

Modeling Univariate and Multivariate Stochastic Volatility in R with `stochvol` and `factorstochvol`

Darjus Hosszejni

WU Vienna University
of Economics and Business

Gregor Kastner

University of Klagenfurt

Abstract

Stochastic volatility (SV) models are nonlinear state-space models that enjoy increasing popularity for fitting and predicting heteroskedastic time series. However, due to the large number of latent quantities, their efficient estimation is non-trivial and software that allows to easily fit SV models to data is rare. We aim to alleviate this issue by presenting novel implementations of five SV models delivered in two R packages. Several unique features are included and documented. As opposed to previous versions, `stochvol` is now capable of handling linear mean models, conditionally heavy tails, and the leverage effect in combination with SV. Moreover, we newly introduce `factorstochvol` which caters for multivariate SV. Both packages offer a user-friendly interface through the conventional R generics and a range of tailor-made methods. Computational efficiency is achieved via interfacing R to C++ and doing the heavy work in the latter. In the paper at hand, we provide a detailed discussion on Bayesian SV estimation and showcase the use of the new software through various examples.

Keywords: Bayesian inference, state-space model, heteroskedasticity, dynamic correlation, dynamic covariance, factor stochastic volatility, Markov chain Monte Carlo (MCMC), leverage effect, asymmetric return distribution, heavy tails, financial time series.

Preface

This vignette corresponds to a pre-print of the same name (Hosszejni and Kastner 2019b). At the time of this writing, the vignette excludes the fifth chapter in order to reduce dependency on other packages, and the version at hand might receive minor updates as time goes by. Further information about citing `stochvol` can be obtained in R by installing the package, e.g., through `install.packages("stochvol")`, and calling `citation("stochvol")`.

1. Introduction

Time dependent variance is an indispensable ingredient of financial and economic time series modeling. Already Markowitz (1952) concerns himself with methods that take into account heteroskedasticity in a better way than a rolling window estimation. By 1982, two fundamentally different approaches had been developed to cater to these needs. On the one hand, Engle (1982) lays the groundwork for a family of time varying volatility models, most notably the

generalized autoregressive conditional heteroskedasticity model (GARCH, [Bollerslev 1986](#)). These models feature conditionally deterministic changes in the variance. [Taylor \(1982\)](#), on the other hand, addresses heteroskedasticity in his seminal work with a non-linear latent state space model, later coined the stochastic volatility (SV) model. There, the volatility process evolves in a stochastic manner. Despite some empirical evidence in favor of SV models over their corresponding GARCH counterparts ([Jacquier, Polson, and Rossi 1994](#); [Ghysels, Harvey, and Renault 1996](#); [Kim, Shephard, and Chib 1998](#); [Nakajima 2012](#)), SV and its variants enjoy little publicity among practitioners. As [Bos \(2012\)](#) underlines, one reason for this might be the lack of standard software. In response, [Kastner \(2016\)](#) provides a first version of the R ([R Core Team 2020](#)) package **stochvol** but fails to feature conditional non-Gaussianity, asymmetry (the so-called leverage effect), and multivariate generalizations.

We address these shortcomings in the manuscript at hand. First, we extend **stochvol** ([Hosszejni and Kastner 2020](#)) with several practically relevant univariate methods. Second, we introduce the new companion package **factorstochvol** ([Kastner and Hosszejni 2020](#)) which focuses on the multivariate case. The extended **stochvol** now provides the means for the Bayesian estimation of vanilla SV, heavy-tailed SV, SV with leverage, and heavy-tailed SV with leverage ([Harvey and Shephard 1996](#); [Omori, Chib, Shephard, and Nakajima 2007](#); [Nakajima and Omori 2012](#)). Moreover, the package also handles these models naturally when embedded into a linear model or an autoregressive (AR) context. The **factorstochvol** package implements an efficient method for the Bayesian estimation of the factor SV model ([Kastner, Frühwirth-Schnatter, and Lopes 2017](#)). Among other features, the package provides several automatic factor identification schemes, hierarchical shrinkage priors (variations of the normal gamma prior, [Griffin and Brown 2010](#)), and an array of intuitive visualization methods for the high-dimensional posteriors.

The remainder of this paper is structured as follows. We formally introduce the univariate and the multivariate models in Sections 2 and 3, respectively, including a discussion about prior distributions and a brief overview of the estimation methods. In Section 4, we unveil the new samplers of the **stochvol** package through three example models. We describe the **factorstochvol** package in Section 5, and then we conclude.

2. Univariate SV models

We begin by introducing the vanilla SV model with linear regressors, henceforth simply called the SV model. This is a minor but important extension of the SV model without regressors. We also settle the notation and establish a baseline model that we generalize and reuse throughout the manuscript. Consequently, we proceed with three generalized models: the SV model with Student's t errors (SVt), the SV model with leverage (SVl), and their combination, the SV model with Student's t errors and leverage (SVtl). Finally, we close the section after discussing prior distributions and Markov chain Monte Carlo (MCMC) sampling.

2.1. Model specifications

The key feature of the SV model is its stochastic and time-varying specification of the variance evolution. In particular, the log-variance is assumed to follow an AR(1) process. This feature unites the following models.

Vanilla SV with linear regressors

Let $\mathbf{y} = (y_1, \dots, y_n)^\top$ denote a vector of observations. The SV model assumes the following structure for \mathbf{y} ,

$$\begin{aligned} y_t &= \mathbf{x}_t \boldsymbol{\beta} + \exp(h_t/2) \varepsilon_t, \\ h_{t+1} &= \mu + \varphi(h_t - \mu) + \sigma \eta_t, \\ \varepsilon_t &\sim \mathcal{N}(0, 1), \\ \eta_t &\sim \mathcal{N}(0, 1), \end{aligned} \tag{1}$$

where $\mathcal{N}(b, B)$ denotes the normal distribution with mean $b \in \mathbb{R}$ and variance $B \in \mathbb{R}^+$, and ε_t and η_t are independent. The log-variance process $\mathbf{h} = (h_1, \dots, h_n)^\top$ is initialized by $h_0 \sim \mathcal{N}(\mu, \sigma^2/(1 - \varphi^2))$. $\mathbf{X} = (\mathbf{x}_1^\top, \dots, \mathbf{x}_n^\top)^\top$ is an $n \times K$ matrix containing in its t th row the vector of K regressors at time t . The K regression coefficients are collected in $\boldsymbol{\beta} = (\beta_1, \dots, \beta_K)^\top$. We refer to $\boldsymbol{\vartheta} = (\mu, \varphi, \sigma)$ as the SV parameters: μ is the level, φ is the persistence, and σ (also called *volvol*) is the standard deviation of the log-variance.

SV with Student's t errors

Several authors have suggested to use non-normal conditional residual distributions for stochastic volatility modeling. Examples include the Student's t distribution (Harvey, Ruiz, and Shephard 1994), the extended generalized inverse Gaussian (Silva, Lopes, and Migon 2006), (semi-)parametric residuals (Jensen and Maheu 2010; Delatola and Griffin 2011), or the generalized hyperbolic skew Student's t distribution (Nakajima and Omori 2012). We implement Student's t errors for the observation equation in **stochvol**. Formally,

$$\begin{aligned} y_t &= \mathbf{x}_t \boldsymbol{\beta} + \exp(h_t/2) \varepsilon_t, \\ h_{t+1} &= \mu + \varphi(h_t - \mu) + \sigma \eta_t, \\ \varepsilon_t &\sim t_\nu(0, 1), \\ \eta_t &\sim \mathcal{N}(0, 1), \end{aligned} \tag{2}$$

where ε_t and η_t are independent. $t_\nu(a, b)$ is the Student's t distribution with ν degrees of freedom, mean a , and variance b . The single difference between Equation 1 and Equation 2 is that here the observations are conditionally t distributed. Hence, Equation 2 generalizes Equation 1 through the new parameter ν as the Student's t distribution converges in law to the standard normal distribution when ν goes to infinity.

SV with leverage

Propositions for asymmetric innovations include non-parametric distributions (Jensen and Maheu 2014), skewed distributions (Nakajima and Omori 2012), and distributions featuring correlation with the variance process, also called the leverage effect (Harvey and Shephard 1996; Jacquier, Polson, and Rossi 2004). We implement the leverage effect in the **stochvol** package. Formally,

$$\begin{aligned} y_t &= \mathbf{x}_t \boldsymbol{\beta} + \exp(h_t/2) \varepsilon_t, \\ h_{t+1} &= \mu + \varphi(h_t - \mu) + \sigma \eta_t, \\ \varepsilon_t &\sim \mathcal{N}(0, 1), \\ \eta_t &\sim \mathcal{N}(0, 1), \end{aligned} \tag{3}$$

where the correlation matrix of (ε_t, η_t) is

$$\Sigma^\rho = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}. \quad (4)$$

The vector $\zeta = (\mu, \varphi, \sigma, \rho)^\top$ collects the SV parameters. The new parameter compared to Equation 1 is a correlation term ρ which relates the residuals of the observations to the innovations of the variance process. Equation 1 is therefore a special case of Equation 3 with a pre-fixed $\rho = 0$.

