



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/94382/>

Version: Accepted Version

Article:

Langrock, R., Michelot, T., Sohn, A. et al. (2014) Semiparametric stochastic volatility modelling using penalized splines. *Computational Statistics*, 30 (2). pp. 517-537. ISSN: 0943-4062

<https://doi.org/10.1007/s00180-014-0547-5>

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Semiparametric stochastic volatility modelling using penalized splines

Roland Langrock¹, Théo Michelot², Alexander Sohn³, and Thomas Kneib³

¹University of St Andrews, UK.

²INSA de Rouen, France

³Georg August University of Göttingen, Germany

Abstract

Stochastic volatility (SV) models mimic many of the stylized facts attributed to time series of asset returns, while maintaining conceptual simplicity. The commonly made assumption of conditionally normally distributed or Student-t-distributed returns, given the volatility, has however been questioned. In this manuscript, we introduce a novel maximum penalized likelihood approach for estimating the conditional distribution in an SV model in a nonparametric way, thus avoiding any potentially critical assumptions on the shape. The considered framework exploits the strengths both of the powerful hidden Markov model machinery and of penalized B-splines, and constitutes a powerful and flexible alternative to recently developed Bayesian approaches to semiparametric SV modelling. We demonstrate the feasibility of the approach in a simulation study before outlining its potential in applications to three series of returns on stocks and one series of stock index returns.

Keywords: B-splines; cross-validation; forward algorithm; hidden Markov model; numerical integration; penalized likelihood

1 Introduction

Stochastic volatility (SV) models are immensely popular tools for the analysis of financial time series. This subclass of state-space models (SSMs) constitutes one of the two most widely

used approaches for modelling stock market volatility, the other one being ARCH/GARCH type models. The structure of the standard discrete-time SV model, labeled SV_0 in the following, for log-returns y_1, \dots, y_T on an asset, is as follows:

$$y_t = \varepsilon_t^{(0)} \beta \exp(g_t/2), \quad g_t = \phi g_{t-1} + \sigma \eta_t, \quad (1)$$

where $\beta, \sigma > 0$, and where $\{\varepsilon_t^{(0)}\}$ and $\{\eta_t\}$ are independent sequences of independent standard normal random variables (see, e.g., Shephard, 1996). Stationarity is obtained for $|\phi| < 1$. The unobserved sequence $\{g_t\}$, commonly referred to as the log-volatility process, represents the time-varying “nervousness” of the market. This model captures several of the stylized facts attributed to asset returns, including positive autocorrelation of squared returns (indicating a volatility that slowly varies over time and hence volatility clustering), zero autocorrelation of the unsquared returns, and a kurtosis in excess of 3. However, the basic model tends to underestimate the probability of relatively extreme returns, such that it is often more adequate to consider a Student-t distribution with ν degrees of freedom for $\varepsilon_t^{(0)}$ (Chib et al., 2002); we label this second model SV_t .

In the existing literature, several different model formulations have been considered that extend the flexibility of the log-volatility process, $\{g_t\}$ (e.g., Gallant et al., 1997; Abraham et al., 2006; Langrock et al., 2012). Yet Durham (2006) found “no evidence that even simple single factor models are unable to capture the dynamics of the volatility process” (p.276). Instead, he considers the shape of the conditional distribution in SV models — i.e., of the conditional distribution of y_t , given g_t — to be “the more critical problem” (p.304). In addition to heavy tails, which are accounted for in the SV_t formulation, evidence of asymmetries has been found (see, e.g., Gallant et al., 1997; Harvey and Siddique, 2000; Jondeau and Rockinger, 2003; Durham, 2006). To get the shape right, and in particular to accurately estimate the tails of the conditional distribution, is of high importance for example in risk management. While parametric models can be constructed that enable the inclusion of heavy tails and skewness, nonparametric approaches have the considerable advantage that no restriction to a particular class of distributions is made a priori.

