

Article

Bayesian Inference for Long Memory Stochastic Volatility Models

Pedro Chaim ^{1,†}  and Márcio Poletti Laurini ^{2,*,†} 

¹ Department of Economics, Federal University of Santa Catarina, Centro Socioeconômico (CSE)–Bloco F, Florianópolis 88035-972, SC, Brazil; pedro.chaim@ufsc.br

² Department of Economics, School of Economics, Business Administration and Accounting at Ribeirão Preto, University of São Paulo, Av. dos Bandeirantes 3900, Ribeirão Preto 14040-905, SP, Brazil

* Correspondence: laurini@fearp.usp.br; Tel.: +55-16-3329-0867

† These authors contributed equally to this work.

Abstract: We explore the application of integrated nested Laplace approximations for the Bayesian estimation of stochastic volatility models characterized by long memory. The logarithmic variance persistence in these models is represented by a Fractional Gaussian Noise process, which we approximate as a linear combination of independent first-order autoregressive processes, lending itself to a Gaussian Markov Random Field representation. Our results from Monte Carlo experiments indicate that this approach exhibits small sample properties akin to those of Markov Chain Monte Carlo estimators. Additionally, it offers the advantages of reduced computational complexity and the mitigation of posterior convergence issues. We employ this methodology to estimate volatility dependency patterns for both the SP&500 index and major cryptocurrencies. We thoroughly assess the in-sample fit and extend our analysis to the construction of out-of-sample forecasts. Furthermore, we propose multi-factor extensions and apply this method to estimate volatility measurements from high-frequency data, underscoring its exceptional computational efficiency. Our simulation results demonstrate that the INLA methodology achieves comparable accuracy to traditional MCMC methods for estimating latent parameters and volatilities in LMSV models. The proposed model extensions show strong in-sample fit and out-of-sample forecast performance, highlighting the versatility of the INLA approach. This method is particularly advantageous in high-frequency contexts, where the computational demands of traditional posterior simulations are often prohibitive.



Citation: Chaim, Pedro, and Márcio Poletti Laurini. 2024. Bayesian Inference for Long Memory Stochastic Volatility Models. *Econometrics* 12: 35. <https://doi.org/10.3390/econometrics12040035>

Academic Editor: Luis Alberiko Gil-Alana

Received: 2 October 2024

Revised: 6 November 2024

Accepted: 19 November 2024

Published: 27 November 2024

Keywords: long memory; Gaussian Markov random fields; Laplace approximations; volatility forecasting

1. Introduction

Conditional volatility models are an essential tool in the fields of financial econometrics and risk management, and because of this importance, the literature related to the development of econometric models and inference methods to capture volatility patterns over time is wide-ranging. The main problem is that conditional volatility is a latent process, which adds a considerable degree of difficulty to traditional inference methods, such as maximum likelihood estimators.

Stochastic volatility (SV) models play a crucial role in understanding the dynamic behavior of financial time series, especially in capturing the inherent variability in returns over time. A common and effective approach in these models is to use a first-order autoregressive (AR(1)) structure for the latent log-variance, which is the natural logarithm of the conditional variance of returns. This AR(1) specification, as proposed in Taylor (1986), models the evolution of log-variance as a process that depends linearly on its past value plus a stochastic innovation. The formulation is particularly attractive due to its simplicity and ability to reproduce well-known stylized facts related to financial time series, such as volatility clustering, heavy tails, and the non-Gaussian nature of return distributions.



Copyright: © 2024 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/>).

The AR(1) structure assumes a mean-reverting process where volatility tends to revert to a long-term average level, capturing the short-memory property of financial volatility. This property aligns with the observation that shocks to volatility have a significant but temporary impact, decaying over time. Practical advantages of this structure lie in its tractability and the relative ease of estimation using state-space methods such as the Kalman filter or Bayesian inference techniques like Markov Chain Monte Carlo (MCMC).

Despite these advantages, the short-memory assumption inherent in AR(1) models limits their ability to capture certain empirical characteristics of volatility. Specifically, financial markets often exhibit persistent volatility patterns that decay more slowly for very distant lags than an AR(1) process can represent. This phenomenon, known as long-range dependence or long memory, implies that past volatility shocks can have prolonged effects that influence future volatility over extended periods. Short-memory models, which rely on exponential decay, fail to account for this feature, potentially leading to biased estimates and suboptimal forecasting performance when applied to assets with strong persistence in terms of volatility.

To address these limitations, long-memory structures for volatility, such as those based on fractional Brownian motion or fractional Gaussian noise, have been proposed. These models exhibit hyperbolic rather than exponential decay, capturing the persistence seen in empirical studies of financial time series. The Long-Memory Stochastic Volatility (LMSV) model is a prominent example, where the log-variance process follows a fractional integration parameter that allows for gradual decays in dependence. This modification enables the model to better represent the observed behavior of volatility, particularly in markets with persistent volatility clusters.

The importance of long-memory structures is reinforced by studies demonstrating their relevance for various asset classes and volatility measures. For instance, evidence from stock market indices, commodities, and even cryptocurrency markets indicates that incorporating long memory into volatility modeling can significantly improve in-sample fit and out-of-sample forecasting performance. Additionally, the use of long-memory models helps capture the slow mean reversion observed in realized volatility measures obtained from high-frequency data, aligning with empirical findings in financial econometrics (Christensen and Nielsen 2007; Maasoumi and McAleer 2008).

In summary, while first-order autoregressive structures are useful for modeling stochastic volatility due to their simplicity and alignment with short-term volatility dynamics, they fall short when addressing the long-term persistence seen in financial markets. Introducing long memory into volatility models provides a more flexible framework that aligns with the empirical properties of asset returns, enhancing both the modeling accuracy and the practical application of these models in risk management and financial forecasting.

The modeling of stochastic volatility with the use of a long-memory structure presents an additional complexity in relation to the models that assume an autoregressive structure. The key point is that the long-memory structure based on Fractional Brownian Motion/Fractional Gaussian Noise is a non-Markovian and non-semimartingale process, violating the usual assumptions regarding Markovian structures and the Martingale difference innovation processes used in the estimation of these models. As a direct example of the difficulties generated by the introduction of long-memory structures, note that a good portion of the methods used in the estimation of stochastic volatility models is based on a linear Gaussian state space representation, whose fundamental assumption is the use of a Markovian structure of dependence for latent states. Due to this difficulty, alternative forms were proposed for the estimation of this class of models.

In this paper, we employed an alternative Bayesian estimation for long-memory stochastic volatility models using an approximation of a Fractional Gaussian Noise process that allows us to represent this process as a Gaussian Markov Random Field, as proposed by in Sørbye et al. (2019). From this representation, we can use Integrated Nested Laplace Approximations (INLA) to perform the Bayesian estimation of parameters and latent log-variances (Rue et al. 2009). This formulation is interesting since it is computationally

efficient, allowing for the fast and accurate estimation of this class of models without necessitating the use of simulation-based methods, which can become computationally prohibitive when the sample size is very large; this, for example, occurs in the estimation using high-frequency data to realize volatility measurements. An application of this methodology in modeling interest rates in multifactor models can be found in ?, showing the good properties of this methodology for the approximation of long-memory structures.

The introduced methodology is also interesting in that it allows for modifications to be introduced to long-memory models involving variance. In this paper, we propose two extensions of the stochastic volatility model with long memory. The first extension is a two-factor model, where the first factor is the long-memory component, and the second component can be thought of as a smooth variation process in the long-term average of the volatility process, similar to the formulation of a Spline-GARCH model (e.g., [Engle and Rangel 2008](#)). The second modification is a stochastic volatility model with long memory that is adapted to the persistence patterns observed in high-frequency data. For this model, we employed a structure with two latent factors, with the first latent factor being a long-memory process based on a Fractional Gaussian Noise representatio, and the second latent factor allowing us to incorporate patterns of intraday seasonality into the stochastic volatility structure. This structure allows for the simultaneous estimation of the long-memory process and intraday seasonality pattern, avoiding the problems associated with the two-stage estimation proposed in [Deo et al. \(2006\)](#).

We can summarize the main research hypotheses as follows: the INLA methodology can estimate latent parameters and log-variances in LMSV models with a comparable accuracy to traditional MCMC methods, while significantly reducing computational complexity and alleviating convergence issues. By approximating the Fractional Gaussian Noise process as a Gaussian Markov Random Field, the approach is hypothesized to be computationally efficient, making it suitable for high-frequency financial data analysis. Furthermore, it is expected that multi-factor extensions of LMSV models using INLA will demonstrate a strong in-sample fit and out-of-sample forecasting performance, particularly in applications such as 5 min Bitcoin returns with integrated intra-day seasonal components.

Our simulation results suggest the INLA methodology can provide estimates of latent parameters and log-variances in LMSV models, which are comparable in terms of accuracy to traditional MCMC methods. The proposed extensions display good properties in terms of their in-sample adjustment and out-of-sample forecast, showcasing the flexibility of the INLA approach. The computational burden that arises from a large number of latent variables often prohibits the application of traditional posterior simulation methods in high-frequency settings. This is especially true for LMSV models. The LMSV specification we apply to 5 min Bitcoin returns incorporates the simultaneous estimation of additive intra-day seasonal components and performs well in the construction of realized volatility measures.

This paper is structured as follows. Section 3 exposes the formulation proposed for the stochastic volatility models and discusses the estimation methods. Section 2 presents a literature review of the estimation of SV models. In Section 4, we present a Monte Carlo study of the small sample properties of the estimators. In Section 5, we apply the methodology to the daily returns of major cryptocurrencies and S&P 500 Index, then compare the in-sample fit and out-of-sample forecast performance to popular alternatives. The method's application to high-frequency Bitcoin data is presented in Section 6. Section 7 concludes the paper.

2. Literature Review

The literature on econometric models of conditional volatility can be divided into two main classes of models based on the form of treatment applied for the latent volatility process. The first important class is the class of observation-driven models, which considers that the process of latent variance can be approximated by a structure that depends only on observable processes. In this class of models, the main reference is the GARCH (Generalized

Autoregressive Conditional Heteroskedasticity) family of models, introduced in [Engle \(1982\)](#) and generalized in [Bollerslev \(1986\)](#), and later extended in several other directions.

In this structure, latent variance is considered a deterministic function of the past variance itself and the past squared returns or residues. A general framework for observation-driven process modeling is the generalized score model class proposed in [Creal et al. \(2013\)](#), which uses the model's past score as the update mechanism for time-varying model parameters. The use of this observation-driven structure is interesting because, considering that the variance process depends only on observable components in the current period, it avoids the inference problems generated by the presence of latent variables in the model, allowing for the direct use of maximum likelihood estimators.

The second class of models is based on the treatment of conditional volatility as a truly stochastic dependent process, with the main reference in discrete time being the so-called log-normal model of stochastic volatility (SV) introduced in [Taylor \(1986\)](#). In this model, the dynamics of the conditional log-variance are derived through a first-order autoregressive process. This structure is very convenient since this formulation is consistent with the main stylized facts of financial series, such as volatility clusters, heavy tails, and the non-Gaussianity of returns, and avoids the artificial assumption that volatility is deterministic.

However, assuming that the process of latent volatility is stochastic adds significant difficulties in the inference procedures for parameters and for the filtering of the conditional volatility itself. When assuming a stochastic structure, it is not possible to directly use maximum likelihood estimators, since these require the marginalization of each unobserved volatility, thus necessitating the calculation of a multiple integral with a dimension equal to the sample size, which is not computationally feasible. Due to this difficulty, a series of frequentist and Bayesian methods was created to estimate stochastic volatility models. The estimation method proposed in [Taylor \(1986\)](#) is based on the method of moments, which does not allow for an estimation of the latent volatility process itself, only the parameter vector of the model.

Among the frequentist methods, the main reference is the Quasi-Maximum Likelihood method based on the decomposition of the forecast error through the Kalman Filter independently proposed in [Harvey et al. \(1992\)](#) and [Nelson \(1988\)](#). In this structure, the log of the squared returns is composed of a mean component plus a latent autoregressive process, which is the log-variance of the returns, formulated as a state-space process that allows for the use of the Kalman Filter to filter this latent component. As this estimation is based on a linearization of the model, the estimator is approximated and is therefore a Quasi-Maximum Likelihood estimator. Although this estimator is computationally simple, its performance may be sub-optimal due to the bias introduced by the linearization of the model, as discussed in [Andersen and Sørensen \(1996\)](#). To circumvent this limitation, other frequentist estimation methods use maximum simulated likelihood ([Sandmann and Koopman 1998](#)), efficient method of moments ([Andersen et al. 1999](#)), empirical likelihood, and minimum discrepancy ([Laurini and Hotta 2017](#)), among other possible treatments.

Several Bayesian estimation methods have been proposed for the class of stochastic volatility models. Bayesian estimation is particularly interesting for the problem at hand because latent processes can be treated as additional components of the model, and thus additional parameters to be estimated. This allows for the use of traditional methods of Bayesian inference without additional complications. Among the main works proposing the use of Markov Chain Monte Carlo (MCMC) for estimating SV models, we have [Kim et al. \(1998\)](#), [Chib et al. \(2002\)](#) and [Asai \(2005\)](#), and as examples of recent developments, we have [Kastner and Frühwirth-Schnatter \(2014\)](#) and [Kastner \(2019\)](#).