SV with Student's t errors and leverage

Some authors have proposed the combination of t errors with the leverage effect (Jacquier *et al.* 2004; Omori *et al.* 2007; Nakajima and Omori 2009). We implement the common generalization of Equation 2 and Equation 3. Formally,

$$\begin{aligned} y_t &= \mathbf{x}_t \boldsymbol{\beta} + \exp(h_t/2) \varepsilon_t, \\ h_{t+1} &= \mu + \varphi(h_t - \mu) + \sigma \eta_t, \\ \varepsilon_t &\sim t_\nu(0, 1), \\ \eta_t &\sim \mathcal{N}(0, 1), \end{aligned} \quad (5)$$

where the correlation matrix of (ε_t, η_t) is Σ^ρ as in Equation 4.

2.2. Prior distributions

We a priori assume $\boldsymbol{\beta} \sim \mathcal{N}_K(\mathbf{b}_\beta, \mathbf{B}_\beta)$, where $\mathcal{N}_l(\mathbf{b}, \mathbf{B})$ is the l -dimensional normal distribution with mean vector \mathbf{b} and variance-covariance matrix \mathbf{B} . For small values in the diagonal of \mathbf{B}_β , this prior enforces shrinkage towards \mathbf{b}_β ; for large values in the diagonal, the prior turns rather uninformative. By setting \mathbf{b}_β to the zero vector and \mathbf{B}_β to a scaled identity matrix, the prior distribution becomes the Bayesian analogue to ridge regression (see, e.g., Park and Casella 2008, for a discussion of this and other shrinkage priors).

The level $\mu \in \mathbb{R}$ is unrestricted, hence we can apply the common $\mu \sim \mathcal{N}(b_\mu, B_\mu)$ prior. Depending on the application, a fairly uninformative distribution is the usual choice, e.g., setting $b_\mu = 0$ and $B_\mu \geq 100$ for daily asset log returns. In our experience, the exact values of the prior mean and prior variance of μ do not strongly affect the estimation results unless B_μ is small.

To achieve stationarity in the variance process, a restricted persistence $\varphi \in (-1, 1)$ is needed. To this end, we assume $(\varphi + 1)/2 \sim \mathcal{B}(a_\varphi, b_\varphi)$, where $\mathcal{B}(a_\varphi, b_\varphi)$ is the beta distribution with shape parameters a_φ and b_φ . The selection of the shape parameters may be relatively influential with many data sets. In financial applications with daily asset log returns, the variance tends to be highly persistent, i.e., $\varphi \approx 1$. Such domain knowledge can be used as prior information by allocating more probability to positive high values of φ , e.g., by setting $a_\varphi \gtrsim 5$ and $b_\varphi \approx 1.5$. As an alternative, when stationarity is not assumed, the untruncated prior $\varphi \sim \mathcal{N}(b_\varphi, B_\varphi)$ can also be applied.

The volvol is positive but we would like allow σ to approach 0 as closely as needed – that allows us to be less informative and to improve the estimates. Following Frühwirth-Schnatter and Wagner (2010) and Kastner and Frühwirth-Schnatter (2014), we advocate

$\sigma \sim |\mathcal{N}(0, B_\sigma)|$ instead, where $|\mathcal{N}(0, B_\sigma)|$ denotes the half normal distribution. It corresponds to $\sigma^2 \sim \mathcal{G}(1/2, 1/(2B_\sigma))$, where $\mathcal{G}(a, b)$ is the gamma distribution with shape parameter a and rate parameter b . As an alternative, the commonly applied and convenient conjugate gamma prior on σ^{-2} can be assumed. However, it bounds σ away from 0 and it is therefore in our view an unsatisfactory choice.

As a last step in fully specifying the vanilla SV model in Equation 1, the variance process is initialized a priori with its stationary distribution, i.e., $h_0 \sim \mathcal{N}(\mu, \sigma^2/(1 - \varphi^2))$. This consistently extends our prior assumptions about \mathbf{h} following a stationary AR(1) process. As an alternative, when stationarity is not assumed, $h_0 \sim \mathcal{N}(\mu, B_h)$ can be applied with a constant variance B_h .

The SV models with Student's t errors additionally require the prior specification of the degrees of freedom parameter ν . To ascertain interpretability of the scaling $\exp(h_t/2)$, we ensure finite second moments of \mathbf{y} by enforcing $\nu > 2$. As a reviewer recommended, we follow (Geweke 1993) and equip ν with an exponential prior $\nu - 2 \sim \mathcal{E}(\lambda_\nu)$, where λ_ν is the rate of the exponential distribution.

Finally, in the case of the SV models with leverage, we employ the translated and scaled beta distribution for $\rho \in (-1, 1)$ as in Omori *et al.* (2007), i.e., $(\rho + 1)/2 \sim \mathcal{B}(a_\rho, b_\rho)$. We find that the posterior estimates of ρ can be sensitive to its prior distribution, thus, some care is needed when setting the hyperparameters in practice. In our experience, slightly informative choices such as $a_\rho = b_\rho \approx 4$ work well in financial applications.

2.3. Estimation

All methods implemented in **stochvol** and **factorstochvol** rely on the Bayesian paradigm.¹ Bayesian analysis aims to estimate model parameters through Bayesian updating. By using probability distributions to represent information, Bayes' theorem can be employed to update the prior information to the posterior information by incorporating the observations. This approach has the advantage of providing full uncertainty quantification in a probabilistic framework without relying on asymptotic results; moreover, so-called *shrinkage priors* can be used to regularize the posterior and guard against overfitting. For an introductory textbook on Bayesian statistics, see, for instance, McElreath (2015).

When the posterior distribution is not available analytically, one customarily resorts to approximations such as perfect simulation (Huber 2015), approximate Bayesian computation (Sisson, Fan, and Beaumont 2018), adaptive Monte Carlo methods (Roberts and Rosenthal 2007), or MCMC methods. When computationally feasible, MCMC is a valuable tool that provides draws from the posterior distribution in question. That way, MCMC approximates the posterior distribution similarly to a histogram approximating a density. For a more in-depth introduction on MCMC methods, see, for instance, Brooks, Gelman, Jones, and Meng (2011).

The estimation algorithm of SV, SVt, SVl, and SVtl all resemble the original methodology developed in Kastner and Frühwirth-Schnatter (2014) for the vanilla SV model. Namely, to draw from the posterior distribution of \mathbf{h} efficiently, the MCMC sampler resorts to approximate mixture representations of Equations 1, 2, 3 and 5 similar to the ones in Kim *et al.*

¹At this point we would also like to point out works aiming at estimating stochastic volatility and related models within the frequentist framework, see, e.g., Abanto-Valle, Langrock, Chen, and Cardoso (2017); Creal (2017), and in particular the recent **stochvolTMB** package (Wahl 2020).

(1998) and Omori *et al.* (2007). Doing so yields a conditionally Gaussian state space model for which efficient sampling methods are available (Frühwirth-Schnatter 1994; Carter and Kohn 1994). Following Rue (2001) and McCausland, Miller, and Pelletier (2011), we draw the full vector \mathbf{h} “all without a loop” (AWOL).

When Student’s t errors with unknown degrees of freedom are used, we handle the added complication through the well-known representation of the t distribution as a scale mixture of Gaussians. This requires additional Gibbs and independence Metropolis-Hastings steps documented in Kastner (2015). Furthermore, we deal with the increased complexity in the posterior space of the leverage case by repeated ancillarity-sufficiency interweaving strategies (ASIS, Yu and Meng 2011) steps in the sampling scheme, see Hosszejni and Kastner (2019a) for details.

To verify the correctness of the implementation, unit tests are included in the package which can be run by `devtools::test()` (Wickham, Hester, and Chang 2020). In particular, a variant of Geweke’s test (Geweke 2004) is part of the test suite. In this test, we exploit that the sampling distribution of the model parameters during the Geweke test is identical to their preset prior distribution. Therefore, the cumulative distribution function maps the sample to a uniform distribution, which in turn is mapped to a normal distribution using the normal distribution’s quantile function. If the user chooses to execute the automated unit tests in `stochvol`, the system evaluates the thinned and transformed sample using the `shapiro.test()` function, where the thinning of the sample is done to approximate independent sampling.

For maximal computational effectiveness, all sampling algorithms are implemented in the compiled language C++ (ISO/IEC 2017) with the help of the R package `Rcpp` (Eddelbuettel and François 2011). Matrix computations make use of the efficient C++ template library `Armadillo` (Sanderson and Curtin 2016) through the R package `RcppArmadillo` (Eddelbuettel and Sanderson 2014).² After sampling, the resulting R objects make use of plotting and summary functions of the R package `coda` (Plummer, Best, Cowles, and Vines 2006).

3. Multivariate SV models

A key difficulty accompanying dynamic covariance estimation is the relatively high number of unknowns compared to the number of observations. More precisely, letting m denote the cross-sectional dimension, the corresponding covariance matrix Σ_t contains $m(m+1)/2$ degrees of freedom, a quadratic term in m . Table 1 illustrates the “curse of dimensionality” for various values of m . One way to break this curse is to use latent factors and thereby achieve a sparse representation of Σ_t .