Much recent work in this direction has been conducted in the Bayesian context, where the normal distribution can be used as a building block to formulate more complex models that still utilize the benefits of the normal formulation for constructing convenient update schemes

in a Markov chain Monte Carlo simulation. Abanto-Valle et al. (2010) consider scale mixtures of normals for the conditional distribution, where the variance of the normal distribution is supplemented with suitable prior specifications that yield a larger class of potential marginal distributions after integrating out the mixing distribution for the variance. Nonparametric specifications relying on an infinite-dimensional mixture of normals, generated by a Dirichlet process mixture prior, have been developed in Jensen and Maheu (2010) and Delatola and Griffin (2011). While Jensen and Maheu (2010) directly tackle the conditional distribution of the returns, Delatola and Griffin (2011) employ a different representation where the logarithm of the squared noise in the return process is considered for the analysis. Delatola and Griffin (2013) extend the latter approach by including a leverage effect, allowing for the potential correlation of the two error terms, in the return and in the log-volatility process, respectively.

In this manuscript, we develop a novel frequentist approach for nonparametrically estimating the conditional distribution in an SV model. The proposed maximum penalized likelihood approach exploits the strengths both of likelihood-based hidden Markov model (HMM) machinery and of penalized B-splines (i.e., P-splines). The former is employed to deal with a well-known difficulty with SV models, which is that their likelihood is given by a high-order multiple integral that is analytically intractable. It has however been shown that methods available for HMMs — which have the same dependence structure as SV models and constitute another subclass of SSMs, with finite state space — can be applied in order to perform a fast and accurate numerical integration of the SV model likelihood. More specifically, such a numerical integration corresponds to a fine discretization of the support of the log-volatility process. The associated transformation of the continuous support of $\{g_t\}$ to a finite support renders the powerful HMM forward algorithm applicable, making it feasible to evaluate an arbitrarily accurate approximation to the SV model likelihood (Fridman and Harris, 1998; Bartolucci and De Luca, 2001, 2003; Langrock et al., 2012). We extend this likelihood-based approach to allow for a nonparametric estimation of the conditional distribution, by representing the density of this distribution as a linear combination of a large number of standardized B-spline basis functions, including a roughness penalty in the likelihood in order to arrive at an appropriate balance between goodness of fit and smoothness for the fitted density. Since we still model the log-volatility process in a parametric way, we use the label SV_{sp} to refer to the resulting *semiparametric* SV model with nonparametrically modelled conditional distribution.

The paper is structured as follows. In Section 2, we begin by describing the HMM-based likelihood evaluation, then introducing the B-spline-based representation of the conditional distribution and discussing associated inferential issues. The performance of the suggested approach is investigated in a simulation study in Section 3. In Section 4, we apply the approach to real data related to three stocks and one stock index, comparing the predictive performance of our model to popular parametric counterparts.

2 Semiparametric SV modelling

2.1 SV model likelihood

We consider a model SV_{sp} of the form

$$y_t = \varepsilon_t \exp(g_t/2), \quad g_t = \phi g_{t-1} + \sigma \eta_t, \quad (2)$$

with the η_t iid standard normal, but where, in contrast to the models SV_0 and SV_t , we do not make any assumptions on the distributional form of the random variables ε_t . However, we do assume these variables to be iid, and to be independent of $\{\eta_t\}$. Our aim is to nonparametrically estimate the probability density function (pdf) f_ε of the variables ε_t . Compared to model SV_0 , as given in (1), we have omitted the parameter β in (2), since otherwise the semiparametric model would not be identifiable; in the SV_{sp} model the effect of β will be absorbed within f_ε . Before we introduce our strategy for estimating f_ε in a nonparametric way (alongside the other model parameters), we will derive a tractable likelihood function for general f_ε , including the nonparametric case, but also those of a normal distribution and of a Student-t distribution. To formulate the likelihood, we will require the conditional pdfs of the random variables y_t , given g_t ($t = 1, \dots, T$). We denote these conditional pdfs by $f(y_t|g_t)$, for $t = 1, \dots, T$. These pdfs are simple transformations of the density f_ε :