[Martino et al. \(2011\)](#) proposed an alternative form of the Bayesian estimation of SV models based on the Integrated Nested Laplace Approximations (INLA) framework of [Rue et al. \(2009\)](#). The INLA methodology allows for one to perform Bayesian estimation for parameters and latent factors for a wide class of models that can be represented or approximated by Gaussian Markov Random Fields. Because INLA is based on an analytical approximation, it does not require simulation procedures and thus is not affected by the

chain convergence issues that sometimes effect MCMC algorithms. In addition, these methods are very accurate and are generally computationally fast and efficient, since, in conjunction with the INLA methodology, it is possible to use a series of sparse matrix representations to represent the Gaussian Markov Random Field. The properties of this method for the estimation of SV models are analyzed in [Ehlers and Zavallos \(2015\)](#), which shows that the method is very accurate and has good properties for the calculation of measures derived from conditional volatility, estimated as the Value at Risk. A multifactor formulation of multivariate stochastic volatility models is presented in [Nacinben and Laurini \(2024a\)](#).

An important point is that most of the models proposed for stochastic volatility modeling are based on a short-memory structure for the dependency process, assuming a mean reversion velocity compatible with an autoregressive process. An alternative possibility is to assume that the process of dependence on volatility is provided by a long-memory process, also known as a long-range dependence process. In this structure, the most used form is an innovation structure based on a discrete version of the fractional Brownian Motion, a representation based on a fractional Gaussian Noise, as described in [Taqqu \(2003\)](#). In this case, the innovations of the dependency process follow a hyperbolic decay process, rather than the standard exponential decay of short-memory models. References to the use of long-memory structures in stochastic volatility models can be found in [Harvey \(1998\)](#), [Breidt et al. \(1998\)](#), and [Hurvich and Soulier \(2009\)](#). The main model used is the Long-Memory Stochastic Volatility (LMSV), where the log variance follows a long-memory process.

The use of long-memory models for variance may be interesting as there is empirical evidence showing that this structure may be appropriate for classes of assets and measures of volatility. [Christensen and Nielsen \(2007\)](#), [Scharth and Medeiros \(2009\)](#), [Hillebrand and Medeiros \(2016\)](#), [Shackleton et al. \(2008\)](#), and [Fleming and Kirby \(2011\)](#) discuss the importance of a long memory in conditional stock market volatility. Another important point is the impact of long memory on the calculation of realized volatility measurements using high-frequency data, as discussed in [Maasoumi and McAleer \(2008\)](#), [McAleer and Medeiros \(2008\)](#), and [Lieberman and Phillips \(2008\)](#). There is recent interest in the use of these models because of the ample evidence of the presence of long-range dependence features in cryptocurrency markets, as shown by [Mensi et al. \(2019\)](#), [Phillip et al. \(2019\)](#), [Bouri et al. \(2019\)](#), and [Chaim and Laurini \(2019\)](#), among other works.

Many methods have been proposed for the estimation of LMSV models. [Harvey \(1998\)](#) and [Breidt et al. \(1998\)](#) employed a spectral approximation for the likelihood function, while [Arteche \(2004\)](#) specified a semi-parametric form based on a local Whittle estimator. Different formulations can be used using an approximation of the state space representation through a truncated Autoregressive Fractionally Integrated Moving Average (ARFIMA) process, as used in [Chan and Palma \(1998\)](#), [Basak et al. \(2001\)](#) and [Ferraz and Hotta \(2007\)](#) in frequentist formulations, and [Chan and Petris \(2000\)](#) for the Bayesian estimation in a Markov Chain Monte Carlo algorithm. Comparative reviews of estimation methods of stochastic volatility models with long memory can be found in [Perez and Ruiz \(2001\)](#) and [Crato and Ray \(2002\)](#).

3. Long-Memory Stochastic Volatility Models

The standard stochastic volatility (SV) model, introduced in [Taylor \(1986\)](#), is a log-normal mixture model in which the log normal variance component follows a stationary autoregressive process of the first order. Let $\{r_t\}_{t=1}^N$ be an equally spaced sequence of asset returns; then, an SV model can be written as

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

$$h_t = \mu + \phi(h_{t-1} - \mu) + z_t, \quad z_t \sim N\left(0, ((1 - \phi^2)\tau_h)^{-1}\right). \quad (2)$$

This model has three parameters that dictate the AR(1) dynamics of latent log-variance h_t . The long-run mean is μ , autoregressive persistence depends on ϕ , and the latent log-

variance negatively depends on the (marginal) precision τ_h . Due to the internal INLA conventions, in this paper, we report the marginal precision of log-variance h_t , τ_h . It is more common to speak in terms of the precision of the i.i.d. error term z_t , which can be easily obtained by the relation $\tau = \tau_h(1 - \phi^2)$.

Because log-variance is not an observed variable, maximum likelihood estimation requires evaluating a multiple integral with a dimension equal to the sample size, which is often computationally unfeasible. By employing the method of moments, originally proposed by Taylor (1986), one can estimate parameters μ , ϕ , and τ , but is unable to recover latent log-variance h_t . A traditional frequentist solution is the Quasi-Maximum Likelihood (QML) method, based on the forecast error decomposition obtained using Kalman Filter (Harvey 1989). Here, log-squared returns are linearized and decomposed into a mean. Additionally, a first-order autoregressive process is used, such that the model has a linear state space representation, and is then amenable to Kalman filtering. Alternative frequentist methods which can circumvent these issues include simulated maximum likelihood (Sandmann and Koopman 1998), the efficient method of moments (Andersen et al. 1999), empirical likelihood, and generalized minimum discrepancy (Laurini and Hotta 2017), to mention a few.

There are several Bayesian estimation methods for the class of stochastic volatility models. Because latent variables are treated as additional parameters, traditional Bayesian inference methods are directly applicable to SV models. Markov Chain Monte Carlo algorithms are commonly used for posterior sampling, e.g., Gamerman and Lopes (2006) and Johanes and Polson (2005). It is known that due to the high correlation within latent variables and between latent variables and deep parameters, MCMC methods can require long chains to provide a good posterior characterization of SV models (Kastner and Frühwirth-Schnatter 2014; Gong and Stoffer 2020).

Although estimation using posterior simulation methods is feasible in a wide range of practical applications, MCMC methods still can be computationally demanding, especially models involving latent variables, whose computational burden increases more than proportionally with larger sample sizes. Martino et al. (2011) show how an SV model can be estimated using Integrated Nested Laplace Approximations (INLA), a class of methods introduced by Rue et al. (2009), which can be applied to the Bayesian estimation of parameters and latent factors of models that can be represented by, or approximated by, a Gaussian Markov Random Field (GMRF). Because INLA calculations are analytic, this avoids issues inherent to posterior simulation MCMC methods, such as slow chain convergence.

Following Martino et al. (2011), if we assign mean parameter μ a Gaussian prior with zero mean and a large known variance, the standard stochastic volatility model (1) and (2) can be seen as a latent Gaussian model with latent field x :

$$x = \{h_1, \dots, h_n, \mu\} \sim N(\mathbf{0}, \mathbf{Q}^{-1}(\theta_1)), \quad (3)$$

where $\theta_1 = \{\tau_h, \phi\}$ are the parameters driving the log-variance. The $|x| = n$ -dimensional latent Gaussian field is partially observed through the n_d conditionally independent data $r = \{r_1, \dots, r_{n_d}\}$ with the following likelihood:

$$\pi(r|x, \theta_2) = \prod_{t=1}^{n_d} \pi(r_t|h, \theta_2),$$

where θ_2 are parameters of the ε_t return process, which, in our case, is a precision of one, since $\varepsilon_t \sim N(0, 1)$.

Let $\theta = \{\theta_1, \theta_2\}$. The main goal in estimating SV models is to evaluate the marginal distributions

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

The INLA procedure is a computationally efficient method to compute these marginal distributions. At its core is a very fast Gaussian approximation to densities of the following form:

$$\pi(\mathbf{x}|\mathbf{r}, \boldsymbol{\theta}) \propto \exp\left\{-\frac{1}{2}\mathbf{x}^T \mathbf{Q} \mathbf{x} + \sum g_t(h_t)\right\},$$

where $g_t(h_t) = \log \pi(r_t|h_t, \boldsymbol{\theta})$. The Gaussian approximation $\bar{\pi}_G(\mathbf{x}|\mathbf{r}, \boldsymbol{\theta})$ is obtained by matching the curvature at the posterior mode value \mathbf{m} , which is computed iteratively using a Newton–Raphson algorithm:

$$\bar{\pi}_G(\mathbf{x}|\mathbf{r}, \boldsymbol{\theta}) = K_1 \exp\left\{-\frac{1}{2}(\mathbf{x} - \mathbf{m})^T (\mathbf{Q} + \text{diag}(\mathbf{c}))(\mathbf{x} - \mathbf{m})\right\},$$

where K_1 is a normalizing constant, and the vector \mathbf{c} is a correction term given by second-order terms in the Taylor expansion of $\sum g_t(h_t)$ at the modal value \mathbf{m} .

3.1. Long-Memory Stochastic Volatility Models

Although SV models capture key aspects of returns series such as there being almost no autocorrelation in terms of levels and excess kurtosis, for some assets, the decay in the autocorrelation of distant lags of squared returns can be slower than implied by the autoregressive structure in Equation (2).

An alternative is to model log-variance as a stationary process with long memory. Following [Beran \(2017\)](#), a weakly stationary process has long memory if its autocovariance function $\gamma(k)$ for distant lags k satisfies

$$\gamma(k) \sim C_1 k^{2(H-1)}, \quad \text{as } k \rightarrow \infty,$$

for $C_1 > 0$, with $1/2 < H < 1$. Or, equivalently, in the frequency domain, if its spectral density $f(\omega)$ for frequencies $\omega \in [-\pi, \pi]$ close to zero satisfies

$$f(\omega) \sim C_2 |\omega|^{-2H+1}, \quad \text{as } \omega \rightarrow 0,$$

for $C_2 > 0$ and $1/2 < H < 1$.

If the log-variance component h_t has long-range dependence properties, then we have the long-memory stochastic volatility model (LMSV), as introduced by [Harvey \(1998\)](#) and [Breidt et al. \(1998\)](#).

The traditional approach when introducing long-range dependence to SV models involves characterizing latent log-variance dynamics as an ARFIMA process. In its simplest form, log-variance h_t follows an ARFIMA(0,d,0) and this model can be written as

$$r_t = \exp(h_t/2)\varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (4)$$

$$(1 - B)^d h_t = \eta_t, \quad \eta_t \sim N(0, \tau_\eta^{-1}), \quad (5)$$

where B is the backshift operator and parameter $d \in (-0.5, 0.5)$ determines the fractional integration order.

Because long-memory processes do not have a finite state space representation, the estimation of LMSV models is not straightforward. Notice that many of the presented methods for estimating SV models take advantage of some sort of state-space representation of the model. Estimations of LMSV models have traditionally been carried out in the frequency domain using spectral likelihood estimators ([Breidt et al. 1998](#)), and in the time domain using simulation-based Bayesian methods and Quasi-Maximum likelihood estimators. The latter approaches are often based on a truncated state-space representation of the long-memory process, as discussed in [Chan and Palma \(1998\)](#), [Chan and Petris \(2000\)](#) and [Ferraz and Hotta \(2007\)](#). A popular approach is to deal with the approximate representation of an ARFIMA process as a long lagged Autoregressive Moving Average

(ARMA) model. A simple representation of this is the expansion of an ARFIMA(0, d , 0) through an autoregressive process in the following form:

$$\sum_{i=0}^P \Pi_i h_{t-i} = \eta_t \quad (6)$$

with $\Pi_0 = 1$, $\Pi_j = [(j - d - 1)/j]\Pi_{j-1}$, P as the order of approximation, h_t as the latent process, and η_t as a white noise process. This representation is quite interesting since it is sufficient to rewrite the representation in state space as that of an AR(P) model, in this case using the coefficients determined by the above expansion. The usual quasi-maximum likelihood and MCMC estimators can easily be adapted to this representation.

The approach we use here involves introducing long-range dependence through a fractional Gaussian noise (fGn) process rather than an ARFIMA. Conceptually, an ARFIMA comes from the fractional difference of a discrete ARMA process, while an fGn comes from the fractional differentiation of a continuous Brownian motion (Hosking 1981). The two processes are closely related, especially if the autoregressive and moving average order of the ARFIMA are 0. The relationship between the Hurst exponent H and the fractional integration order d is such that $H = d + 0.5$.

The LMSV specification we explore here can be written as

$$r_t = \exp(h_t/2)\varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (7)$$

$$h_t = \mu + x_t, \quad (8)$$

where μ is a mean parameter, and x_t is an fGn with parameters τ and H .