3.1. The factor SV model

Latent factor models embody the idea that even high dimensional systems are driven by only a few sources of randomness. These few sources of randomness control a few factors, which in turn account for the interactions between the observations. Moreover, latent factor models provide an efficient tool for dynamic covariance matrix estimation. They allow for a reduction in the number of unknowns. A conventional latent factor model with r factors implies the

²For explicit run time discussions please see Kastner and Frühwirth-Schnatter (2014) and Hosszejni and Kastner (2019a). For the possibility to use multi-core computing within a single MCMC chain and potential speed gains when doing so, please see Kastner (2019).

m	free elements of Σ_t	free elements of Σ_t per data point
1	1	1
10	55	5.5
100	5050	50.5
1000	500500	500.5

Table 1: Absolute and relative numbers of free elements of the time-varying covariance matrix Σ_t for different numbers of component series m .

m	free elements of Σ_t	free elements of Σ_t per data point
10	44	4.4
100	494	4.94
1000	4994	4.994

Table 2: Absolute and relative numbers of free elements of the time-varying covariance matrix Σ_t in a factor model for different numbers of component series m and number of factors $r = 4$.

decomposition

$$\Sigma_t = \check{\Sigma}_t + \bar{\Sigma}_t, \quad (6)$$

where $\text{rank}(\check{\Sigma}_t) = r < m$, and $\bar{\Sigma}_t$ is the diagonal matrix containing the variances of the idiosyncratic errors. The rank assumption on the symmetric $\check{\Sigma}_t$ gives rise to the factorization $\check{\Sigma}_t = \Psi\Psi^\top$, where $\Psi \in \mathbb{R}^{m \times r}$ contains $mr - r(r-1)/2$ free elements (see, e.g., the pivoted Cholesky algorithm in Higham 1990). Hence, $m(r+1) - r(r-1)/2$ free elements remain in Σ_t , now only linear in m . Table 2 illustrates the ‘‘broken curse of dimensionality’’ for various values of m and $r = 4$.

In the following, we describe the factor SV model employed in the `factorstochvol` package. We model the observations $\mathbf{y}_t = (y_{t1}, \dots, y_{tm})^\top$ as follows.

$$\begin{aligned} \mathbf{y}_t \mid \boldsymbol{\beta}, \boldsymbol{\Lambda}, \mathbf{f}_t, \bar{\Sigma}_t &\sim \mathcal{N}_m(\boldsymbol{\beta} + \boldsymbol{\Lambda}\mathbf{f}_t, \bar{\Sigma}_t), \\ \mathbf{f}_t \mid \check{\Sigma}_t &\sim \mathcal{N}_r(\mathbf{0}, \check{\Sigma}_t), \end{aligned} \quad (7)$$

where $\mathbf{f}_t = (f_{t1}, \dots, f_{tr})^\top$ is the vector of factors, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_m)^\top$ is an observation-specific mean, and $\boldsymbol{\Lambda} \in \mathbb{R}^{m \times r}$ is a tall matrix holding the factor loadings. The covariance matrices $\bar{\Sigma}_t$ and $\check{\Sigma}_t$ are both diagonal representing independent vanilla SV processes.

$$\begin{aligned} \bar{\Sigma}_t &= \text{diag}(\exp(\bar{h}_{t1}), \dots, \exp(\bar{h}_{tm})), \\ \check{\Sigma}_t &= \text{diag}(\exp(\tilde{h}_{t1}), \dots, \exp(\tilde{h}_{tr})), \\ \bar{h}_{ti} &\sim \mathcal{N}(\bar{\mu}_i + \bar{\varphi}_i(\bar{h}_{t-1,i} - \bar{\mu}_i), \bar{\sigma}_i^2), \quad i = 1, \dots, m, \\ \tilde{h}_{tj} &\sim \mathcal{N}(\tilde{\mu}_j + \tilde{\varphi}_j(\tilde{h}_{t-1,j} - \tilde{\mu}_j), \tilde{\sigma}_j^2), \quad j = 1, \dots, r. \end{aligned} \quad (8)$$

For a more theoretical treatment of factor SV from a Bayesian point of view, the reader is referred to, e.g., Pitt and Shephard (1999), Aguilar and West (2000), Chib, Nardari, and Shephard (2006), and Han (2006).

Based on Equation 7, we can reformulate Equation 6 as

$$\Sigma_t = \boldsymbol{\Lambda}\check{\Sigma}_t\boldsymbol{\Lambda}^\top + \bar{\Sigma}_t, \quad (9)$$

from which several identification issues are apparent: the order, the sign, and the scale of the factors is unidentified. More specifically, for any generalized permutation matrix³ \mathbf{P} of size $r \times r$, we find another valid decomposition $\boldsymbol{\Sigma}_t = \boldsymbol{\Lambda}' \tilde{\boldsymbol{\Sigma}}_t (\boldsymbol{\Lambda}')^\top + \bar{\boldsymbol{\Sigma}}_t$, where $\boldsymbol{\Lambda}' = \boldsymbol{\Lambda} \mathbf{P}^{-1}$ and $\tilde{\boldsymbol{\Sigma}}_t = \mathbf{P} \boldsymbol{\Sigma}_t \mathbf{P}^\top$. We resolve the ambiguity in the scale of the factors by fixing the level of their log-variance to zero, i.e., $\tilde{\mu}_j = 0$ for $j = 1, \dots, r$. Sign and order identification can be enforced through restrictions on the factor loadings matrix $\boldsymbol{\Lambda}$. Several options are available in **factorstochvol** for restricting $\boldsymbol{\Lambda}$, for details see Section ??.

3.2. Prior distributions

Priors need to be specified for the mean, the latent log-variance processes, and for the factor loadings matrix $\boldsymbol{\Lambda}$. We choose $\beta_j \sim \mathcal{N}(b_\beta, B_\beta)$, independently for $j = 1, \dots, m$. For small values of B_β , this shrinks β_j toward b_β ; for large values of B_β , the prior is fairly uninformative.

The log-variance processes have the same prior specification as in the univariate case in Section 2.2. For $\boldsymbol{\Lambda}$, three types of priors are currently implemented in **factorstochvol**. All three can be written in the form $\Lambda_{ij} \sim \mathcal{N}(0, \tau_{ij}^2)$ independently for each applicable $i \in \{1, \dots, m\}$ and $j \in \{1, \dots, r\}$. First, one can fix all the τ_{ij}^2 s – not necessarily to the same value – a priori. This results in a normal prior for each element of the loadings matrix.

The second type is a hierarchical prior which has been developed to induce more flexible and potentially stronger shrinkage,

$$\Lambda_{ij} \mid \tau_{ij}^2 \sim \mathcal{N}(0, \tau_{ij}^2), \quad \tau_{ij}^2 \mid \lambda_i^2 \sim \mathcal{G}(a, a\lambda_i^2/2). \quad (10)$$

This distribution is termed normal gamma prior by Griffin and Brown (2010) and implies a conditional variance $\mathbb{V}(\Lambda_{ij} \mid \lambda_i^2)$ of $2/\lambda_i^2$ and an unconditional excess kurtosis of $3/a$. The value of a is treated as a structural parameter to be fixed by the user, where choosing a small ($\lesssim 1$) enforces strong shrinkage towards zero, while choosing a large ($\gtrsim 1$) imposes little shrinkage. The case $a = 1$ is a special case termed the Bayesian Lasso prior (Park and Casella 2008). The parameter λ_i^2 is estimated from the data with $\lambda_i^2 \sim \mathcal{G}(c, d)$.

The third type is a slight modification of the second. Because variances in each row of the factor loadings matrix $\boldsymbol{\Lambda}$ can be seen as “random effects” from the same underlying distribution, the prior in Equation 10 induces row-wise shrinkage with element-wise adaption. Analogously, one could also consider column-wise shrinkage with element-wise adaption, i.e.,

$$\Lambda_{ij} \mid \tau_{ij}^2 \sim \mathcal{N}(0, \tau_{ij}^2), \quad \tau_{ij}^2 \mid \lambda_j^2 \sim \mathcal{G}(a, a\lambda_j^2/2), \quad (11)$$

with the corresponding prior $\lambda_j^2 \sim \mathcal{G}(c, d)$.

3.3. Estimation

Bayesian estimation in the factor SV model builds on the univariate vanilla SV implementations in **stochvol** and features several levels of efficiency boosting. To alleviate the problem of potentially slow convergence in high dimensions, it is carried out via a sampler that utilizes several variants of ASIS. The sampling details implemented in **factorstochvol** are described in

³A generalized permutation matrix has the zero–non-zero pattern of a permutation matrix, but it is allowed to have any non-zero values instead of just ones. Hence, a generalized permutation matrix not only permutes but also scales and switches the sign of its multiplier.

Kastner *et al.* (2017, using Gaussian priors for the factor loadings) as well as Kastner (2019, using hierarchical shrinkage priors for the factor loadings).

Similarly to **stochvol** and in an attempt to make computation time bearable even in higher dimensions, **factorstochvol**'s main sampler is written in C++. It uses the R package **Rcpp** to ease communication between R and C++. The univariate SV parts are borrowed from **stochvol** and interfaced through its C/C++-level updating function `update_fast_sv()`. In doing so, moving between interpreted R code and compiled C++ code at each MCMC iteration is avoided.