$$f(y_t|g_t) = \exp(-g_t/2) f_\varepsilon(y_t \exp(-g_t/2)).$$

For any f_ε , the likelihood of the model defined by (2) can then be derived as

$$\begin{aligned}
\mathcal{L} &= \int \dots \int f(y_1, \dots, y_T, g_1, \dots, g_T) dg_T \dots dg_1 \\
&= \int \dots \int f(y_1, \dots, y_T | g_1, \dots, g_T) f(g_1, \dots, g_T) dg_T \dots dg_1 \\
&= \int \dots \int f(g_1) f(y_1 | g_1) \prod_{t=2}^T f(g_t | g_{t-1}) f(y_t | g_t) dg_T \dots dg_1 \\
&= \int \dots \int f(g_1) \exp(-g_1/2) f_\varepsilon(y_1 \exp(-g_1/2)) \\
&\quad \times \prod_{t=2}^T f(g_t | g_{t-1}) \exp(-g_t/2) f_\varepsilon(y_t \exp(-g_t/2)) dg_T \dots dg_1. \tag{3}
\end{aligned}$$

In the second last step, we exploited the dependence structure that is characteristic of SV models, HMMs and general SSMs. Hence, the likelihood is a high-order multiple integral that cannot be evaluated directly. Via numerical integration, using a simple rectangular rule based on m equidistant intervals, $B_i = (b_{i-1}, b_i)$, $i = 1, \dots, m$, with midpoints b_i^* and of length b , the likelihood can be approximated as follows:

$$\begin{aligned}
\mathcal{L} &\approx b^T \sum_{i_1=1}^m \dots \sum_{i_T=1}^m f(b_{i_1}^*) \exp(-b_{i_1}^*/2) f_\varepsilon(y_1 \exp(-b_{i_1}^*/2)) \\
&\quad \times \prod_{t=2}^T f(b_{i_t}^* | b_{i_{t-1}}^*) \exp(-b_{i_t}^*/2) f_\varepsilon(y_t \exp(-b_{i_t}^*/2)) = \mathcal{L}_{\text{approx}}. \tag{4}
\end{aligned}$$

This approximation can be made arbitrarily accurate by increasing m , and in fact virtually exact for m around 100, provided that the interval (b_0, b_m) covers the essential range of the log-volatility process (more details below). In the given form, the approximate likelihood (4) is usually computationally intractable, since it involves m^T function evaluations. However, an efficient recursive scheme can be used to evaluate (4). To see this, note that the numerical integration essentially corresponds to a discretization of the state space, i.e., of the support of the log-volatility process $\{g_t\}$. Therefore, the approximate likelihood given in (4) can be evaluated using the well-developed and powerful machinery of the subclass of SSMs given by HMMs, which are SSMs with a finite state space (cf. Langrock, 2011; Langrock et al., 2012). We sketch the relevant HMM methodology in the appendix to this manuscript. In particular, in the appendix we highlight a key property of HMMs, which is that the likelihood can be evaluated efficiently using the so-called forward algorithm (Zucchini and MacDonald, 2009). For an HMM, applying the forward algorithm results in a matrix product expression for the

likelihood, and this is exactly what we obtain also in the present context:

$$\mathcal{L}_{\text{approx}} = \boldsymbol{\delta} \mathbf{P}(y_1) \boldsymbol{\Omega} \mathbf{P}(y_2) \boldsymbol{\Omega} \mathbf{P}(y_3) \cdots \boldsymbol{\Omega} \mathbf{P}(y_{T-1}) \boldsymbol{\Omega} \mathbf{P}(y_T) \mathbf{1}. \quad (5)$$

Here, the $m \times m$ -matrix $\boldsymbol{\Omega} = (\omega_{ij})$ is the analogue to the transition probability matrix in case of an HMM (see the appendix), defined as $\omega_{ij} = f(b_j^* | b_i^*) \cdot b$. Furthermore, the vector $\boldsymbol{\delta}$ is the analogue to the Markov chain initial distribution in case of an HMM, here defined such that δ_i , $i = 1, \dots, m$, is the density of the normal distribution with mean zero and standard deviation $\sigma / \sqrt{1 - \phi^2}$ — the stationary distribution of the autoregressive process used to model the log-volatility — evaluated at b_i^* and multiplied by b . Finally, $\mathbf{P}(y_t)$ is an $m \times m$ diagonal matrix with i th diagonal entry $\exp(-b_i^*/2) f_\varepsilon(y_t \exp(-b_i^*/2))$, hence the analogue to the matrix comprising the state-dependent probabilities in case of an HMM. Using the matrix product expression given in (5), the computational effort required to evaluate the approximate likelihood is linear in the number of observations, T , and quadratic in the number of intervals used in the discretization, m . In practice, this means that the likelihood of an SV model can typically be calculated in a fraction of a second, even for T in the thousands and say $m = 150$, a value which renders the approximation virtually exact. Furthermore, $\mathcal{L}_{\text{approx}} \rightarrow \mathcal{L}$ as $b_m, m \rightarrow \infty$ and $b_0 \rightarrow -\infty$.