Specifically, $\mathbf{x}' = (x_1, \dots, x_n)$ is a zero-mean multinormal vector such that

$$\mathbf{x} \sim N(0, \tau^{-1}\Sigma). \quad (9)$$

The covariance matrix Σ has a Toeplitz structure, with the first-row elements given by the autocorrelation function:

$$\gamma_{fGn}(k) = \frac{1}{2} \left(|k-1|^{2H} - 2|k|^{2H} + |k+1|^{2H} \right), \quad (10)$$

where $k = 0, 1, \dots, n-1$. Note that the autocorrelation function is indeed a form of hyperbolic decay, as $\gamma_{fGn}(k) \sim H(2H-1)k^{2(H-1)}$ when $k \rightarrow \infty$. We will sometimes write $x_t \sim fGn(H, \tau)$ as a shorthand notation when referring to the structure described above.

3.2. Gaussian Markov Random Field Approximation of Fractional Gaussian Noise

Sørbye et al. (2019) took advantage of the known relationship between long memory and cross-sectional aggregation (Granger 1980; Beran et al. 2010) to construct a Gaussian Markov Random Field-based approximation of a fractional Gaussian noise model from weighted sums of independent AR(1) components. They proposed matching the autocorrelation function of this composite autoregressive process to the autocorrelation (10) of a fGn.

Following Sørbye et al. (2019), consider m independent AR(1) processes

$$z_{j,t} = \phi_j z_{j,t-1} + \xi_{j,t}, \quad j = 1, \dots, m, \quad t = 1, \dots, n,$$

where $0 < \phi_j < 1$ is the first-order autoregressive parameter of the j -th process. Also, let $\{\xi_{j,t}\}_{j=1}^m$ be zero-mean independent Gaussian shocks with variance $\sigma_{\xi,j}^2 = 1 - \phi_j^2$. Define the cross-sectional aggregation of the m processes as

$$\bar{z}_m = \sigma \sum_{j=1}^m \sqrt{w_j z^{(j)}}, \quad (11)$$

where $z^{(j)}$ denotes $(z_{j,1}, z_{j,2}, \dots, z_{j,n})$ and the weights w_j sum to one. [Haldrup and Valdés \(2017\)](#) studied the finite sample properties of similar aggregations of AR(1) processes.

The autocorrelation function of (11) is simply

$$\gamma_{z_m}(k) = \sum_{j=1}^m w_j \phi_j^{|k|}, \quad k = 0, 1, \dots, n-1. \quad (12)$$

The idea put forward in [Sørbye et al. \(2019\)](#) is to fit weights $w = \{w_j\}_{j=1}^m$ and coefficients $\phi = \{\phi_j\}_{j=1}^m$ such that the autocorrelation function (12) of the composite AR(1) process matches the autocorrelation function of a true fGn process, as in Equation (10). Values of (w, ϕ) are obtained by minimizing the squared error

$$(w, \phi)_H = \operatorname{argmin}_{(w, \phi)} \sum_{k=1}^{k_{\max}} \frac{1}{k} \left(\gamma_{z_m}(k) - \gamma_{fGn}(k) \right)^2, \quad (13)$$

where k_{\max} is an arbitrary upper limit to the number of lags included. Since the squared error is weighted by k , persistence at distant lags has little impact on the objective function (13).

With this approximation of the fGn process, we can represent the LMSV model (7) and (8) as a latent GMRF using the INLA method, as shown in [Martino et al. \(2011\)](#). Following the recommendations in [Sørbye et al. \(2019\)](#), we use a third-order approximation to represent the fGn process. Further details on the implementation can be found in [Sørbye et al. \(2019\)](#).

3.3. Alternative Specifications

The INLA methodology is not restrictive and SV models can be easily augmented with additional latent factors. As long as the affine structure of log-variance is preserved, the model has a GMRF representation similar to Equation (3), and can be estimated in the same manner. To showcase this flexibility, we consider alternative formulations for both SV and LMSV models, in which log-variance is composed of a smooth spline trend in addition to either the AR(1) or fGn processes. These specifications are similar to the Spline-GARCH model of [Engle and Rangel \(2008\)](#), where conditional volatility is subject to low-frequency variations which work as a time-varying long-run average.

The AR(1) Spline-SV model can be written as

$$r_t = \exp((h_t + s_t)/2) \varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (14)$$

$$h_t = \phi h_{t-1} + z_t, \quad z_t \sim N(0, ((1 - \phi^2) \tau_h)^{-1}), \quad (15)$$

where s_t follows a second-order random walk, which is constructed by assuming independent second-order increments with precision τ_s . That is,

$$\Delta^2 s_t \sim N(0, 1/\tau_s). \quad (16)$$

The use of a second-order random walk structure to represent a smoothly varying spline process is common in statistics (see, for example, [Lindgren and Rue 2008](#)). Note that an analogous relationship exists between the popular Hodrick–Prescott (HP) filter and spline representations, as discussed in [Harvey and Jaeger \(1993\)](#), [Harvey and Trimbur \(2008\)](#) and [Paige and Trindade \(2010\)](#).

Since we use an unrestricted specification for this component, i.e., the component does not sum to zero, it can be used to capture a long-run average of a log-variance that is varying smoothly. Hence, in this specification, it is not necessary to include the μ mean parameter, which is already captured by the dynamics of this process.

Therefore, the AR(1) Spline model has three parameters: the autoregressive coefficient ϕ , the marginal precision of log-volatility τ_h , and the precision of the spline (trend) component τ_s . More details on the specification of the second-order random walk model can be found in the Appendix A.

Likewise, consider the LMSV Spline model

$$r_t = \exp((h_t + s_t)/2)\varepsilon_t, \quad \varepsilon_t \sim N(0,1), \quad (17)$$

$$h_t = x_t, \quad x_t \sim fGn(H, \tau_h), \quad (18)$$

$$\Delta^2 s_t \sim N(0, 1/\tau_s). \quad (19)$$

This model has three parameters: the Hurst exponent H , which determines the temporal persistence of the fGn process, the marginal precision of fGn component τ_h , and the precision τ_s of the second-order random walk spline.

Another alternative specification explored here is the inclusion of an additive seasonal component $seas_t$ in conditional log-variance dynamics. The seasonal component $seas_t$ has a set periodicity M and is restricted such as the sum of the m individual components is zero. This model can be written as

$$r_t = \exp(h_t + seas_t)/2)\varepsilon_t, \quad \varepsilon_t \sim N(0,1), \quad (20)$$

$$h_t = \mu + x_t, \quad x_t \sim fGn(H, \tau), \quad (21)$$

$$seas_t + seas_{t-1} + \dots + seas_{t-m+1} \sim N(0, 1/\tau_{seas}). \quad (22)$$

This seasonal component is of special interest when we are dealing with high-frequency data. As discussed in Deo et al. (2006), there is evidence of periodic patterns in intraday returns volatility. The periodicity M , and thus the number of seasonal components, depend on the aggregation used to compute intraday returns. For example, in Section 6 we use prices 5 min apart, and this defines the number of individual seasonal components.

The direct incorporation of this periodicity structure is advantageous since it avoids the use of multi-stage methods for the estimation of the LMSV model for high-frequency data. For example, in Deo et al. (2006) the periodicity pattern is estimated in the first stage, and then returns adjusted for this pattern are modeled as an LMSV process. Our formulation allows for the simultaneous estimation of these components, avoiding the problems associated with estimation at different stages.

A flowchart presenting a summary of the steps required to apply the method is shown in Figure 1.

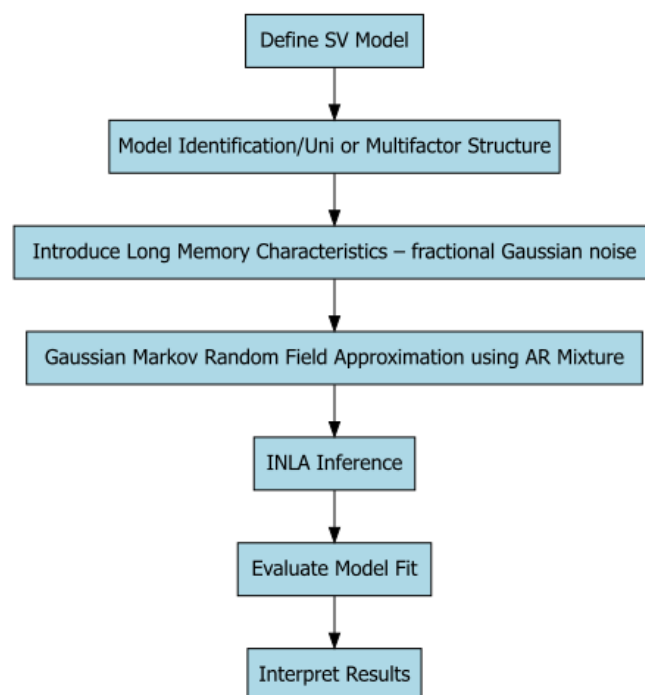


Figure 1. Flowchart of stochastic volatility models' implementation.

4. Monte Carlo Simulation Study

We performed Monte Carlo experiments to evaluate the small sample properties of the INLA fGn approximation and the estimation of LMSV models described in Section 2. INLA estimates were compared to a traditional approach of employing an MCMC posterior sampling algorithm based on a truncated state space representation of long-memory processes, as proposed by Chan and Petris (2000). The long-memory process is approximated as an AR(20). We sample parameters H and τ and the latent volatility h_t using a Random Walk Metropolis scheme, while parameter μ has a conjugate sampler. For a detailed description of this method, see Chan and Petris (2000) or Ferraz and Hotta (2007). Each MCMC estimation was based on a chain of 20,000 samples, after discarding the first 4000. Longer chains did not show significant gains in terms of MCMC accuracy performance in this specific application.

Three prior specifications/estimation methods were compared. The first two had the same prior configuration. For the MCMC estimation and the baseline INLA (INLA 1) implementation, we considered a Gaussian prior with a mean -7 and precision of 0.01 for the average volatility parameter μ , Gamma prior with the parameters $(3, 0.01)$ for marginal precision τ , and Gaussian with mean 0.9 and a precision of 0.01 for parameter H . We also considered a second prior structure (INLA 2), which is based on the penalized complexity priors framework (pc-priors), as introduced by Simpson et al. (2017). These define pc-priors as invariant to reparameterization, having excellent robustness properties, and allowing for a direct comparison between different conditional models. Sørbye and Rue (2018) showed how to construct such priors for precision τ and persistence H parameters describing an fGn process. See specific details in the Appendix A.

We chose two parameter configurations for the data-generating processes describing different levels of volatility. The first one, $\theta_1 = \{\mu = -9.90, H = 0.931, \tau = 0.962\}$, represents the standard market conditions. The second parameter vector, $\theta_2 = \{\mu = -7.182, H = 0.938, \tau = 0.294\}$, represents more volatile conditions, such as cryptocurrency markets. For both parameter vectors, we simulated 1000 samples from an LMSV model with 500 and 1000 observations. In each experiment, we compared the point estimates to the true parameter value in terms of mean error (ME), root mean squared error (RMSE), and mean absolute error (MAE). Tables 1 and 2 present the results for each parameter vector, with sample sizes of 500 and 1000. It can be observed that the estimation of INLA presents better results in terms of mean error (ME), root mean squared error (RMSE), and mean absolute error (MAE) for parameters μ and τ for the two parameter configurations and sample sizes. The different parameter configurations used in the estimation of the INLA do not seem to have relevant impacts on the estimation of these parameters. When looking at the Hurst exponent H , we find rather mixed results. For the first parameter vector in Table 1, INLA fares better with respect to all three accuracy measures. Interestingly, the opposite is true for the second parameter configuration in Table 2, whose values represent more turbulent market conditions.