4. The stochvol package

The **stochvol** package provides means for fitting univariate SV, SVt, SVI, and SVtl models via its sampling routines `svsample()`, `svtsample()`, `svlsample()`, and `svtlsample()`, respectively. In the following, we describe a recommended workflow with **stochvol**. First, we discuss estimation, visualization, and prediction using default settings. Then, we show how to adapt the values of the prior hyperparameters and how to configure the sampling mechanism.

4.1. Preparing the data and running the MCMC sampler

We estimate three models that exemplify the features and the user interface of **stochvol**. Using the `exrates` data found in the package, we model the EURCHF exchange rate (the price of 1 euro in Swiss franc) in the period between March 1, 2008 and March 1, 2012 (1028 data points) in three different ways.

AR(1) model with SV residuals

The first example is an AR(1) model with SV residuals, i.e., Equation 1 turns into

$$\begin{aligned} y_t \mid y_{t-1}, \beta_0, \beta_1, h_t &\sim \mathcal{N}(\beta_0 + \beta_1 y_{t-1}, \exp(h_t)), \\ h_{t+1} \mid \boldsymbol{\vartheta}, h_t &\sim \mathcal{N}(\mu + \varphi(h_t - \mu), \sigma^2). \end{aligned}$$

Using this model, we test whether the exchange rate follows a random walk with SV. In this case, we expect the posteriors of β_0 and β_1 to concentrate around 0 and 1, respectively.

In order to estimate this AR(1)-SV model, we need to prepare the input \mathbf{y} as a numeric sequence of length n and pass it as the first input argument to `svsample()` as follows:

```
R> set.seed(1)
R> library("stochvol")
R> data("exrates")
R> ind <- which(exrates$date >= as.Date("2008-03-01") &
+   exrates$date <= as.Date("2012-03-01"))
R> CHF_price <- exrates$CHF[ind]
R> res_sv <- svsample(CHF_price, designmatrix = "ar1")
```

We set `designmatrix = "ar1"` to use the AR(1) specification. More generally, `designmatrix` can take character values of the form "ar0" for a constant mean model, or "ar1", "ar2", etc., for AR(1), AR(2), and so on.

Constant mean model with SVt residuals

The second example is a constant mean model with SVt residuals, i.e., Equation 2 becomes

$$\begin{aligned} y_t \mid \beta_0, h_t, \nu &\sim t_\nu(\beta_0, \exp(h_t/2)), \\ h_{t+1} \mid \vartheta, h_t &\sim \mathcal{N}(\mu + \varphi(h_t - \mu), \sigma^2). \end{aligned}$$

If the returns are heavy-tailed, most of the posterior mass of ν concentrates on low values, e.g., smaller than 20. Otherwise, there is little evidence for high kurtosis.

We compute the log returns by applying `logret()` on the previously calculated `CHF_price`. Then, to estimate the constant mean model with heavy tailed SV residuals, we pass the vector of log returns to `svtsample()` with `designmatrix` set to "ar0".

```
R> set.seed(2)
R> CHF_logret <- 100 * logret(CHF_price)
R> res_svt <- svtsample(CHF_logret, designmatrix = "ar0")
```

Multiple regression with SVI residuals

The third example is a multiple regression model with an intercept, two regressors, and SVI residuals; that is, Equation 3 turns into

$$\begin{aligned} \begin{pmatrix} y_t \\ h_{t+1} \end{pmatrix} \mid h_t, \zeta, \begin{pmatrix} x_{t1} \\ x_{t2} \end{pmatrix}, \beta_0, \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} &\sim \mathcal{N}_2 \left(\begin{pmatrix} \beta_0 + \beta_1 x_{t1} + \beta_2 x_{t2} \\ \mu + \varphi(h_t - \mu) \end{pmatrix}, \Sigma^\rho \right), \\ \Sigma^\rho &= \begin{pmatrix} \exp(h_t) & \rho\sigma \exp(h_t/2) \\ \rho\sigma \exp(h_t/2) & \sigma^2 \end{pmatrix}. \end{aligned}$$

For illustration, we regress EURCHF log returns onto the contemporaneous log returns on EURUSD and EURJPY, the value of 1 euro per US dollar and Japanese yen, respectively.

To estimate a multiple regression model using `stochvol`, we need to prepare a `numeric` matrix \mathbf{X} of dimension $n \times K$, where rows correspond to time points and columns to covariates. We create an intercept as the first column of \mathbf{X} , and we set the second and the third columns to the EURUSD log returns and the EURJPY log returns, respectively; finally, we use the columns of \mathbf{X} as covariates in the multiple regression.

```
R> set.seed(3)
R> X <- cbind(constant = 1,
+ 100 * logret(exrates$USD[ind]),
+ 100 * logret(exrates$JPY[ind]))
R> res_svl <- svlsample(CHF_logret, designmatrix = X, thin = 10)
```

4.2. Visualizing the results

Often, the joint posterior distribution of model parameters and latent quantities mark the goal of a Bayesian analysis. To inspect it, one can look at summary statistics and various types of visualizations of marginal posterior distributions. Also, it is recommended to examine the

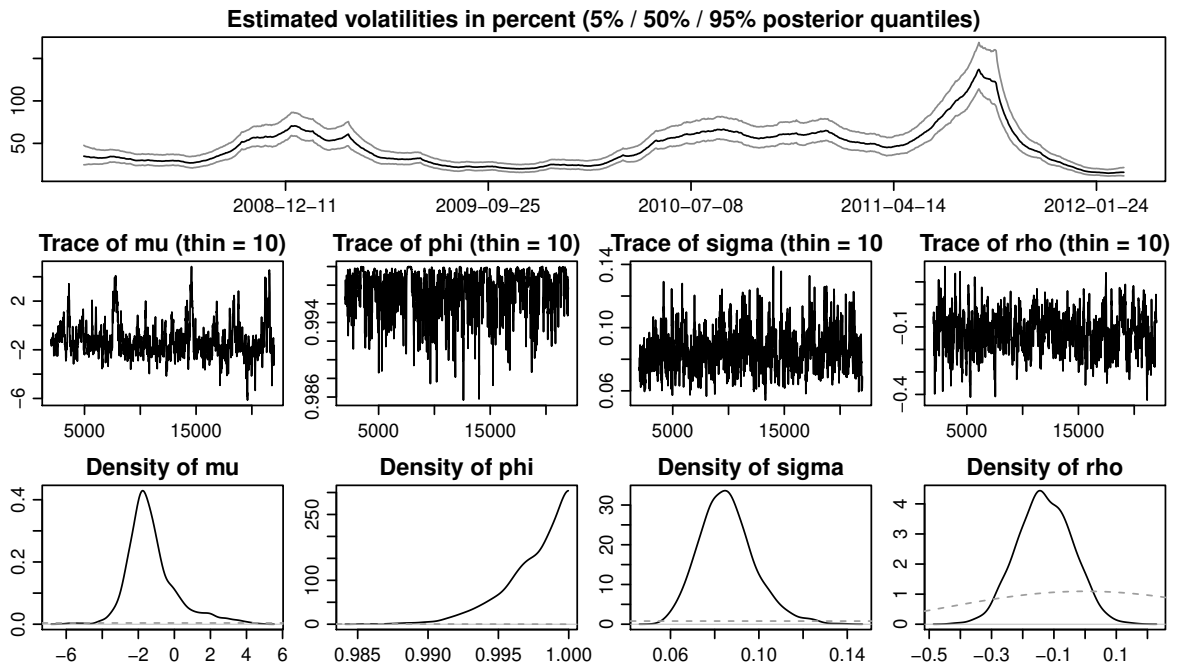


Figure 1: The default plot of an estimated model. The top row shows a summary of the posterior of the daily volatility (in percent) $100 \exp(\mathbf{h}/2)$ through its median (black) and 5% and 95% quantiles (gray). The remaining panels summarize the Markov chains of the parameters μ , φ , σ , and ρ . In particular, the middle row presents trace plots and the bottom row shows prior (gray, dashed) and posterior (black, solid) densities.

Markov chain for possible convergence issues – this happens usually by investigating trace plots of posterior quantities. For this reason, inspired by the `coda` package, `stochvol` provides its own instances of the R generic functions `plot()` and `summary()`. In order to introduce the tools that `stochvol` provides for analyzing MCMC output, we briefly examine the results of the third example (multiple regression with SVI errors) in the remaining part of the section. First, we plot the output of the estimation.

```
R> plot(res_svl, showobs = FALSE,
+       dates = exrates$date[ind[-1]])
```

The result is shown in Figure 1. We see in the first row a summary of the posterior density of the volatility. Apart from its median, we also receive a quantification of the uncertainty through the 5% and the 95% quantiles at each time point. In the second row, we can follow the evolution of the Markov chain of the SV parameters. In this example, they are μ , φ , σ , and ρ . Lastly, we see prior and posterior density plots of the parameters in the third row in gray and black, respectively. They show high persistence and significant leverage.

