, which serves as a proxy for the true unobserved volatility. In this table the bold values represent the better fit, and we can observe that the use of non-Gaussian dynamics implies a better in-sample performance in the three analyzed measures, indicating adjustment gains in relation to the usual specification of a Gaussian dynamics for the SV model in the series and in the analyzed period.

Table 4: In-Sample Error Measures

	ME	**RMSE**	**MAE**
LGM	0.00149	0.01826	0.01258
LnGM	**0.00040**	**0.01654**	**0.01150**
ratio	3.725	1.10399	1.093913

We also evaluate out-of-sample predictive performance by making 1-, 5-, 10-, and 20-steps-ahead forecasts for the last sixty observations in the sample, again comparing forecasts for conditional volatility from the Gaussian and non-Gaussian specifications of the SV model

with the Bitcoin series absolute returns. Forecasts are made using a rolling sample using information up to the last date before the forecast. The results of this analysis are presented in Table 5.

Table 5: Out-of-Sample Error Measures

		ME	**RMSE**	**MAE**
1-step	LGM	−0.00507	0.01890	**0.01174**
	LnGM	**−0.00506**	**0.01889**	0.01175
	ratio	1.00197	1.00052	0.99914
5-steps	LGM	**−0.00524**	**0.01971**	0.01252
	LnGM	−0.00551	0.01976	**0.01247**
	ratio	0.95099	0.99746	1.00401
10-steps	LGM	0.00140	0.01868	0.01409
	LnGM	**0.00019**	**0.01863**	**0.01364**
	ratio	7.368421	1.00268	1.03299
20-steps	LGM	**−0.00340**	**0.01581**	0.01137
	LnGM	−0.00364	0.01586	**0.01134**
	ratio	0.93406	0.9968474	1.002646

We can observe that the best predictive performance results depend on the forecast horizon and the metric used. For the 1-step forecasting horizon, both models perform similarly in terms of ME and RMSE, but LnGM outperforms LGM slightly in terms of MAE. As the forecasting horizon increases to 5 steps, LGM shows slightly lower ME and RMSE compared to LnGM, but LnGM has a lower MAE. At the 10-step horizon, LGM displays a higher ME and RMSE compared to LnGM, while LnGM achieves a lower MAE. When forecasting 20 steps ahead, LGM again shows lower ME and RMSE, but LnGM demonstrates a lower MAE.

But in general, the values are similar between specifications in out-of-sample prediction. Note that this result is partly expected, since the use of a distribution with heavy tails for latent volatility is a way of explaining the greater amplitudes realized in the return series. These events can be associated with jump processes, which are usually assumed to be independent processes and in practice are very difficult to predict. In this way, a better in-sample fit of the model with NIG dynamics for latent volatility is an expected result, even though the predictive potential is similar to the standard model for out-of-sample forecasting, due to the independence in the jump processes.

5 Final remarks

In this work we present a new version of the stochastic volatility model where we assume that the dynamics of the latent log-variance process follows a non-Gaussian dynamic, approximated by a mixing process using an NIG-distributed driving noise. We make use of a novel approach for faster and computationally efficient inference of latent non-Gaussian models, proposed by Cabral et al. (2024), using variational Bayes and Laplace approximations as a Bayesian alternative to traditional Markov chain Monte Carlo (MCMC) algorithms.

This new representation of stochastic volatility models is interesting as it manages to capture behaviors that are not covered by the standard Gaussian process assumed in SV models. In particular, this mixture model representation is capable of approximating heavy-tailed processes in the dynamics of latent variables.

Latent processes with heavy tails are useful in analyzing the conditional volatility of financial series, since in these series the occurrence of extreme variations in prices and volatilities is more likely to occur than would be predicted by Gaussian densities. The use of mixing processes with heavy tails is a parsimonious way of approximating the dynamics of abrupt variations in the conditional variance process, being an alternative to the use of models with jumps.

We perform the empirical analysis by comparing Gaussian and non-Gaussian specifications of the SV model for Bitcoin return series. This asset is characterized by extreme variations in its prices, and relevant to abrupt changes in the conditional variance process (Chaim and Laurini (2018)), and thus is an excellent candidate for the application of SV models with non-Gaussian dynamics for the latent conditional variance. The results obtained in the empirical analysis indicate gains in adjusting this series with the use of the new specification proposed in the article, indicating the advantages of using this new methodology for this asset class.

References

Ando, T. (2006). Bayesian inference for nonlinear and non-Gaussian stochastic volatility model with leverage effect. *Journal of the Japan Statistical Society 36*(2), 173–197.

Barndorff-Nielsen, O. (1997). Normal inverse Gaussian distributions and stochastic volatility modelling. *Scandinavian Journal of Statistics 24*(1), 1–13.

Cabral, R., D. Bolin, and H. Rue (2022). Controlling the flexibility of non-Gaussian processes through shrinkage priors. *Bayesian Analysis (in print), arXiv preprint arXiv:2203.05510*.

Cabral, R., D. Bolin, and H. Rue (2024). Fitting latent non-Gaussian models using variational Bayes and Laplace approximations. *Journal of the American Statistical Association*, 1–13.

Chaim, P. and M. P. Laurini (2018). Volatility and return jumps in Bitcoin. *Economics Letters 173*, 158–163.

Chaim, P. and M. P. Laurini (2019). *Essays in Financial Econometrics*. PhD Dissertation in Applied Economics, Faculdade de Economia, Administração e Contabilidade de Ribeirão Preto, Ribeirão Preto.

Jordan, M. I., Z. Ghahramani, T. S. Jaakkola, and L. K. Saul (1999). An introduction to variational methods for graphical models. *Machine Learning 37*(2), 183–233.

Martino, S., K. Aas, O. Lindqvist, L. Neef, and H. Rue (2011). Estimating stochastic volatility models using integrated nested Laplace approximations. *European Journal of Finance 17*(7), 487–503.

Nacinben, J. P. C. d. S. M. and M. Laurini (2024). Multivariate stochastic volatility modeling via integrated nested Laplace approximations: A multifactor extension. *Econometrics 12*(1).

Phillip, A., J. Chan, and S. Peiris (2018). A new look at cryptocurrencies. *Economics Letters 163*, 6–9.

Rue, H., S. Martino, and N. Chopin (2009). Approximate Bayesian inference for latent gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society 71*(2), 319–392.

Rue, H., A. Riebler, S. H. Sørbye, J. B. Illian, D. P. Simpson, and F. K. Lindgren (2017). Bayesian computing with inla: a review. *Annual Review of Statistics and Its Application 4*, 395–421.

Taylor, S. (1986). *Modelling Financial Time Series*. Hoboken: John Wiley & Sons.

Wang, Y. and D. M. Blei (2019). Frequentist consistency of variational Bayes. *Journal of the American Statistical Association 114*(527), 1147–1161.

Zhang, C., J. Bütepage, H. Kjellström, and S. Mandt (2018). Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence 41*(8), 2008–2026.