Table 1. Monte Carlo Experiment LMSV Model—First parameter set ($\mu = -9.9, H = 0.931, \tau = 0.962$)—Sample sizes 500 and 1000. Note—INLA 1—Gamma/Gaussian Priors. INLA 2—Penalized Complexity Priors.

| | | Sample Size 500 | | | Sample Size 1000 | | |
|-------|--------|------------------------|--------|--------|------------------------|--------|--------|
| | | ME | RMSE | MAE | ME | RMSE | MAE |
| μ | MCMC | -0.0640 | 0.2280 | 0.1840 | -0.0380 | 0.1310 | 0.1050 |
| | INLA 1 | 0.0430 | 0.2140 | 0.1730 | 0.0090 | 0.1130 | 0.0900 |
| | INLA 2 | -7.00×10^{-4} | 0.2010 | 0.1630 | -5.20×10^{-4} | 0.1130 | 0.0900 |
| H | MCMC | -0.1220 | 0.1530 | 0.1240 | -0.0430 | 0.0700 | 0.0520 |
| | INLA 1 | -0.0810 | 0.1010 | 0.0818 | -0.0450 | 0.0610 | 0.0470 |
| | INLA 2 | -0.0840 | 0.1100 | 0.0864 | -0.0380 | 0.0620 | 0.0450 |
| h_t | MCMC | 3.35×10^{-4} | 0.0031 | 0.0020 | -3.50×10^{-4} | 0.0020 | 0.0010 |
| | INLA 1 | -2.20×10^{-4} | 0.0020 | 0.0010 | -4.00×10^{-4} | 0.0020 | 0.0010 |
| | INLA 2 | -4.40×10^{-4} | 0.0020 | 0.0010 | -4.00×10^{-4} | 0.0020 | 0.0010 |

Table 2. Monte Carlo experiment LMSV model; second parameter set ($\mu = -7.182$, $H = 0.938$, $\tau = 0.294$); sample sizes 500 and 1000. Note—INLA 1; Gamma/Gaussian priors. INLA 2 penalized complexity priors.

| | | Sample Size 500 | | | Sample Size 1000 | | |
|--------|--------|-----------------|-------|--------|------------------|-------|-------|
| | | ME | RMSE | MAE | ME | RMSE | MAE |
| μ | MCMC | −0.050 | 0.379 | 0.306 | −0.031 | 0.211 | 0.168 |
| | INLA 1 | 0.023 | 0.342 | 0.276 | 0.016 | 0.188 | 0.150 |
| | INLA 2 | 0.027 | 0.348 | 0.281 | 0.0174 | 0.206 | 0.161 |
| H | MCMC | −0.042 | 0.066 | 0.048 | −0.002 | 0.033 | 0.025 |
| | INLA 1 | −0.052 | 0.068 | 0.053 | −0.026 | 0.049 | 0.031 |
| | INLA 2 | −0.045 | 0.065 | 0.048 | −0.022 | 0.050 | 0.030 |
| τ | MCMC | 0.258 | 0.319 | 0.260 | 0.435 | 0.475 | 0.435 |
| | INLA 1 | 0.250 | 0.302 | 0.255 | 0.133 | 0.182 | 0.150 |
| | INLA 2 | 0.183 | 0.239 | 0.193 | 0.102 | 0.164 | 0.125 |
| h_t | MCMC | −0.002 | 0.020 | 0.012 | −0.002 | 0.019 | 0.012 |
| | INLA 1 | −0.004 | 0.018 | 0.011 | −0.003 | 0.018 | 0.011 |
| | INLA 2 | −0.004 | 0.018 | 0.0116 | −0.003 | 0.018 | 0.011 |

We also compared the performance of these methods in estimating the paths of latent log-variance h_t . In this case, we compute the ME, RMSE, and MAE between the true and estimated log-variance for each replication of the experiment, and report the average values between all simulations in Tables 1 and 2. The performances of the MCMC and INLA methods in estimating latent log-variance are quite similar, with a marginally better estimation performance being obtained using INLA for both parameter settings and sample sizes. This result is very relevant, since the main objective in risk measurements is to accurately recover the true volatility of the process, which is the necessary input for the calculation of risk management measures such as Value at Risk. In this regard, the proposed methodology has a satisfactory performance for empirical applications, showing the validity of the method for real problems of risk measurement and management.

A major advantage of the INLA approach comes from its lower computational cost. We performed the MCMC estimations using a compiled C++ code though the R package nimble and INLA with the r-inla¹ package. The average time required for MCMC estimations with a sample size of 1000 was 38.5 s, increasing to 106.8 s if we include compilation time. Each INLA estimation was performed, on average, in 4.805 s. This very simple comparison suggests an increase in speed of about eight times. The results of the Monte Carlo experiment indicate that the INLA estimation presents a performance that is generally equal or superior to the MCMC method. Thus, we have evidence of INLA's good performance in LMSV model estimations, in terms of both statistical performance and computational time.


5. Application to Daily Data

Analyses of cryptocurrencies' trading data have attracted increased research attention in recent years. Many studies point to observed patterns of long-range dependence in the volatility of cryptocurrencies' returns (see, for example, [Urquhart 2016](#); [Bariviera et al. 2017](#); [Mensi et al. 2019](#); [Cheah et al. 2018](#); [Phillip et al. 2019](#); [Chaim and Laurini 2019](#); [Bouri et al. 2019](#); [Charfeddine and Maouchi 2019](#); [Tan et al. 2021](#); [Assaf et al. 2022](#); [Arouxet et al. 2022](#)). This mounting evidence motivated us to choose the two major cryptocurrencies as case studies to acquire daily data: Bitcoin (btc) and Ethereum (eth). Our sample consists of daily data from 1 January 2020 until 31 July 2024, which corresponds to 1672 observations. We also repeated our analysis for daily returns of the S&P500 (sp500) index within the same time period.

Table 3 reports descriptive statistics of the daily log returns and absolute returns of Bitcoin and Ethereum, as well as the S&P500 index. The trajectories are depicted in Figure 2.

Table 3. Descriptive statistics—Bitcoin (btc), Ethereum (eth), and S&P 500 (sp500).

| | Returns | | | Abs Returns | | |
|-----------|---------|---------|---------|-------------|--------|--------|
| | Btc | Eth | Sp500 | Btc | Eth | Sp500 |
| Min. | −47.317 | −56.791 | −12.788 | 0.000 | 0.000 | 0.001 |
| 0.25q | −1.500 | −1.964 | −0.567 | 0.648 | 0.869 | 0.265 |
| Median | −0.144 | −0.063 | 0.066 | 1.457 | 2.010 | 0.651 |
| Mean | −0.128 | −0.036 | 0.023 | 2.277 | 2.992 | 0.907 |
| 0.75q | 1.345 | 2.081 | 0.715 | 3.112 | 4.024 | 1.184 |
| Max. | 16.504 | 24.254 | 8.947 | 47.317 | 56.791 | 12.788 |
| Std. Dev. | 3.462 | 4.485 | 1.387 | 2.611 | 3.340 | 1.049 |
| Skew. | −1.547 | −1.326 | −0.795 | 4.589 | 4.353 | 4.097 |
| Kurt. | 25.639 | 21.506 | 16.766 | 58.885 | 48.948 | 31.824 |

**Figure 2.** Returns—Bitcoin (btc), Ethereum (eth), and S&P 500 (sp500).

For each asset, we experimented with four specifications. The first one, denoted Ar1SV, is a traditional stochastic volatility model with AR(1) conditional variance, as in Equations (1) and (2). Our second specification is the LMSV model with volatility persistence given by a fractional Gaussian noise, as described in Equations (7) and (8). Since we employed penalized complexity priors for the fGn process, the LMSV specification is equivalent to INLA 2 from the Monte Carlo simulation tables of Section 4. As discussed in Section 3, to showcase the flexibility of the INLA method we augmented both Ar1SV and LMSV with a smooth spline component following a second-order random walk, which functions as a slow-moving unconditional average volatility. These specifications, Ar1SV

Spline and LMSV Spline, are presented in Equations (14)–(19). For further details on implementation and prior structure, we point to Appendix A.

5.1. Parameter Estimates

Tables 4 and 5 report descriptive statistics of the posterior distributions of the estimated parameters for Bitcoin and S&P500, respectively. Due to space constraints, we do not present the estimated parameters of Ethereum, but the results can be obtained from the authors.

Table 4. Posterior distribution estimated parameters—Bitcoin (btc).

| | Mean | Sd. | 0.025q | 0.5q | 0.975q |
|-----------------|-------------|------------|------------|-------------|-------------|
| Ar1SV | | | | | |
| μ | −7.348 | 0.083 | −7.510 | −7.348 | −7.185 |
| ϕ | 0.756 | 0.073 | 0.589 | 0.765 | 0.872 |
| τ_h | 0.846 | 0.103 | 0.661 | 0.840 | 1.068 |
| LMSV | | | | | |
| μ | −7.437 | 0.040 | −7.515 | −7.437 | −7.359 |
| H | 0.831 | 0.031 | 0.769 | 0.831 | 0.889 |
| τ_h | 0.760 | 0.088 | 0.601 | 0.755 | 0.947 |
| Ar1SV Spline | | | | | |
| μ | | | | | |
| ϕ | 0.284 | 0.081 | 0.120 | 0.285 | 0.438 |
| τ_h | 1.024 | 0.113 | 0.823 | 1.017 | 1.266 |
| τ_{spline} | 109,670.02 | 40,467.63 | 49,369.00 | 103,358.70 | 206,345.49 |
| LMSV Spline | | | | | |
| μ | | | | | |
| H | 0.691 | 0.050 | 0.600 | 0.690 | 0.792 |
| τ_h | 0.963 | 0.108 | 0.770 | 0.956 | 1.194 |
| τ_{spline} | 122,050.180 | 43,596.506 | 56,537.124 | 115,416.063 | 225,810.715 |

Table 5. Posterior distribution estimated parameters—S&P 500 (sp500).

| | Mean | Sd. | 0.025q | 0.5q | 0.975q |
|-----------------|------------|------------|------------|------------|-------------|
| AR1SV | | | | | |
| μ | −9.270 | 0.230 | −9.721 | −9.270 | −8.820 |
| ϕ | 0.972 | 0.009 | 0.952 | 0.973 | 0.987 |
| τ_h | 1.132 | 0.285 | 0.656 | 1.103 | 1.770 |
| LMSV | | | | | |
| μ | −9.295 | 0.045 | −9.384 | −9.295 | −9.207 |
| H | 0.978 | 0.014 | 0.942 | 0.982 | 0.996 |
| τ_h | 0.356 | 0.211 | 0.089 | 0.309 | 0.891 |
| AR1SV-Spline | | | | | |
| μ | | | | | |
| ϕ | 0.924 | 0.033 | 0.843 | 0.930 | 0.970 |
| τ_h | 2.524 | 0.830 | 1.285 | 2.395 | 4.515 |
| τ_{spline} | 82,149.686 | 39,133.007 | 29,518.874 | 74,417.060 | 179,789.999 |
| LMSV-Spline | | | | | |
| μ | | | | | |
| H | 0.949 | 0.037 | 0.852 | 0.958 | 0.992 |
| τ_h | 1.218 | 0.862 | 0.243 | 0.999 | 3.460 |
| τ_{spline} | 57,595.412 | 31,603.756 | 18,217.443 | 50,535.855 | 138,489.359 |

The qualitative implications of our estimations appear consistent over all three assets. The estimated values for persistence parameters H and ϕ in models Ar1SV and LMSV are large, but not so close to unity as to suggest the nonstationarity of latent volatility, which is typical in applications to financial data.

Incorporating a spline component into our stochastic volatility models has the effect of diminishing the persistence of latent log-variance through lowering the values of parameters ϕ (Ar1SV-Spline) and H (LMSV-Spline). Since we did not include a constant in the expressions for log-variance in our Ar1SV-Spline and LMSV-Spline specifications, the second-order random walk component serves as a smoothly varying unconditional mean log-variance. The effect of reduced persistence is more pronounced for Bitcoin than for the S&P 500, which could be taken as a suggestion that mean log-variance itself varies more over time for Bitcoin. This contributes to the body of empirical evidence showing that cryptocurrencies display higher overall levels of volatility, but also a relatively higher variability in the volatility level over time (e.g., Ghosh et al. 2023; Ahmed et al. 2024).

5.2. Model Selection

Table 6 reports the estimated log marginal likelihood (mlik) and widely applicable information criterion (waic) measures for all assets and models. Vehtari et al. (2016) introduced waic as an adequate method for choosing between Bayesian models of different complexities with an interpretation analogous to traditional information criteria (the lower the better).

The boldface entries in Table 6 indicate the best specification for each asset according to waic. For the two cryptocurrencies, the LMSV spline model was chosen as the best specification, while for S&P 500 the Ar1SV model was selected. These results seem to indicate that the LMSV specification is suitable for increased volatility and persistent cryptocurrencies, as the inclusion of a smoothly varying component in the unconditional average log-variance seems to better capture the volatility dynamics of this market. The inclusion of this component provides a way of incorporating a change in parameters into the unconditional volatility process, and the results are consistent with the many changes that occurred in these markets in the analyzed period.

Table 6. Model Fit Measures—marginal likelihood and waic.

| | Btc Mlik | Waic | Eth Mlik | Waic | Sp500 Mlik | Waic |
|--------------|-------------|---------|-------------|---------|---------------|---------|
| Ar1SV | 3500.0 | −7019.3 | 3095.9 | −6259.5 | 3597.6 | −7283.0 |
| LMSV | 3542.2 | −7041.8 | 3129.0 | −6297.5 | 3594.6 | −7262.4 |
| Ar1SV Spline | 3513.7 | −7088.7 | 3096.2 | −6288.2 | 3583.6 | −7282.3 |
| LMSV Spline | 3523.7 | −7122.7 | 3107.7 | −6315.9 | 3585.1 | −7271.9 |

Note: Bold values denote the selected model.

In-sample fit measures, ME, RMSE, and MAE are reported in Table 7. Absolute returns were taken as a proxy for the true unobserved volatility. We observe that the LMSV Spline model was the best model in terms of in-sample fit for S&P500, achieving the best results for all criteria, whereas for Bitcoin and Ethereum, the Ar1SV Spline model was chosen by ME, RMSE, and MAE. An out-of-sample forecast analysis is presented in the next subsection.

Table 7. In-sample error measures.

| | Btc ME | RMSE | MAE | Eth ME | RMSE | MAE | Sp500 ME | RMSE | MAE |
|--------------|-----------|---------|----------|-----------|----------|----------|-------------|---------|---------|
| Ar1SV | −0.00255 | 0.01927 | 0.01364 | −0.00519 | 0.027109 | 0.01916 | −0.00164 | 0.00771 | 0.00552 |
| LMSV | −0.00344 | 0.01847 | 0.01327 | −0.00513 | 0.024536 | 0.01745 | −0.00153 | 0.00719 | 0.00516 |
| Ar1SV Spline | −0.00209 | 0.01740 | 0.01208 | −0.00361 | 0.023785 | 0.01655 | −0.00165 | 0.00769 | 0.00551 |
| LMSV spline | −0.00322 | 0.01762 | 0.012549 | −0.00481 | 0.023853 | 0.016890 | −0.00172 | 0.00752 | 0.00540 |

Note: Bold values denote the best value.