Next, we observe the AR coefficients.

```
R> for (i in seq_len(3)) {
+   coda::traceplot(svbeta(res_svl)[, i])
+   coda::densplot(svbeta(res_svl)[, i], show.obs = FALSE)
+ }
```

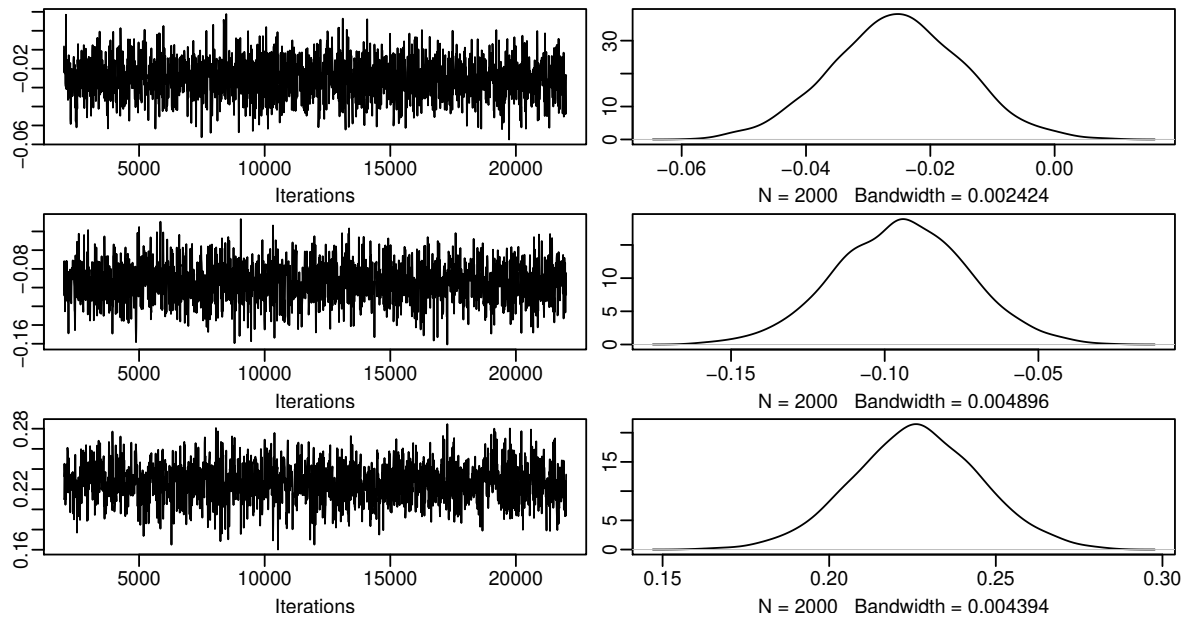


Figure 2: Trace plots and estimated kernel densities of posterior draws from $p(\boldsymbol{\beta} \mid \mathbf{y})$.

The result is shown in Figure 2. On the left hand side, we do not spot any signs of convergence or mixing problems in the trace plots. On the right hand side, we see that none of the posterior densities of β_0 , β_1 , and β_2 concentrate around 0, hence the covariates seem to have an impact on the dependent variable.

As the final step, we print a numeric summary of the estimation results.

```
R> summary(res_svl, showlatent = FALSE)
```

Summary of 'svdraws' object

Prior distributions:

```
mu          ~ Normal(mean = 0, sd = 100)
(phi+1)/2  ~ Beta(a = 5, b = 1.5)
sigma^2    ~ Gamma(shape = 0.5, rate = 0.5)
nu         ~ Infinity
rho        ~ Beta(a = 4, b = 4)
beta       ~ MultivariateNormal(...)
```

Stored 2000 MCMC draws after a burn-in of 2000.

Thinning: 10.

Posterior draws of SV parameters (thinning = 10):

	mean	sd	5%	50%	95%	ESS
mu	-1.2530	1.4071	-2.9699	-1.5307	1.752	99
phi	0.9973	0.0023	0.9928	0.9979	1.000	292

```

sigma      0.0855 0.0124  0.0670  0.0847 0.108 311
rho        -0.1291 0.0877 -0.2735 -0.1313 0.014 236
exp(mu/2)  0.7461 0.9622  0.2265  0.4652 2.401  99
sigma^2    0.0075 0.0022  0.0045  0.0072 0.012 311

```

Posterior draws of regression coefficients (thinning = 10):

	mean	sd	5%	50%	95%	ESS
beta_0	-0.025	0.010	-0.042	-0.025	-0.0082	2109
beta_1	-0.095	0.021	-0.129	-0.094	-0.0595	1545
beta_2	0.226	0.019	0.195	0.226	0.2577	1206

For brevity, we set `showlatent = FALSE` in order not to print all the 1027 latent states. The output shows the length of the burn-in and the number of draws, the prior specification of the parameters, and a concise summary of the marginal posterior distributions of the parameters μ , φ , σ , and ρ , and additionally of the level of the volatility $\exp(\mu/2)$ and of σ^2 , and of the vector of regression coefficients β . This posterior summary is a table consisting of columns for the posterior mean and standard deviation, the 5%, 50%, and 95% quantiles. The user can influence the shown quantiles by passing a sequence of values between 0 and 1 to `svsample()`, `svtsample()`, `svlsample()`, or `svtllsample()` via the argument `quantiles`.

The last column in the table depicts the so-called effective sample size (ESS), a measure of the quality of a converged MCMC chain. Formally, ESS of a Markov chain C is defined through $M/(1 + 2\sum_{s=1}^{\infty} \rho^{\text{eff}}(s))$, where M is the length of C and $\rho^{\text{eff}}(s)$ denotes the autocorrelation function for lag s among the elements of C . In principle, ESS is the sample size of a serially uncorrelated chain bearing the same Monte Carlo error as our (marginal) chain. Intuitively speaking, this means that ESS is the number of independent and identically distributed draws that were acquired and gives a sense of how well our chain has explored the posterior space. Higher values of ESS indicate better mixing.

4.3. Prediction with `stochvol`

We employ our estimated model to predict log returns for the remaining days in the data set. To do so, we first prepare the covariates for the next 24 days and pass them via the argument `newdata` of the generic `predict()` function along with the estimation output. Note that we need 25 days of price data to obtain 24 returns.

```

R> set.seed(4)
R> pred_ind <- seq(tail(ind, 1), length.out = 25)
R> pred_X <- cbind(constant = 1,
+   100 * logret(exrates$USD[pred_ind]),
+   100 * logret(exrates$JPY[pred_ind]))
R> pred_svl <- predict(res_svl, 24, newdata = pred_X)

```

As we have access to the entire distribution of future log returns, we can quantify the uncertainty around our predictions through quantiles. In the following code snippet, we visualize the k -step-ahead predictive distributions for $k = 1, \dots, 24$, along with the truly observed values. The result is in Figure 3.

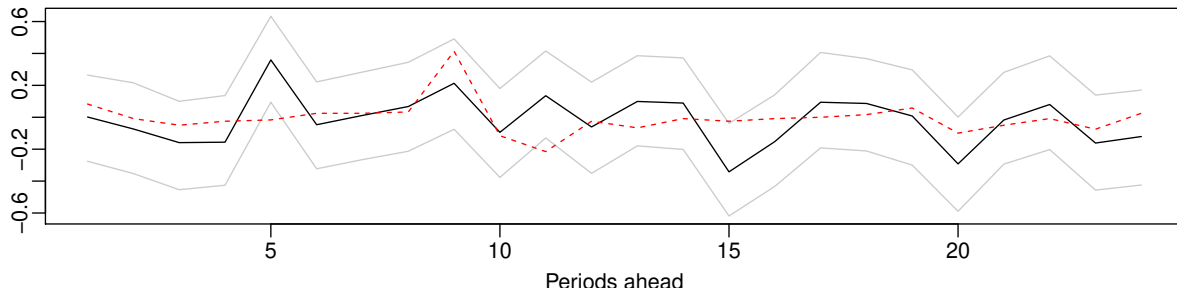


Figure 3: Multi-step ahead predictive distributions (solid, gray and black) and observations (dashed, red).

```
R> opar <- par(mgp = c(1.7, 0.5, 0))
R> obs_CHF <- 100 * logret(exrates$CHF[pred_ind])
R> ts.plot(cbind(t(apply(predy(pred_svl), 2, quantile, c(0.05, 0.5, 0.95))),
+   obs_CHF), xlab = "Periods ahead", lty = c(rep(1, 3), 2),
+   col = c("gray80", "black", "gray80", "red"))

R> par(opar)
```

4.4. Rolling window estimation

Inspired by `ugarchroll()` in the R package **rugarch** (Ghalanos 2020), we introduce the suite of wrapper functions `svsample_roll()`, `svtsample_roll()`, `svlsample_roll()`, and `svtlsample_roll()`, built around their corresponding routines `svsample()`, `svtsample()`, `svlsample()`, and, respectively, `svtlsample()`, to simplify rolling window estimation of SV models. In this estimation method, either a fixed width time window is *moving* through the time series or a sequence of *expanding* time windows with the same starting time point covers larger and larger chunks of the observations, and the same model is estimated in all time windows independently. Next, each estimated model is employed for out-of-sample prediction, typically one day to one week ahead of the time window. Lastly, the set of predicted values might be used to evaluate the model fit.

In Bayesian statistics, a natural approach for assessing the predictive power of a model is through its *posterior predictive distribution*. Its density, also called the *predictive density*, is defined as

$$p(y_{t+1} \mid \mathbf{y}_{[1:t]}^o) = \int_{\mathcal{K}} p(y_{t+1} \mid \mathbf{y}_{[1:t]}^o, \boldsymbol{\kappa}) p(\boldsymbol{\kappa} \mid \mathbf{y}_{[1:t]}^o) d\boldsymbol{\kappa}, \quad (12)$$

where $\boldsymbol{\kappa}$ collects all unobserved variables, i.e. $\boldsymbol{\kappa} = (\mu, \varphi, \sigma, \rho, \nu, \mathbf{h}, \boldsymbol{\beta})^\top$ in the most general case of SVtl, and the domain of integration \mathcal{K} is the set of all possible values for $\boldsymbol{\kappa}$. We follow Geweke and Amisano (2010) in our notation by using a superscript \mathbf{o} for the vector of observed values $\mathbf{y}_{[1:t]}^o = (y_1, y_2, \dots, y_t)^\top$. Equation 12 can be seen as the integration of the predictive likelihood over the posterior distribution of all parameters and therefore it accounts for posterior parameter uncertainty for the predicted values.

The integral in Equation 12 has no closed form and its dimensionality increases with t ; it is intractable. Hence, we rely on Monte Carlo integration and we simulate from the posterior

predictive distribution. For the evaluation of the predictive density at an observation $x = y_{t+1}^o$, called the *predictive likelihood*, we apply the computation

$$p(x \mid \mathbf{y}_{[1:t]}^o) \approx \frac{1}{M} \sum_{m=1}^M p(x \mid \mathbf{y}_{[1:t]}^o, \boldsymbol{\kappa}^{(m)}), \quad (13)$$

where $\boldsymbol{\kappa}^{(m)}$ denotes the m th posterior sample from the estimation procedure of the SV model. For other applications, the quantiles of the posterior predictive distribution, henceforth the *predictive quantiles*, might be of interest. We estimate the $q\%$ quantile through random variates simulated from $y_{t+1} \sim p(y_{t+1} \mid \mathbf{y}_{[1:t]}^o)$, which we acquire by repeating two steps for $m = 1, \dots, M$:

Step 1. Simulate $\boldsymbol{\kappa}^{(m)}$ from the SV posterior $p(\boldsymbol{\kappa} \mid \mathbf{y}_{[1:t]}^o)$, and

Step 2. Simulate $y_{t+1}^{(m)}$ from $p(y_{t+1} \mid \mathbf{y}_{[1:t]}^o, \boldsymbol{\kappa}^{(m)})$.

Lastly, we take the $q\%$ quantile of the sample vector $(y_{t+1}^{(1)}, y_{t+1}^{(2)}, \dots, y_{t+1}^{(M)})^\top$ as the approximate $q\%$ quantile of the predictive density. We implement the estimation of both the predictive likelihood and predictive quantiles in **stochvol**.

All four rolling window routines `svsample_roll()`, `svtsample_roll()`, `svlsample_roll()`, and `svtlsample_roll()` bear the same programming interface. They expect as their first argument the input data $\mathbf{y}_{[1:L]}$, which is of length L . For estimating the SV model in each time window $j = 1, \dots, J$ in the moving or expanding window scheme, the sub-vector $\mathbf{y}_{[j:(t+j-1)]}$, or, respectively, $\mathbf{y}_{[1:(t+j-1)]}$, is taken as data and is used to predict $n_{\text{ahead}} \geq 1$ time steps ahead. The width t of the first time window can be determined from L , J , and n_{ahead} . The following example demonstrates how the rolling window sampling routines can be called in **stochvol**.

```
R> set.seed(5)
R> res <- svsample_roll(CHF_logret, n_ahead = 1,
+   forecast_length = 30,
+   refit_window = "moving",
+   calculate_quantile = c(0.01, 0.05),
+   calculate_predictive_likelihood = TRUE)
```

Argument `n_ahead` is used to set n_{ahead} , `forecast_length` is used to set J , and `refit_window` expects either "moving" or "expanding" to set the rolling window scheme to moving or expanding, respectively. Argument `calculate_quantile` expects a vector of numbers between 0 and 1; the numbers are interpreted as the quantiles to be predicted. Furthermore, if `calculate_predictive_likelihood` is set to `TRUE`, the function estimates the predictive likelihood. Lastly, the output `res` is a list of length J , i.e. one element for each time window. It contains the respective posterior quantile and predictive likelihood results together with all posterior parameter draws for $\boldsymbol{\kappa}$.

4.5. Specifying the prior hyperparameters

As discussed in Section 2.2, the prior distributions need to be specified before the estimation process can start. Concerning the common model parameters μ , φ , and σ , all of

`svsample()`, `svtsample()`, `svlsample()`, and `svtlsample()` expect through their input arguments `priormu`, `priorphi`, and `priorsigma` values for $(b_\mu, \sqrt{B_\mu})$, (a_φ, b_φ) , and B_σ , respectively. Furthermore, all sampling functions accept the argument `priorbeta` to set an independent prior for the regression coefficients by providing (b_β, s_β) , where b_β and s_β are the common mean and, respectively, the common standard deviation. For a general multivariate normal distribution, the `specify_priors()` interface exists, which we detail later in this Section. The prior for ν can be influenced in `svtsample()` and `svtlsample()` by passing λ_ν as the argument `priornu`. Finally, `svlsample()` and `svtlsample()` take the numeric sequence (a_ρ, b_ρ) through the input argument `priorrho`.

The code snippet below shows all the default values of the prior hyperparameters.

```
R> svsample(CHF_logret, priormu = c(0, 100), priorphi = c(5, 1.5),
+   priorsigma = 1, priorbeta = c(0, 10000))
R> svtsample(CHF_logret, priormu = c(0, 100), priorphi = c(5, 1.5),
+   priorsigma = 1, priorbeta = c(0, 10000), priornu = 0.1)
R> svlsample(CHF_logret, priormu = c(0, 100), priorphi = c(5, 1.5),
+   priorsigma = 1, priorbeta = c(0, 10000), priorrho = c(4, 4))
R> svtlsample(CHF_logret, priormu = c(0, 100), priorphi = c(5, 1.5),
+   priorsigma = 1, priorbeta = c(0, 10000), priornu = 0.1,
+   priorrho = c(4, 4))
```

As an alternative to the concise interface above, a broader set of prior distributions can be specified via an object created by the `specify_priors()` function. The function has an input argument for each model parameter: `mu`, `phi`, `sigma2`, `nu`, `rho`, `beta`, and additionally one for the variance of h_0 called `latent0_variance`. There is a list of accompanying functions that create distributions in `stochvol`: `sv_beta()` has arguments `shape1` and `shape2` and it is accepted for `phi` and `rho`; `sv_constant()` has argument `value` and it is accepted for `mu`, `phi`, `sigma2`, `nu`, `rho`, and `latent0_variance`; `sv_normal()` has arguments `mean` and `sd` and it is accepted for `mu` and `phi`; `sv_multinormal()` has arguments `mean` and either `sd` and `dim` or `precision`, and it is accepted for `beta`; `sv_exponential()` has argument `rate` and it accepted for `nu`; `sv_gamma()` has arguments `shape` and `rate` and it is accepted for `sigma2`; `sv_inverse_gamma()` has arguments `shape` and `scale` and it is accepted for `sigma2`; and `sv_infinity()` has no arguments and it is accepted for `nu` hence turning the Student's t distribution into a normal distribution. Additionally, `latent0_variance` accepts the character value "stationary". All four sampling methods accept the prior specification object through the input argument `priorspec`.

All input arguments for `specify_priors` are optional, their default values and how they are used is seen below.

```
R> ps <- specify_priors(
+   mu = sv_normal(mean = 0, sd = 100),
+   phi = sv_beta(shape1 = 5, shape2 = 1.5),
+   sigma2 = sv_gamma(shape = 0.5, rate = 0.5),
+   nu = sv_infinity(),
+   rho = sv_constant(0),
+   latent0_variance = "stationary",
+   beta = sv_multinormal(mean = 0, sd = 10000, dim = 1))
R> svsample(CHF_logret, priorspec = ps)
```

4.6. Setting up the Markov chain

When conducting Bayesian inference using an MCMC sampling scheme, the number of draws from the posterior distribution, the length of the so-called burn-in phase, the initial values of the Markov chain, and the various strategies of storing the results are all of general interest. The input arguments `draws` and `burnin` settle the first two points. A sample size of `burnin + draws` is acquired from the posterior distribution out of which the first `burnin` number of draws are thrown away. The default is to draw 10000 elements after a burnin of 1000 for SV models without leverage, and draw 20000 elements after a burnin of 2000 for SV models with leverage, which in our experience is enough for most applications.