Figures 3–6 present a comparison of the volatility implied by each model specification and the observed absolute returns for Bitcoin and S&P 500. We can see that the adjustment of estimated models is consistent with the pattern of volatility observed in all the analyzed series.




Figure 3. Bitcoin. (btc)—Fitted volatility and VaR (5%)—AR1SV and LMSV models.

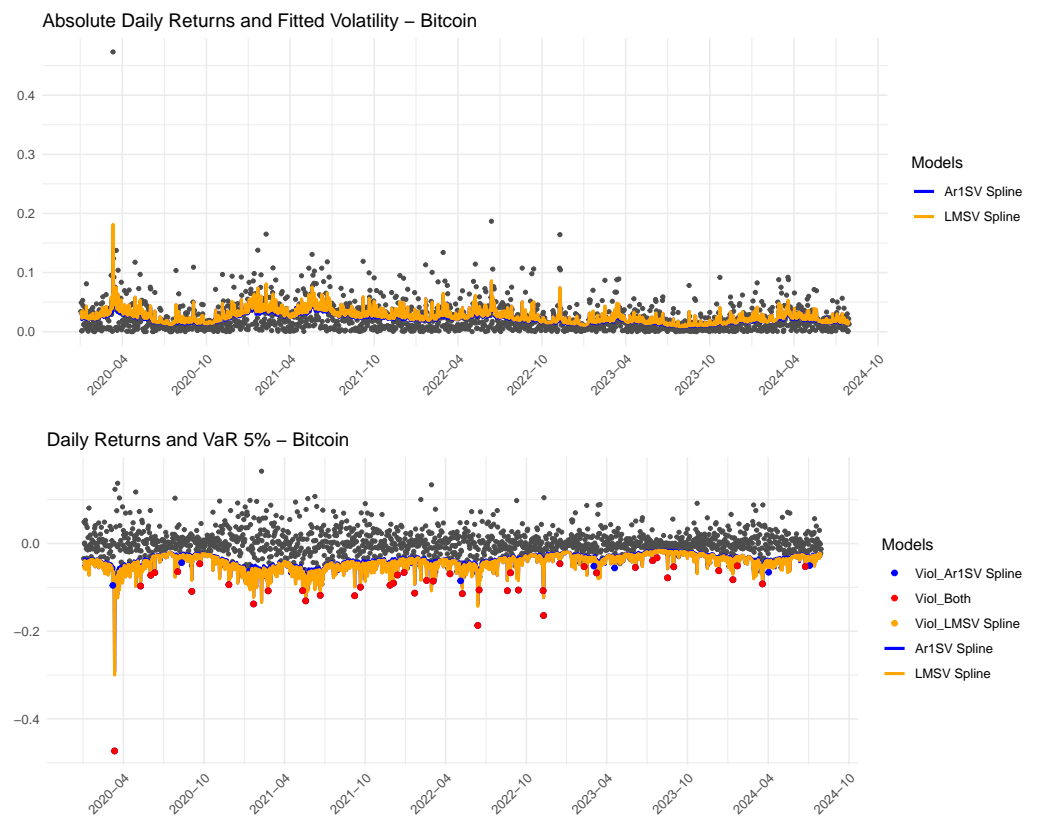


Figure 4. Bitcoin (btc)—Fitted volatility and VaR (5%)—Ar1SV spline and LMSV spline models.

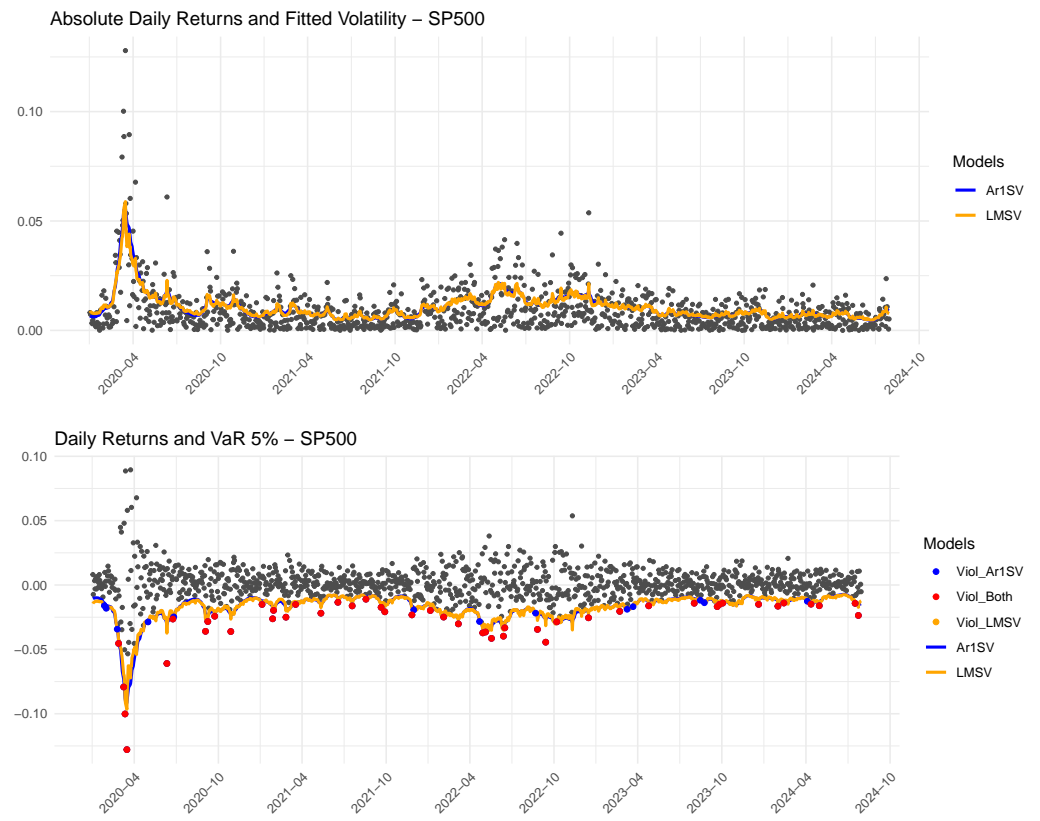


Figure 5. S&P 500 (sp500)—fitted volatility and VaR (5%)—AR1SV and LMSV models.

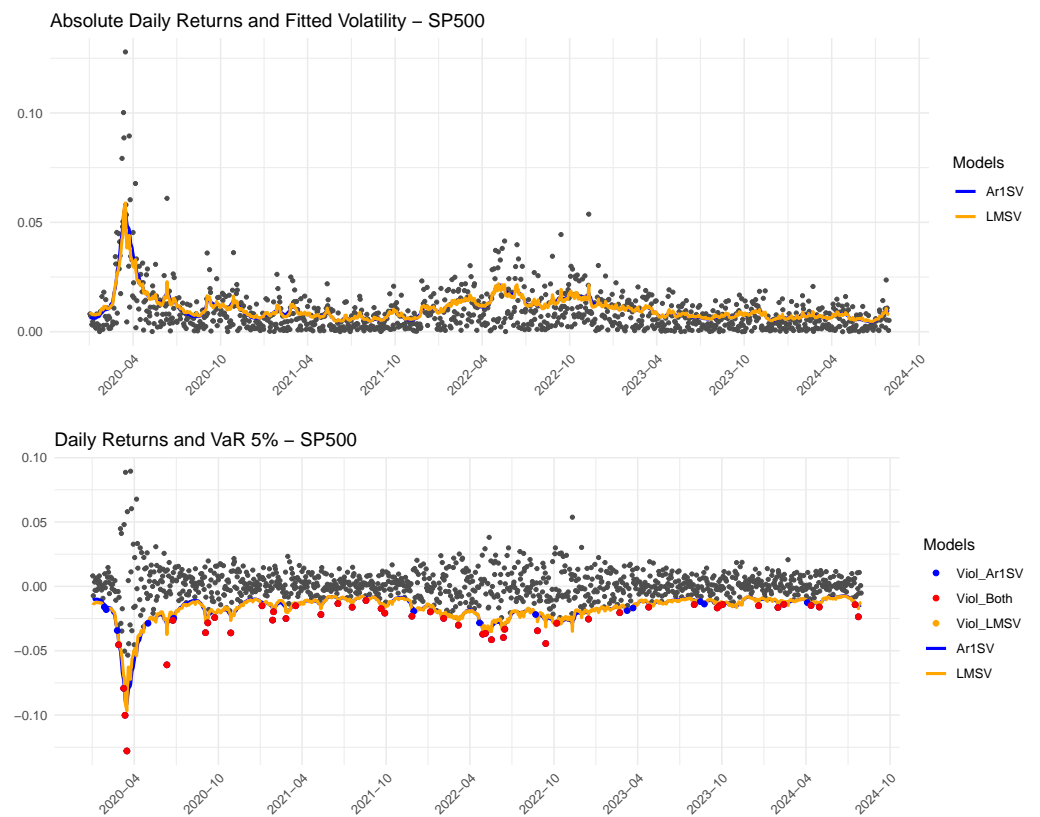


Figure 6. S&P 500 (sp500)—fitted volatility and VaR (5%)—Ar1SV spline and LMSV spline models.

5.3. Dynamic Value at Risk Estimation

Dynamic Value at Risk (VaR) measurements, calculated using different specifications, are presented in the bottom panels of Figures 3–6. The values were computed using the approximation $\text{VaR}(\alpha) = \hat{\mu} - q_\alpha \times \sigma_t$, where $\hat{\mu}$ is the sample mean of the process, q_α is the α -quantile of a standard Normal distribution and σ_t is the estimated conditional volatility. In this experiment we show the results for the calculation of a VaR with $\alpha = 0.05$, a usual measure of tail risk. The results for other VaR levels are available from the authors upon request. In the same Figures we also show the dynamic Value at Risk (VaR) measurements calculated using the different volatility components estimated using the four analyzed specifications.

Broadly, dynamic VaR measurements appear to follow the observed tail risk for all assets. In order to more carefully verify the performance of the different specifications for conditional volatility, we show, in Table 8, the proportion of observed violations within the sample (prop. viol.) and the p-values from the backtesting method for the VaR proposed in Christoffersen (1998) to verify if the empirical coverage of the estimated VaR is statistically equal to the VaR nominal value. This is compared against an alternative hypothesis of empirical coverage that is distinct from the nominal, indicating a problem in the VaR estimation.

Table 8. Value At Risk (VaR) statistics.

| | Btc | | Eth | | Sp500 | |
|--------------|-------------|-----------------------|-------------|---------|-------------|---------|
| | Prop. Viol. | p.Value | Prop. Viol. | p.Value | Prop. Viol. | p.Value |
| Ar1SV | 0.04425 | 0.31218 | 0.05322 | 0.53688 | 0.05212 | 0.73500 |
| LMSV | 0.03110 | 0.00020 | 0.03588 | 0.00697 | 0.04083 | 0.17557 |
| Ar1SV Spline | 0.03050 | 0.00012 | 0.03887 | 0.03763 | 0.05212 | 0.73500 |
| LMSV Spline | 0.02631 | 1.67×10^{-6} | 0.03528 | 0.00493 | 0.04517 | 0.49878 |

For the calculation of VaR for BTC, the Ar1SV and LMSV spline models present different performances. Ar1SV comes closest to the expected value of violations (5%), while the LMSV spline tends to be conservative, underestimating the violations. For ETH, Ar1SV is closest to the expected performance with respect to the VaR of 5%. The LMSV and LMSV spline models consistently underestimate the violations. Finally, for the S&P 500 series, both Ar1SV and Ar1SV spline are well-calibrated with the VaR of 5%, with violation proportions very close to 5%. The LMSV spline also presents a good adherence.

5.4. Out-of-Sample Forecast

Our previous analyses compared the in-sample performance of the conditional volatility models. To verify the predictive performance for conditional volatility out-of-sample, we performed a forecast experiment using a rolling sample for the last 22 observations in the sample, which were left out of the estimation procedure.

In this experiment, we employed rolling samples with a size equal to the size of the original series minus 22, and with this new series, we estimate the four specifications analyzed above. From these estimates, we predict future volatility for 1, 5, and 22 steps ahead, and compare this with the observed values of absolute returns (proxy for true unobserved volatility). Then, we added one more observation in the estimation sample, and replicated the forecasting process until the end of this augmented sample. The prediction results for each analyzed series, four different specifications, and the three forecast horizons are presented in Table 9. The best results in each category are highlighted in bold.

Table 9. Out-of-sample Forecast Measures.

| | ME | RMSE | MAE |
|----------------------|----------------------|----------------|----------------|
| Btc | | | |
| Ar1SV 1-step | 0.00402 | 0.01681 | 0.01328 |
| LMSV 1-step | 0.00023 | 0.01558 | 0.01200 |
| Ar1SV spline 1-step | −0.0012 | 0.01577 | 0.01214 |
| LMSV spline 1-step | 1×10^{-4} | 0.01595 | 0.01229 |
| Ar1SV 5-steps | 0.00474 | 0.01441 | 0.01197 |
| LMSV 5-steps | − 0.00056 | 0.01375 | 0.01070 |
| Ar1SV spline 5-steps | −0.00281 | 0.01438 | 0.01074 |
| LMSV spline 5-steps | −0.0011 | 0.0142 | 0.0111 |
| Ar1 22-steps | 0.00996 | 0.01669 | 0.01452 |
| LMSV 22-steps | 0.00431 | 0.01419 | 0.01196 |
| Ar1 spline 22-steps | − 0.00055 | 0.01429 | 0.01123 |
| LMSV spline 22-steps | 0.00101 | 0.01441 | 0.01164 |
| Eth | | | |
| Ar1SV 1-step | 0.00435 | 0.01683 | 0.01334 |
| LMSV 1-step | 0.00081 | 0.01576 | 0.01227 |
| Ar1SV spline 1-step | − 7×10^{-5} | 0.01571 | 0.01219 |
| LMSV spline 1-step | 0.00147 | 0.01573 | 0.01225 |
| Ar1SV 5-steps | 0.00841 | 0.01961 | 0.01644 |
| LMSV 5-steps | 0.00323 | 0.01700 | 0.01338 |
| Ar1SV spline 5-steps | 0.00143 | 0.01701 | 0.01320 |
| LMSV spline 5-steps | 0.0037 | 0.01766 | 0.01419 |
| Ar1 22-steps | 0.0134 | 0.02142 | 0.01953 |
| LMSV 22-steps | 0.00471 | 0.01719 | 0.01412 |
| Ar1 spline 22-steps | − 0.00017 | 0.01738 | 0.01283 |
| LMSV spline 22-steps | 0.00218 | 0.01773 | 0.01371 |
| SP500 | | | |
| Ar1SV 1-step | 0.0022 | 0.00396 | 0.00371 |
| LMSV 1-step | 0.00222 | 0.00392 | 0.00366 |
| Ar1SV spline 1-step | 0.00169 | 0.00380 | 0.00359 |
| LMSV spline 1-step | 0.00186 | 0.00381 | 0.00361 |
| Ar1SV 5-steps | 0.00276 | 0.00386 | 0.00356 |
| LMSV 5-steps | 0.00283 | 0.0039 | 0.0036 |
| Ar1SV spline 5-steps | 0.00187 | 0.00334 | 0.00298 |
| LMSV spline 5-steps | 0.0021 | 0.00342 | 0.00309 |
| Ar1 22-steps | 0.00089 | 0.00504 | 0.00383 |
| LMSV 22-steps | 5×10^{-4} | 0.00507 | 0.00376 |
| Ar1 spline 22-steps | −0.00172 | 0.00507 | 0.00346 |
| LMSV spline 22-steps | −0.00142 | 0.00507 | 0.00345 |

Note: Bold values denote the best value.

To compare the predictive results, we evaluated three error indicators for each time series (btc, eth, and sp500) and three forecast horizons (1-step, 5-step, and 22-step) for the four forecast models: Ar1SV, LMSV, Ar1SV Spline, and LMSV Spline. For the btc series and the one-step-ahead forecast, in terms of average error, the LMSV spline and Ar1SV spline stand out as they present the values closest to zero (0.0001 and −0.0012), respectively. In terms of RMSE, LMSV has the best performance (0.01558), which is repeated for the MAE in this precision horizon (0.0120), showing that the LMSV stands out in this configuration. For the five-step horizon, the LMSV model presents the best result in terms of ME (−0.00056), RMSE (0.01375), and MAE (0.01070). For the 22-step horizon, the Ar1 spline (−0.00055) and LMSV spline (0.00101) have the lowest mean errors, while the LMSV model has the lowest RMSE (0.01419), and the Ar1 spline model has the lowest MAE (0.01123).

Analyzing the predictive results for eth, for one-step-ahead, the Ar1SV spline model stands out in terms of having the best ME (-7×10^{-5}), RMSE (0.01571), and MAE (0.012190). For the five-step horizon, the Ar1SV spline is again better in terms of ME (0.00143) and MAE (0.01320), and the LMSV presents the lowest RMSE (0.01700). The results are similar for the 22-step horizon, with a better ME (-0.00017) and MAE (0.01283) for the Ar1 spline model and a better RMSE (0.017190) being obtained by the LMSV. The results for sp500 indicate that the AR1SV spline obtained the best performance for one-step-ahead, with an ME, RMSE, and MAE of 0.001690, 0.00380, and 0.00359, respectively. Again, the best performance for five-steps-ahead was achieved with an ME, RMSE and MAE of 0.00187, 0.00334, and 0.002980, and for 22-steps, the LMSV had the best ME (5×10^{-4}), and Ar1 presented with the lowest RMSE (0.00504) and Ar1 spline with the lowest MAE (0.00346).

In summary, for btc, the LMSV model is consistently the best in terms of RMSE and MAE, except for 22-steps, where the Ar1 spline excels in terms of MAE. For eth, the Ar1SV spline is the most effective for 1-step, while for 5- and 22 steps, the LMSV and Ar1 spline models present the best results in terms of RMSE and MAE, respectively, and for sp500, the Ar1SV spline has a better performance for 1- and 5-steps, while for 22-steps, the Ar1 and Ar1 spline stand out.

6. Application to High Frequency Data

A natural application of long-memory models is to model conditional volatility in high-frequency return data (Deo et al. 2006). This type of application is important since it is possible to use the estimated conditional variance for high-frequency data to calculate the realized volatility measurements (e.g., Maasoumi and McAleer 2008; McAleer and Medeiros 2008; Lieberman and Phillips 2008) for daily data. The estimation methodology based on INLA and the approximation of an fGn by a Gaussian Markov Random Field is especially interesting. It is computationally efficient in terms of computational speed and memory usage, and thus can be used for the large number of observations observed in high-frequency data. The second advantage is the possibility of directly incorporating a latent variable into the model to capture patterns of intraday seasonality in market volatility without the need for multi-step estimation procedures, as proposed in Deo et al. (2006). The model uses the sum of fGN and zero-sum intraday periodic components, as provided by Equation (20).

We estimated an LMSV model using 5 min intraday Bitcoin data from Bitstamp exchange for the period from 1 January 2021 until 20 September 2024, corresponding to a sample with 390,544 intraday observations. Figure 7 shows the intraday Bitcoin returns that were analyzed. As described in Section 3, additive intraday seasonality was included. For each 5 min window, one component was included, totaling 288 different intraday seasonal fixed effects.

Table 10 shows the parameters estimated using the LMSV model for this series, and Figure 8 shows the estimated conditional volatility compared to the absolute intraday returns. We can observe that the model closely fits the temporal variation observed in the intraday returns.

Table 10. Posterior Estimated Paramaters—LMSV for Bitcoin 5 min returns.

| | Mean | Sd. | 0.025q | 0.5q | 0.975q |
|---------------|-----------------------|-----------------------|-----------------------|-----------------------|----------|
| μ | −6.01 | 0.0892 | −6.190 | −6.010 | −5.840 |
| H | 0.902 | 0.00199 | 0.898 | 0.902 | 0.906 |
| τ_h | 0.581 | 0.00776 | 0.566 | 0.581 | 0.596 |
| τ_{seas} | 8.81×10^{-5} | 1.64×10^{-5} | 6.12×10^{-5} | 8.64×10^{-5} | 0.000125 |

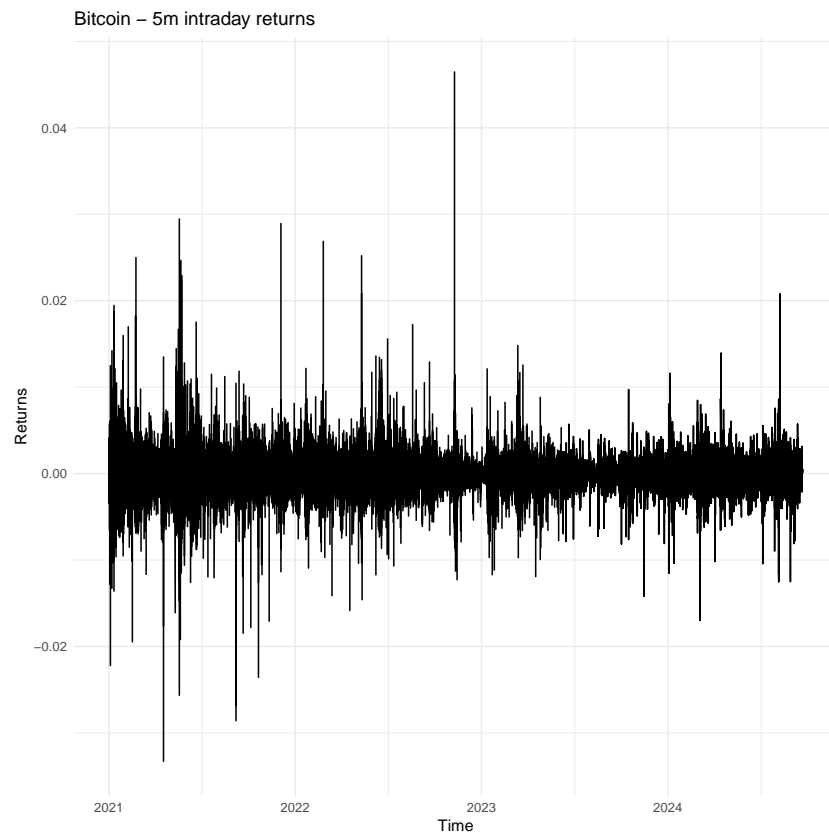


Figure 7. 10-min intraday Bitcoin returns.

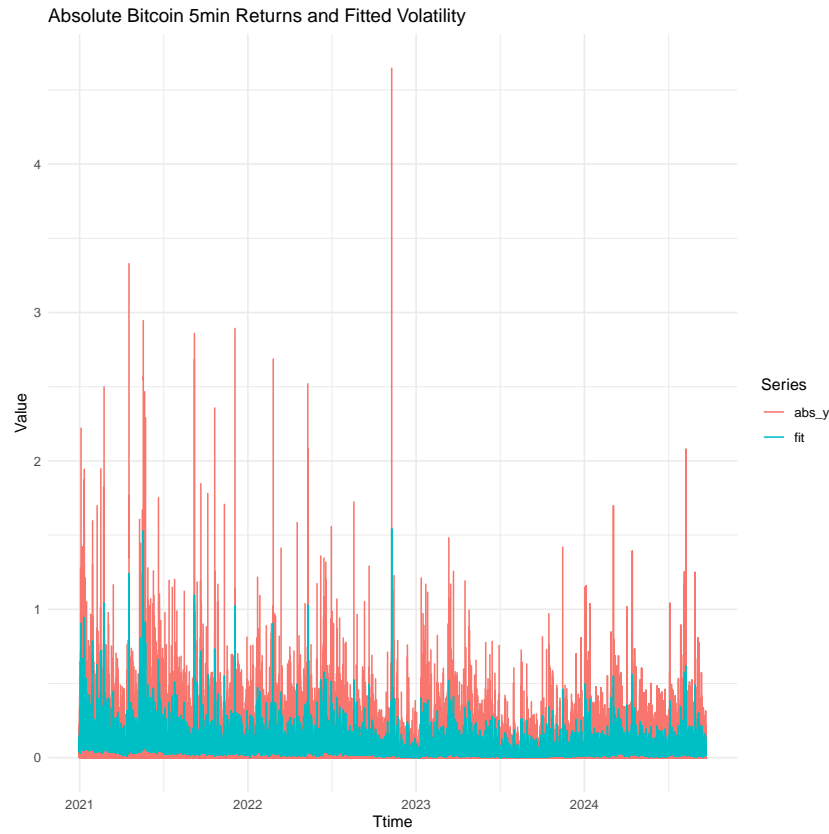


Figure 8. Fitted Intraday Volatility.

The patterns of intraday seasonality estimated by the model can be seen in Figure 9, which shows that there is substantial variation during the day. Log-variance is higher during the night time (with a notable increase starting from 18:00 UTC) and lower during the day, from 6:00 until 17:00. Shaded bands represent 95% credible intervals.




Figure 9. Fitted Intraday Volatility Seasonality.

The mean error in the estimation of this model, using absolute returns as the metric of true log-variance, was calculated as 0.00229, indicating a minimum bias in this estimation. The root mean squared error was calculated with the value of 0.04588, and the mean absolute error was calculated with the value of 0.02710, indicating an adequate adjustment to the observed intraday data.

As previously discussed, one of the main advantages of using INLA in the Bayesian estimation procedure is its computational efficiency, a problem that is a fundamental limitation in the estimation of complex dynamic models with latent variables using intraday data. The total computational time involved in estimating this model was 3905 s, which is quite reduced compared to the complexity of a model with long memory and this high number of observations, emphasizing the computational gains of this approach compared to alternative methods, such as MCMC.