As for the initial values, `startpara` and `startlatent` provide a way to set them. The argument `startpara` is expected to be a named list mapping parameter names to starting values, and `startlatent` must be a sequence of length m that contains starting values for \mathbf{h} . Default values are set to be the prior mean for φ , σ , ν , and ρ , these have only minor influence on the Markov chain. The default value for β is the ordinary least squares estimator $(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}$, where \mathbf{X} denotes the regression design matrix and \mathbf{y} denotes the vector of observations. After setting β , the level of log-variance μ is initialized according to the Bayesian linear regression

$$\begin{aligned} \log(y_t^2) &= \mu + \xi_t, \\ \mu &\sim \mathcal{N}(b_\mu, B_\mu), \end{aligned} \tag{14}$$

where $\xi_t \sim \mathcal{N}(-1.27, 4.934)$. Equation 14 results from the first line of Equation 1 by fixing h_t at its stationary expected value μ and then taking $x \mapsto \log(x^2)$ of both sides. The homoskedastic error term ξ_t is acquired as the Laplace approximation to $\log(\varepsilon_t^2)$ (Harvey and Shephard 1996). At the end, by default all values of the vector `startlatent` are set to the initial value of μ .

It is customary to start independent Markov chains in parallel and `stochvol` provides facilities for that in all of its sampling procedures. The argument `n_chains` is expected to be a positive integer, it sets the number of independent chains. Additionally, arguments `parallel`, `n_cpus`, and `cl` can be used to control parallelism used by `stochvol`. To overwrite the default sequential execution strategy, `parallel` is to be set either to "snow", to employ the so-called "SNOW" clusters, or to "multicore" to use the "multicore" type computation (R Core Team 2020). Next, argument `n_cpus` should be set to the physical number of parallel processing units to be used. Finally, in case "SNOW" is applied, the sampling routines optionally accept an already running "SNOW" cluster through argument `cl`.

As mentioned earlier, the sampling algorithms for the latent states \mathbf{h} in `stochvol` rely on a Gaussian mixture approximation as in Omori *et al.* (2007) and Kastner and Frühwirth-Schnatter (2014). The approximation tends to be very good, therefore the default setting is not to correct for model misspecification. However, this correction can be enabled in all of the sampling routines through the `expert` argument as shown for `svsample()` in the following.

Lastly, `stochvol` provides three ways to economize storage during and after the execution of the sampler. Setting the `integer` argument `thinpara` to ι tells the sampler to store only every ι th draw of the vector of parameters, and supplying a value for `thinlatent` does the same for \mathbf{h} . Finally, one has the opportunity not to store the full vector \mathbf{h} but only its last value by setting `keeptime = "last"`. The default behavior is to store every draw after the burn-in phase.

5. The factorstochvol package

For more information on **factorstochvol**, please see the vignette attached to the package (Kastner and Hosszejni 2020).

6. Summary and discussion

We extended the work of Kastner (2016) to other SV models, including the univariate heavy-tailed SV, the SV model with leverage, and the multivariate factor SV model. We showcased the features that are the most important to end users in R: estimation through the sampler functions, visualization, summary, and prediction methods. Due to its more involved nature, however, we did not include the description of the C++ interface. Two functions called `update_fast_sv()` and `update_general_sv()` are exported and programmers have the possibility to access the samplers in **stochvol** directly from C++ after linking to the compiled package. For usage examples, see the implementations of **factorstochvol** or **shrink-TVP** (Knaus, Bitto-Nemling, Cadonna, and Frühwirth-Schnatter 2020).

References

- Abanto-Valle CA, Langrock R, Chen MH, Cardoso MV (2017). “Maximum likelihood estimation for stochastic volatility in mean models with heavy-tailed distributions.” *Applied Stochastic Models in Business and Industry*, **33**(4), 394–408. doi:10.1002/asmb.2246.
- Aguilar O, West M (2000). “Bayesian Dynamic Factor Models and Portfolio Allocation.” *Journal of Business and Economic Statistics*, **18**(3), 338–357. doi:10.1080/07350015.2000.10524875.
- Bollerslev T (1986). “Generalized Autoregressive Conditional Heteroskedasticity.” *Journal of Econometrics*, **31**(3), 307–327. doi:10.1016/0304-4076(86)90063-1.
- Bos CS (2012). “Relating Stochastic Volatility Estimation Methods.” In L Bauwens, C Hafner, S Laurent (eds.), *Handbook of Volatility Models and Their Applications*, pp. 147–174. John Wiley & Sons. doi:10.1002/9781118272039.ch6.
- Brooks S, Gelman A, Jones G, Meng XL (eds.) (2011). *Handbook of Markov Chain Monte Carlo*. Handbooks of Modern Statistical Methods, 1st edition. Chapman and Hall, CRC Press. ISBN 9781420079418.
- Carter CK, Kohn R (1994). “On Gibbs Sampling for State Space Models.” *Biometrika*, **81**(3), 541–553. doi:10.1093/biomet/81.3.541.
- Chib S, Nardari F, Shephard N (2006). “Analysis of High Dimensional Multivariate Stochastic Volatility Models.” *Journal of Econometrics*, **134**(2), 341–371. doi:10.1016/j.jeconom.2005.06.026.
- Creal DD (2017). “A Class of Non-Gaussian State Space Models With Exact Likelihood Inference.” *Journal of Business & Economic Statistics*, **35**(4), 585–597. doi:10.1080/07350015.2015.1092977.

- Delatola EI, Griffin JE (2011). “Bayesian Nonparametric Modeling of the Return Distribution with Stochastic Volatility.” *Bayesian Analysis*, **6**, 901–926. doi:10.1214/11-BA632.
- Eddelbuettel D, François R (2011). “**Rcpp**: Seamless R and C++ Integration.” *Journal of Statistical Software*, **40**(8), 1–18. doi:10.18637/jss.v040.i08.
- Eddelbuettel D, Sanderson C (2014). “**RcppArmadillo**: Accelerating R with High-Performance C++ Linear Algebra.” *Computational Statistics and Data Analysis*, **71**, 1054–1063. doi:10.1016/j.csda.2013.02.005.
- Engle RF (1982). “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation.” *Econometrica*, **50**(4), 987–1007. doi:10.2307/1912773.
- Frühwirth-Schnatter S (1994). “Data Augmentation and Dynamic Linear Models.” *Journal of Time Series Analysis*, **15**(2), 183–202. doi:10.1111/j.1467-9892.1994.tb00184.x.
- Frühwirth-Schnatter S, Wagner H (2010). “Stochastic Model Specification Search for Gaussian and Partial Non-Gaussian State Space Models.” *Journal of Econometrics*, **154**(1), 85–100. doi:10.1016/j.jeconom.2009.07.003.
- Geweke J (1993). “Bayesian Treatment of the Independent Student-*t* Linear Model.” *Journal of Applied Econometrics*, **8**(S1), S19–S40. doi:10.1002/jae.3950080504.
- Geweke J (2004). “Getting It Right: Joint Distribution Tests of Posterior Simulators.” *Journal of the American Statistical Association*, **99**, 799–804.
- Geweke J, Amisano G (2010). “Comparing and Evaluating Bayesian Predictive Distributions of Asset Returns.” *International Journal of Forecasting*, **26**(2), 216–230. doi:10.1016/j.ijforecast.2009.10.007.
- Ghalanos A (2020). **rugarch**: *Univariate GARCH Models*. R package version 1.4-4, URL <https://CRAN.R-project.org/package=rugarch>.
- Ghysels E, Harvey AC, Renault E (1996). “Stochastic Volatility.” In GS Maddala, CR Rao (eds.), *Handbook of Statistics*, volume 14, chapter 5, pp. 119–191. Elsevier. doi:10.1016/s0169-7161(96)14007-4.
- Griffin JE, Brown PJ (2010). “Inference with Normal-Gamma Prior Distributions in Regression Problems.” *Bayesian Analysis*, **5**(1), 171–188. doi:10.1214/10-BA507.
- Han Y (2006). “Asset Allocation with a High Dimensional Latent Factor Stochastic Volatility Model.” *Review of Financial Studies*, **19**(1), 237–271. doi:10.1093/rfs/hhj002.
- Harvey AC, Ruiz E, Shephard N (1994). “Multivariate Stochastic Variance Models.” *The Review of Economic Studies*, **61**(2), 247–264. doi:10.2307/2297980.
- Harvey AC, Shephard N (1996). “Estimation of an Asymmetric Stochastic Volatility Model for Asset Returns.” *Journal of Business & Economic Statistics*, **14**(4), 429–434. doi:10.1080/07350015.1996.10524672.
- Higham NJ (1990). “Analysis of the Cholesky Decomposition of a Semi-Definite Matrix.” *Technical report*, Manchester Institute for Mathematical Sciences. MIMS EPrint 2008.56, URL <http://eprints.maths.manchester.ac.uk/1193/>.