7. Conclusions

In this work, we explored the use of an estimation methodology for long-memory stochastic volatility models using an approximation of a fractional Gaussian Noise process as a Gaussian Markov Random Field, and with this approach we used the Integrated Nested Laplace Approximations methodology to perform a Bayesian estimation of parameters and latent variables.

This methodology is a useful addition to the set of tools used in the modeling of conditional volatility in time series, since it is computationally efficient when compared to traditional posterior simulation methods and the model specification can easily be extended to include additional latent factors. In this work, we show that a simple extension, with the addition of a smooth variation component to model the variation in the long-run average of the conditional variance, analogous to a spline, allows for gains in terms of in-sample fit

and out-of-sample forecasting. Another extension is that of an LMSV model that includes seasonal patterns, which is especially important in the modeling of intraday returns.

The computational efficiency of this method is especially important in the modeling of high-frequency data, which are characterized by a large number of observations. Sample size may be an important limitation for other methods of conditional variance estimation, such as quasi-maximum likelihood estimators and MCMC-based methods. We show how the proposed estimator can be used to construct measures of realized variance, and also to make out-of-sample forecasts for this measure, incorporating the entire dependence structure in the conditional variance observed in the intraday returns.

The Integrated Nested Laplace Approximations (INLA) methodology, despite its notable computational efficiency and effectiveness, presents several limitations that warrant consideration. Primarily, INLA excels in the estimation of latent Gaussian models; however, its applicability may be constrained when addressing models that do not conform to this framework. More intricate non-Gaussian or non-linear models can pose challenges for estimation through INLA. While it is feasible to incorporate alternative distributions as innovation processes within the observation equation, the reliance on Gaussian distributions for latent variables constitutes a fundamental assumption of the INLA approach. Consequently, stochastic volatility models that incorporate non-Gaussian distributions often necessitate modifications, such as the implementation of variational approximations (Cabral et al. 2024; Van Niekerk et al. 2023), which can complicate the inference process and increase computational demands.

Additionally, INLA employs numerical approximations to compute posterior marginals, which may not consistently achieve the precision of fully simulated methods like the Markov Chain Monte Carlo (MCMC), particularly in scenarios characterized by highly skewed or multi-modal posterior distributions. Although INLA is well-suited to hierarchical and structured models, its flexibility may be more limited compared to MCMC, which can accommodate models with complex dependencies and structures that diverge from conventional latent Gaussian settings. Furthermore, while INLA is adept at managing large datasets, the computational time and memory requirements may significantly escalate when applied to extremely high-frequency data or models with intricate dependencies, while remaining more efficient than traditional MCMC-based methodologies.

This methodology can be extended in several directions. It is possible to construct univariate multifactor models, for example, with permanent and transient components, or factors with short and long memory, in the same model. It is possible to formulate a model with both Ar1 and fGn components, and to verify whether there are different memory patterns in the conditional variance. A model with Ar1 and fGn components would be similar to an ARFIMA (1,d,0) specification for the conditional variance of the process, which may be a useful specification for some time series.

The method presented in this article also can be used to generalize the structure of multivariate stochastic volatility models estimated using INLA, as proposed in Nacinben and Laurini (2024a). In this respect, the long-memory dynamics could be used as an alternative to the first-order autoregressive structure in the construction of multivariate models using a multifactor structure. Another possible application is to compare the fit and predictions of the long memory model on conditional volatility with alternative models using regime shifts and other forms of parameter variation. Since the approximation structure for the long memory process used in our study is itself based on the aggregation of short-memory models, it is expected that this method could be an alternative way to predict volatility in the presence of changes in the process memory.

Another potential extension is the incorporation of non-Gaussian distributions into the innovation process of latent factor dynamics. This approach could enhance the model's robustness in the presence of heavy-tailed distributions that affect the conditional volatility process. Furthermore, variational methods can be integrated with INLA to facilitate inference in stochastic volatility models that employ non-Gaussian distributions, as demonstrated in Cabral et al. (2024) and applied in the estimation of univariate SV models by

Nacinben and Laurini (2024b). By combining long memory structures with non-Gaussian distributions, we can broaden the applicability of stochastic volatility models to encompass the two significant empirical features commonly observed in financial data.

Author Contributions: Conceptualization, P.C. and M.P.L.; methodology, P.C. and M.P.L.; formal analysis, P.C. and M.P.L.; writing—review and editing, P.C. and M.P.L.; funding acquisition, M.P.L. All authors have read and agreed to the published version of the manuscript.

Funding: The authors acknowledge funding from Capes, CNPq (310646/2021-9) and FAPESP (2023/02538-0).

Data Availability Statement: Data from <https://coinmarketcap.com/>, accessed on 10 September 2024.

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

Abbreviations

AR(1)—First-order Autoregressive Process. ARFIMA—Autoregressive Fractionally Integrated Moving Average. ARMA—Autoregressive Moving Average. Ar1SV—Stochastic Volatility Model with AR(1) Dynamics. Ar1SV-Spline—Stochastic Volatility Model with AR(1) and RW2 Dynamics. Btc—Bitcoin. Eth—Ethereum. fBm—Fractional Brownian Motion. fGn—Fractional Gaussian Noise. GARCH—Generalized Autoregressive Conditional Heteroskedasticity. GMRF—Gaussian Markov Random Field. H —Hurst Exponent. HP—Hodrick-Prescott Filter. INLA—Integrated Nested Laplace Approximations. LMSV—Long Memory Stochastic Volatility. LMSV-Spline—Stochastic Volatility Model with fGn and RW2 Dynamics. MAE—Mean Absolute Error. MCMC—Markov Chain Monte Carlo. ME—Mean Error. QML—Quasi-Maximum Likelihood. RMSE—Root Mean Squared Error. SV—Stochastic Volatility. VaR—Value at Risk.

Appendix A. Specification Details

In all models, the mean parameters (that is, μ in SV and LMSV models) have a Normal prior with a large known variance.

Appendix A.1. SV Model

$$r_t = \exp(h_t/2)\varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (\text{A1})$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + z_t, \quad z_t \sim N(0, 1/(1 - \phi^2)\tau). \quad (\text{A2})$$

This model has three parameters: μ , ϕ , and τ .

The AR(1) structure has two hyperparameters θ_1 and θ_2 :

$$\theta_1 = \log(\kappa), \quad \kappa = \tau(1 - \phi^2), \quad (\text{A3})$$

$$\theta_2 = \log\left(\frac{1 + \phi}{1 - \phi}\right). \quad (\text{A4})$$

with the following prior distributions:

$$\theta_1 \sim \text{LogGamma}(1, 5 \times 10^{-5}), \quad (\text{A5})$$

$$\theta_2 \sim \text{Normal}(0, (0.15)^{-1}). \quad (\text{A6})$$

Appendix A.2. LMSV Model

$$r_t = \exp(h_t/2)\varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (\text{A7})$$

$$h_t = \mu + x_t, \quad x_t \sim fGn(H, \tau_h), \quad (\text{A8})$$

There are three estimated parameters: μ , H , and τ_h .

The fGn process h_t has two hyperparameters θ_1 and θ_2 .

$$\tau = \exp(\theta_1), \quad \theta_1 \sim pc.prec(0.2, 0.008) \quad (A9)$$

$$H = \frac{1}{2} \left(\frac{\exp(\theta_2)}{1 + \exp(\theta_2)} \right), \quad \theta_2 \sim pc.cor1(0.9, 0.9) \quad (A10)$$

Appendix A.3. SV Spline Model

In the SV spline model, the log volatility is as follows:

$$r_t = \exp(h_t + s_t)/2 \varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (A11)$$

$$h_t = \phi h_{t-1} + \eta_t, \quad h_t \sim N(0, 1/\tau_h), \quad (A12)$$

$$\Delta^2 s_t = s_t - 2s_{t+1} + s_{t+2}, \quad s_t \sim N(0, 1/\tau_s) \quad (A13)$$

The second-order random walk spline component has one hyperparameter θ_1 , which is the marginal precision of the following component:

$$\theta_1 = \log(\tau_s), \quad (A14)$$

$$\theta_1 \sim \log\text{Gamma}(0, 0.15). \quad (A15)$$

This model has three parameters: τ_s , ϕ , and τ_h .

The AR(1) term has a prior structure similar to the standard SV model, with two hyperparameters: θ_3 and θ_2 .

$$\theta_2 = \log(\kappa), \quad \kappa = (1 - \phi), \quad (A16)$$

$$\theta_3 = \log\left(\frac{1 + \phi}{1 - \phi}\right) \quad (A17)$$

$$\theta_2 \sim pc.prec(0.2, 0.008), \quad (A18)$$

$$\theta_3 \sim pc.cor1(0.9, 0.9). \quad (A19)$$

Appendix A.4. Spline LMSV Model

In the Spline LMSV model log volatility follows

$$r_t = \exp((h_t + s_t)/2) \varepsilon_t, \quad \varepsilon_t \sim N(0, 1), \quad (A20)$$

$$h_t = \mu + fGn(H, \tau), \quad (A21)$$

$$\Delta^2 s_t = s_t - 2s_{t+1} + s_{t+2}, \quad s_t \sim N(0, 1/\tau_s). \quad (A22)$$

The second-order random walk spline component has one hyperparameter, θ_1 , which is the precision of the i.i.d Normal innovations, v_t .

$$\theta_1 = \log(\tau_s), \quad (A23)$$

$$\theta_1 \sim \log\text{Gamma}(0, 0.15). \quad (A24)$$

The fGn process y_t has two hyperparameters: θ_2 and θ_3 .

$$\tau = \exp(\theta_2), \quad \theta_2 \sim pc.prec(0.2, 8 \times 10^{-5}) \quad (A25)$$

$$H = \frac{1}{2} \left(\frac{\exp(\theta_3)}{1 + \exp(\theta_3)} \right), \quad \theta_3 \sim pc.cor1(0.9, 0.9) \quad (A26)$$

Appendix A.5. Seasonal LMSV Model

$$h_t = seas_t + y_t + \eta_t, \quad \eta_t \sim N(0, \tau_h) \quad (A27)$$

$$seas_t + seas_{t-1} + \dots + seas_{t-m+1} \sim N(0, \tau_{seas}) \quad (A28)$$

This model has three estimated parameters: τ_{seas} , τ_h , and H . Periodicity m is 288, which reflects daily frequency with returns of 5 min apart.

The seasonal component $seas_t$ has one hyperparameter: θ_1

$$\theta_1 = \log(\tau_{seas}) \quad (\text{A29})$$

$$\theta_1 \sim \text{logGamma}(0, 0.15). \quad (\text{A30})$$

The fGn component has two hyperparameters: θ_2 and θ_3

$$\theta_1 = \log(\kappa), \quad k = \tau(1 - \rho^2), \quad (\text{A31})$$

$$\theta_2 = \log\left(\frac{1 + \rho}{1 - \rho}\right) \quad (\text{A32})$$

We employ the penalized complexity priors introduced by Sørbye and Rue (2018) for both precision τ and persistence ρ .

Note

- ¹ See <https://r-nimble.org/>, accessed on 20 December 2023 and <https://www.r-inla.org/> package. The simulations here and the estimation in Section 6 were performed in an Intel(R) Xeon(R) W-2265 CPU @ 3.50 GHz with 128gb Ram memory.

References

- Ahmed, Mohamed Shaker, Ahmed A. El-Masry, Aktham I. Al-Maghyreh, and Satish Kumar. 2024. Cryptocurrency volatility: A review, synthesis, and research agenda. *Research in International Business and Finance* 71: 102472. [CrossRef]
- Andersen, Torben G., and Bent E. Sørensen. 1996. Gmm estimation of a stochastic volatility model: A monte carlo study. *Journal of Business & Economic Statistics* 14: 328–52.
- Andersen, Torben G., Hyung-Jin Chung, and Bent E. Sørensen. 1999. Efficient method of moments estimation of a stochastic volatility model: A Monte Carlo study. *Journal of Econometrics* 91: 61–87. [CrossRef]
- Arouxet, M. Belén, Aurelio F. Bariviera, Verónica E. Pastor, and Victoria Vampa. 2022. COVID-19 impact on cryptocurrencies: Evidence from a wavelet-based hurst exponent. *Physica A: Statistical Mechanics and Its Applications* 596: 127170. [CrossRef]
- Arteche, Josu. 2004. Gaussian semiparametric estimation in long memory in stochastic volatility and signal plus noise models. *Journal of Econometrics* 119: 119–31. [CrossRef]
- Asai, Manabu. 2005. Comparison of MCMC methods for estimating stochastic volatility models. *Computational Economics* 25: 281–301. [CrossRef]
- Assaf, Ata, Avishek Bhandari, Husni Charif, and Ender Demir. 2022. Multivariate long memory structure in the cryptocurrency market: The impact of COVID-19. *International Review of Financial Analysis* 82: 102132. [CrossRef]
- Bariviera, Aurelio F., María José Basgall, Waldo Hasperué, and Marcelo Naiouf. 2017. Some stylized facts of the bitcoin market. *Physica A: Statistical Mechanics and Its Applications* 484: 82–90. [CrossRef]
- Basak, Gopal K., Ngai Hang Chan, and Wilfredo Palma. 2001. On the approximation of long memory processes by an ARMA model. *Journal of Forecasting* 20: 367–89. [CrossRef]
- Beran, Jan. 2017. *Statistics for Long-Memory Processes*. London: Routledge.
- Beran, Jan, Martin Schütznner, and Sucharita Ghosh. 2010. From short to long memory: Aggregation and estimation. *Computational Statistics & Data Analysis* 54: 2432–42.
- Bollerslev, Tim. 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31: 307–27. [CrossRef]
- Bouri, Elie, Luis A. Gil-Alana, Rangan Gupta, and David Roubaud. 2019. Modelling long memory volatility in the Bitcoin market: Evidence of persistence and structural breaks. *International Journal of Finance & Economics* 24: 412–26.
- Breidt, F. Jay, Nuno Crato, and Pedro De Lima. 1998. The detection and estimation of long memory in stochastic volatility. *Journal of Econometrics* 83: 325–48. [CrossRef]
- Cabral, Rafael, David Bolin, and Håvard Rue. 2024. Fitting latent non-gaussian models using variational bayes and laplace approximations. *Journal of the American Statistical Association* 1–13. [CrossRef]
- Chaim, Pedro, and Márcio P. Laurini. 2019. Nonlinear dependence in cryptocurrency markets. *The North American Journal of Economics and Finance* 48: 32–47. [CrossRef]
- Chan, Ngai Hang, and Giovanni Petris. 2000. Long memory stochastic volatility: A Bayesian approach. *Communication in Statistics—Theory and Methods* 29: 1367–78. [CrossRef]
- Chan, Ngai Hang, and Wilfredo Palma. 1998. State space modelling of long memory processes. *Annals of Statistics* 26: 719–40. [CrossRef]
- Charfeddine, Lanouar, and Youcef Maouchi. 2019. Are shocks on the returns and volatility of cryptocurrencies really persistent? *Finance Research Letters* 28: 423–30. [CrossRef]

- Cheah, Eng-Tuck, Tapas Mishra, Mamata Parhi, and Zhuang Zhang. 2018. Long memory interdependency and inefficiency in bitcoin markets. *Economics Letters* 167: 18–25. [\[CrossRef\]](#)
- Chib, Siddhartha, Federico Nardari, and Neil Shephard. 2002. Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics* 108: 281–316. [\[CrossRef\]](#)
- Christensen, Bent Jesper, and Morten Orregaard Nielsen. 2007. The effect of long memory in volatility on stock market fluctuations. *The Review of Economics and Statistics* 89: 684–700. [\[CrossRef\]](#)
- Christoffersen, Peter F. 1998. Evaluating interval forecasts. *International Economic Review* 39: 841–42. [\[CrossRef\]](#)
- Crato, Nuno, and Bonnie K. Ray. 2002. Semi-parametric smoothing estimators for long memory processes with added noise. *Journal of Statistical Planning and Inference* 105: 283–97. [\[CrossRef\]](#)
- Creal, Drew, Siem Jan Koopman, and André Lucas. 2013. Generalized autoregressive score models with applications. *Journal of Applied Econometrics* 28: 777–95. [\[CrossRef\]](#)
- Deo, Rohit, Clifford Hurvich, and Yi Lu. 2006. Forecasting realized volatility using a long memory stochastic volatility model: Estimation, prediction and seasonal adjustment. *Journal of Econometrics* 131: 29–58. [\[CrossRef\]](#)
- Ehlers, Ricardo, and Mauricio Zavallos. 2015. Bayesian estimation and prediction of stochastic volatility models via INLA. *Communications in Statistics—Simulation and Computation* 44: 683–93. [\[CrossRef\]](#)
- Engle, Robert F. 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50: 987–1007. [\[CrossRef\]](#)
- Engle, Robert F., and Jose Gonzalo Rangel. 2008. The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *Review of Financial Studies* 21: 1187–222. [\[CrossRef\]](#)
- Ferraz, Rosemeire O., and Luiz K. Hotta. 2007. Quasi-maximum likelihood estimation of long-memory stochastic volatility models. *Brazilian Review of Econometrics* 27: 225–33. [\[CrossRef\]](#)
- Fleming, Jeff, and Chris Kirby. 2011. Long memory in volatility and trading volume. *Journal of Banking & Finance* 35: 1714–26.
- Gamerman, Dani, and Hedibert F. Lopes. 2006. *Markov Chain Monte Carlo: Stochastic Simulation for Bayesian Inference*. New York: Chapman & Hall.
- Ghosh, Bikramaditya, Elie Bouri, Jung Bum Wee, and Noshaba Zulfiqar. 2023. Return and volatility properties: Stylized facts from the universe of cryptocurrencies and nfts. *Research in International Business and Finance* 65: 101945. [\[CrossRef\]](#)
- Gong, Chen, and David S. Stoffer. 2020. A note on efficient fitting of stochastic volatility models. *Journal of Time Series Analysis* 42: 186–200. [\[CrossRef\]](#)
- Granger, Clive W. J. 1980. Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* 14: 227–38. [\[CrossRef\]](#)
- Haldrup, Niels, and J. Eduardo Vera Valdés. 2017. Long memory, fractional integration, and cross-sectional aggregation. *Journal of Econometrics* 199: 1–11. [\[CrossRef\]](#)
- Harvey, Andrew C. 1989. *Forecasting, Structural Time Series and the Kalman Filter*. Cambridge: Cambridge University Press.
- Harvey, Andrew C. 1998. Stochastic volatility models with long memory. In *Forecasting Volatility in Financial Markets*. Edited by J. Knight and E. Satchell. London: Butterworth-Haineman, pp. 307–20.
- Harvey, Andrew C., and A. Jaeger. 1993. Detrending, stylized facts and the business cycle. *Journal of Applied Econometrics* 8: 231–47. [\[CrossRef\]](#)
- Harvey, Andrew, and Thomas Trimbur. 2008. Trend estimation and the Hodrick-Prescott filter. *Journal of the Japan Statistical Society* 38: 41–49. [\[CrossRef\]](#)
- Harvey, Andrew C., Esther Ruiz, and Neil G. Shephard. 1992. *Multivariate Stochastic Variance Model*. Technical Report, Financial Markets Group Discussion Paper. London: London School of Economics.
- Hillebrand, Eric, and Marcelo C. Medeiros. 2016. Nonlinearity, breaks, and long-range dependence in time-series models. *Journal of Business & Economic Statistics* 34: 23–41.
- Hosking, J. R. M. 1981. Fractional differencing. *Biometrika* 68: 165–76. [\[CrossRef\]](#)
- Hurvich, Clifford M., and Philippe Soulier. 2009. Stochastic volatility models with long memory. In *Handbook of Financial Time Series*. Berlin and Heidelberg: Springer, pp. 345–54.
- Johanes, Michael, and Nicholas Polson. 2005. Chapter MCMC for Financial Econometrics. In *Handbook of Financial Econometrics*. Amsterdam: Elsevier-North-Holland, pp. 1–72.
- Kastner, Gregor. 2019. Sparse Bayesian time-varying covariance estimation in many dimensions. *Journal of Econometrics* 210: 98–115. [\[CrossRef\]](#)
- Kastner, Gregor, and Sylvia Frühwirth-Schnatter. 2014. Ancillarity-sufficiency interweaving strategy (asis) for boosting MCMC estimation of stochastic volatility models. *Comput. Stat. Data Anal.* 76: 408–23. [\[CrossRef\]](#)
- Kim, Sangjoon, Neil Shephard, and Siddhartha Chib. 1998. Stochastic volatility: Likelihood inference and comparison with ARCH models. *The Review of Economic Studies* 65: 361–93. [\[CrossRef\]](#)
- Laurini, Márcio Poletti, and Luiz Koodi Hotta. 2017. GMC/GEL estimation of stochastic volatility models. *Communications in Statistics—Simulation and Computation* 46: 6828–44. [\[CrossRef\]](#)
- Lieberman, Offer, and Peter C. B. Phillips. 2008. Refined inference on long memory in realized volatility. *Econometric Reviews* 27: 254–67. [\[CrossRef\]](#)

- Lindgren, Finn, and Håvard Rue. 2008. On the second-order random walk model for irregular locations. *Scandinavian Journal of Statistics* 35: 691–700. [\[CrossRef\]](#)
- Maasoumi, Esfandiari, and Michael McAleer. 2008. Realized volatility and long memory: An overview. *Econometric Reviews* 27: 1–9. [\[CrossRef\]](#)
- Martino, Sara, Kjersti Aas, Ola Lindqvist, Linda R. Neef, and Håvard Rue. 2011. Estimating stochastic volatility models using integrated nested Laplace approximations. *The European Journal of Finance* 17: 487–503. [\[CrossRef\]](#)
- McAleer, Michael, and Marcelo C. Medeiros. 2008. Realized volatility: A review. *Econometric Reviews* 27: 10–45. [\[CrossRef\]](#)
- Mensi, Walid, Khamis Hamed Al-Yahyaee, and Sang Hoon Kang. 2019. Structural breaks and double long memory of cryptocurrency prices: A comparative analysis from bitcoin and ethereum. *Finance Research Letters* 29: 222–30. [\[CrossRef\]](#)
- Nacinben, João Pedro Coli de Souza Monteneri, and Márcio Laurini. 2024a. Multivariate stochastic volatility modeling via integrated nested Laplace approximations: A multifactor extension. *Econometrics* 12: 5. [\[CrossRef\]](#)
- Nacinben, João Pedro Coli de Souza Monteneri, and Márcio Laurini. 2024b. Non-gaussian stochastic volatility models: Laplace-variational Bayes inference. Presented at the XXIV Brazilian Finance Meeting, Curitiba, Brazil, July 4–6.
- Nelson, Daniel B. 1988. The Time Series Behavior of Stock Market Volatility and Returns. Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, USA.
- Paige, Robert L., and A. Alexandre Trindade. 2010. The Hodrick-Prescott Filter: A special case of penalized spline smoothing. *Electronic Journal of Statistics* 4: 856–74. [\[CrossRef\]](#)
- Perez, Ana, and Esther Ruiz. 2001. Finite sample properties of a QML estimator of stochastic volatility models with long memory. *Economics Letters* 70: 157–64. [\[CrossRef\]](#)
- Phillip, Andrew, Jennifer Chan, and Shelton Peiris. 2019. On long memory effects in the volatility measure of cryptocurrencies. *Finance Research Letters* 28: 95–100. [\[CrossRef\]](#)
- Rue, Håvard, Sara Martino, and Nicolas Chopin. 2009. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71: 319–92. [\[CrossRef\]](#)
- Sandmann, Gleb, and Siem Jan Koopman. 1998. Estimation of stochastic volatility models via Monte Carlo maximum likelihood. *Journal of Econometrics* 87: 271–301. [\[CrossRef\]](#)
- Scharth, Marcel, and Marcelo C. Medeiros. 2009. Asymmetric effects and long memory in the volatility of Dow Jones stocks. *International Journal of Forecasting* 25: 304–27. [\[CrossRef\]](#)
- Shackleton, Mark B., Shiuyan Pong, and Stephen J. Taylor. 2008. Distinguishing short and long memory volatility specifications. *The Econometrics Journal* 11: 617–37.
- Simpson, Daniel, Håvard Rue, Andrea Riebler, Thiago G. Martins, and Sigrunn H. Sørbye. 2017. Penalising model component complexity: A principled, practical approach to constructing priors. *Statistical Science* 32: 1–28. [\[CrossRef\]](#)
- Sørbye, Sigrunn Holbek, and Håvard Rue. 2018. Fractional Gaussian noise: Prior specification and model comparison. *Environmetrics* 29: e2457. [\[CrossRef\]](#)
- Sørbye, Sigrunn H, Eirik Myrvoll-Nilsen, and Håvard Rue. 2019. An approximate fractional Gaussian noise model with $O(n)$ computational cost. *Statistics and Computing* 29: 821–33. [\[CrossRef\]](#)
- Tan, Zhengxun, Yilong Huang, and Binuo Xiao. 2021. Value at risk and returns of cryptocurrencies before and after the crash: Long-run relations and fractional cointegration. *Research in International Business and Finance* 56: 101347. [\[CrossRef\]](#)
- Taqqu, Murad S. 2003. Fractional Brownian motion and long-range dependence. In *Theory and Applications of Long-Range Dependence*. Boston: Birkhäuser, pp. 5–38.
- Taylor, Stephen J. 1986. *Modelling Financial Time Series*. Chichester: John Wiley & Sons.
- Urquhart, Andrew. 2016. The inefficiency of bitcoin. *Economics Letters* 148: 80–82. [\[CrossRef\]](#)
- Valente, Fernanda, and Márcio Laurini. 2024. Bayesian inference for long memory term structure models. *Journal of Statistical Computation and Simulation* 94: 1735–59. [\[CrossRef\]](#)
- Van Niekerk, Janet, Elias Krainski, Denis Rustand, and Håvard Rue. 2023. A new avenue for Bayesian inference with INLA. *Computational Statistics & Data Analysis* 181: 107692. [\[CrossRef\]](#)
- Vehtari, Aki, Andrew Gelman, and Jonah Gabry. 2016. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing* 27: 1–20.

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.