- Hosszejni D, Kastner G (2019a). “Approaches Toward the Bayesian Estimation of the Stochastic Volatility Model with Leverage.” In R Argiento, D Durante, S Wade (eds.), *Bayesian Statistics: New Challenges and New Generations*, Springer Proceedings in Mathematics & Statistics, pp. 75–83. Springer-Verlag. doi:10.1007/978-3-030-30611-3_8.
- Hosszejni D, Kastner G (2019b). “Modeling Univariate and Multivariate Stochastic Volatility in R with **stochvol** and **factorstochvol**.” URL <https://arxiv.org/abs/1906.12123>.
- Hosszejni D, Kastner G (2020). **stochvol**: *Efficient Bayesian Inference for Stochastic Volatility (SV) Models*. R package version 3.0.3, URL <http://CRAN.R-project.org/package=stochvol>.
- Huber ML (2015). *Perfect Simulation*. Monographs on Statistics and Applied Probability, 1st edition. Chapman and Hall, CRC Press. ISBN 9781482232448.
- ISO/IEC (2017). *International Standard ISO/IEC 14882:2017(E) – Programming Language C++*. International Organization for Standardization (ISO). URL <https://isocpp.org/std/the-standard>.
- Jacquier E, Polson NG, Rossi PE (1994). “Bayesian Analysis of Stochastic Volatility Models.” *Journal of Business & Economic Statistics*, **20**(1), 69–87. doi:10.1080/07350015.1994.10524553.
- Jacquier E, Polson NG, Rossi PE (2004). “Bayesian Analysis of Stochastic Volatility Models with Fat-Tails and Correlated Errors.” *Journal of Econometrics*, **122**(1), 185–212. doi:10.1016/j.jeconom.2003.09.001.
- Jensen MJ, Maheu JM (2010). “Bayesian Semiparametric Stochastic Volatility Modeling.” *Journal of Econometrics*, **157**(2), 306–316. doi:10.1016/j.jeconom.2010.01.014.
- Jensen MJ, Maheu JM (2014). “Estimating a Semiparametric Asymmetric Stochastic Volatility Model with a Dirichlet Process Mixture.” *Journal of Econometrics*, **178**(3), 523–538. doi:10.1016/j.jeconom.2013.08.018.
- Kastner G (2015). “Heavy-Tailed Innovations in the R Package **stochvol**.” *Technical report*, WU Vienna University of Economics and Business. URL <https://epub.wu.ac.at/4918>.
- Kastner G (2016). “Dealing with Stochastic Volatility in Time Series Using the R Package **stochvol**.” *Journal of Statistical Software*, **69**(5), 1–30. doi:10.18637/jss.v069.i05.
- Kastner G (2019). “Sparse Bayesian Time-Varying Covariance Estimation in Many Dimensions.” *Journal of Econometrics*, **210**(1), 98–115. doi:10.1016/j.jeconom.2018.11.007.
- Kastner G, Frühwirth-Schnatter S (2014). “Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Estimation of Stochastic Volatility Models.” *Computational Statistics and Data Analysis*, **76**, 408–423. doi:10.1016/j.csda.2013.01.002.
- Kastner G, Frühwirth-Schnatter S, Lopes HF (2017). “Efficient Bayesian Inference for Multivariate Factor Stochastic Volatility Models.” *Journal of Computational and Graphical Statistics*, **26**(4), 905–917. doi:10.1080/10618600.2017.1322091.

- Kastner G, Hosszejni D (2020). **factorstochvol**: *Bayesian Estimation of (Sparse) Latent Factor Stochastic Volatility Models*. R package version 0.10.1, URL <https://cran.r-project.org/package=factorstochvol>.
- Kim S, Shephard N, Chib S (1998). “Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models.” *The Review of Economic Studies*, **65**(3), 361–393. doi: [10.1111/1467-937x.00050](https://doi.org/10.1111/1467-937x.00050).
- Knaus P, Bitto-Nemling A, Cadonna A, Frühwirth-Schnatter S (2020). **shrinkTVP**: *Efficient Bayesian Inference for Time-Varying Parameter Models with Shrinkage*. R package version 2.0.1, URL <https://CRAN.R-project.org/package=shrinkTVP>.
- Markowitz H (1952). “Portfolio Selection.” *The Journal of Finance*, **7**(1), 77–91. doi: [10.1111/j.1540-6261.1952.tb01525.x](https://doi.org/10.1111/j.1540-6261.1952.tb01525.x).
- McCausland WJ, Miller S, Pelletier D (2011). “Simulation Smoothing for State-Space Models: A Computational Efficiency Analysis.” *Computational Statistics and Data Analysis*, **55**(1), 199–212. doi:[10.1016/j.csda.2010.07.009](https://doi.org/10.1016/j.csda.2010.07.009).
- McElreath R (2015). *Statistical Rethinking: A Bayesian Course with Examples in R and Stan*. Texts in Statistical Science, 1st edition. Chapman and Hall, CRC Press. ISBN 9781482253443.
- Nakajima J (2012). “Bayesian Analysis of Generalized Autoregressive Conditional Heteroskedasticity and Stochastic Volatility: Modeling Leverage, Jumps and Heavy-Tails for Financial Time Series.” *Japanese Economic Review*, **63**(1), 81–103. doi:[10.1111/j.1468-5876.2011.00537.x](https://doi.org/10.1111/j.1468-5876.2011.00537.x).
- Nakajima J, Omori Y (2009). “Leverage, Heavy-Tails and Correlated Jumps in Stochastic Volatility Models.” *Computational Statistics and Data Analysis*, **53**(6), 2335–2353. doi: [10.1016/j.csda.2008.03.015](https://doi.org/10.1016/j.csda.2008.03.015).
- Nakajima J, Omori Y (2012). “Stochastic Volatility Model with Leverage and Asymmetrically Heavy-Tailed Error Using GH Skew Student’s t Distribution.” *Computational Statistics & Data Analysis*, **56**(11), 3690–3704. doi:[10.1016/j.csda.2010.07.012](https://doi.org/10.1016/j.csda.2010.07.012).
- Omori Y, Chib S, Shephard N, Nakajima J (2007). “Stochastic Volatility with Leverage: Fast and Efficient Likelihood Inference.” *Journal of Econometrics*, **140**(2), 425–449. doi: [10.1016/j.jeconom.2006.07.008](https://doi.org/10.1016/j.jeconom.2006.07.008).
- Park T, Casella G (2008). “The Bayesian Lasso.” *Journal of the American Statistical Association*, **103**(452), 681–686. doi:[10.1198/016214508000000337](https://doi.org/10.1198/016214508000000337).
- Pitt MK, Shephard N (1999). “Time-Varying Covariances: A Factor Stochastic Volatility Approach.” In JM Bernardo, JO Berger, AP Dawid, AFM Smith (eds.), *Bayesian Statistics 6 – Proceedings of the Sixth Valencia International Meeting*, pp. 547–570. Oxford University Press.
- Plummer M, Best N, Cowles K, Vines K (2006). “CODA: Convergence Diagnosis and Output Analysis for MCMC.” *R News*, **6**(1), 7–11. URL https://www.r-project.org/doc/Rnews/Rnews_2006-1.pdf.

- R Core Team (2020). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Roberts GO, Rosenthal J (2007). “Coupling and Ergodicity of Adaptive Markov Chain Monte Carlo Algorithms.” *Journal of Applied Probability*, **44**(2), 458–475. doi:10.1239/jap/1183667414.
- Rue H (2001). “Fast Sampling of Gaussian Markov Random Fields.” *Journal of the Royal Statistical Society B*, **63**(2), 325–338. doi:10.1111/1467-9868.00288.
- Sanderson C, Curtin R (2016). “**Armadillo**: A Template-Based C++ Library for Linear Algebra.” *Journal of Open Source Software*, **1**, 26–18. doi:10.21105/joss.00026.
- Silva RS, Lopes HF, Migon HS (2006). “The Extended Generalized Inverse Gaussian Distribution for Log-Linear and Stochastic Volatility Models.” *Brazilian Journal of Probability and Statistics*, **20**(1), 67–91. URL <https://www.jstor.org/stable/43601074>.
- Sisson SA, Fan Y, Beaumont MA (eds.) (2018). *Handbook of Approximate Bayesian Computation*. Handbooks of Modern Statistical Methods, 1st edition. Chapman and Hall, CRC Press. ISBN 9781439881507.
- Taylor SJ (1982). “Financial Returns Modeled by the Product of Two Stochastic Processes: A Study of Daily Sugar Prices 1961–75.” In OD Anderson (ed.), *Time Series Analysis, Theory and Practice*, pp. 203–226. North-Holland, Amsterdam.
- Wahl JC (2020). *stochvolTMB: Likelihood Estimation of Stochastic Volatility Models*. R package version 0.1.2, URL <https://CRAN.R-project.org/package=stochvolTMB>.
- Wickham H, Hester J, Chang W (2020). *devtools: Tools to Make Developing R Packages Easier*. R package version 2.3.2, URL <https://CRAN.R-project.org/package=devtools>.
- Yu Y, Meng XL (2011). “To Center or not to Center: That is not the Question—An Ancillarity-Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Efficiency.” *Journal of Computational and Graphical Statistics*, **20**(3), 531–570. doi:10.1198/jcgs.2011.203main.

Affiliation:

Darjus Hosszejni
 Institute for Statistics and Mathematics
 Department of Finance, Accounting and Statistics
 WU Vienna University of Economics and Business
 Welthandelsplatz 1 / Building D4 / Level 4
 1020 Vienna, Austria
 E-mail: Darjus.Hosszejni@wu.ac.at
 URL: <http://statmath.wu.ac.at/~hosszejni/>

Gregor Kastner
 Department of Statistics

University of Klagenfurt
Universitätsstraße 65-67
9020 Klagenfurt, Austria
E-mail: Gregor.Kastner@aau.at
URL: <http://statmath.wu.ac.at/~kastner/>