2.2 Nonparametric modelling

We now turn to the nonparametric modelling of the distribution of ε_t . Following Schellhase and Kauermann (2012), we suggest to estimate the pdf of ε_t by considering finite linear combinations of a large number of basis functions:

$$\hat{f}_\varepsilon(x) = \sum_{k=-K}^K a_k \psi_k(x). \quad (6)$$

Here the basis functions ψ_{-K}, \dots, ψ_K are known and fixed pdfs. Clearly, $\hat{f}_\varepsilon(x)$ then is a pdf if $\sum_{k=-K}^K a_k = 1$ and $a_j \geq 0$ for all $j = -K, \dots, K$. To enforce these constraints, the coefficients to be estimated, a_{-K}, \dots, a_K , are transformed using the multinomial logit link function

$$a_k = \frac{\exp(\beta_k)}{\sum_{j=-K}^K \exp(\beta_j)}, \quad (7)$$

where we set $\beta_0 = 0$ for identifiability. In principle, any set of densities ψ_{-K}, \dots, ψ_K can be used to approximate $f_\varepsilon(x)$ as in (6). We follow Schellhase and Kauermann (2012) and use B-splines, in ascending order in the basis used in (6), and standardized such that they integrate to 1. For more details on B-splines, see, e.g., de Boor (1978) and Eilers and Marx (1996). Since each B-spline basis function is associated with a separate parameter, this model formulation in fact leads to a finite-dimensional parameter space. However, the dimensionality is high and the separate parameters are no longer of interest or interpretable — the model is *overparameterized*. We therefore refer to our estimation approach as nonparametric despite the fact that it does rely on a parametric specification with a large number of parameters. This is in line with the standard terminology used in the literature, where (penalized) spline approaches are subsumed under nonparametric approaches (see, e.g., Ruppert et al., 2003).

The approximate likelihood of the resulting SV_{sp} model is given by (5), plugging in \hat{f}_ε for f_ε in the matrices $\mathbf{P}(y_t)$, $t = 1, \dots, T$. Following Eilers and Marx (1996), we modify the (approximate) log-likelihood by including a penalty on (q -th order) differences between coefficients associated with adjacent B-splines, yielding the penalized log-likelihood

$$l_p = \log(\mathcal{L}_{\text{approx}}) - \frac{\lambda}{2} \sum_{k=-K+q}^K (\Delta^q a_k)^2, \quad (8)$$

with a_k parameterized as in (7) and smoothing parameter $\lambda \geq 0$. The penalty term involves the difference operator Δ , where $\Delta a_k = a_k - a_{k-1}$ and $\Delta^q a_k = \Delta(\Delta^{q-1} a_k)$. This results in a penalization of roughness of the estimator, with λ controlling how much emphasis is put on goodness of fit and on smoothness, respectively. In particular, unpenalized estimates are obtained for $\lambda = 0$. For $\lambda \rightarrow \infty$ the penalty will dominate the likelihood, resulting in a sequence of weights a_k that follow a polynomial of order $q - 1$ in k . The difference order therefore also affects the smoothness of the estimates indirectly (and to a much smaller extent than the degree of the spline basis). We will use $q = 2$ in the remainder since this provides an approximation to the integrated squared second derivative penalty that is popular in the context of smoothing splines.

Including the penalty term in the likelihood avoids the problem of selecting an optimal number of basis elements, since the penalty effectively reduces the number of free basis parameters and yields an adaptive fit to the data, provided the smoothing parameter is chosen in a data-driven way. The number of basis elements needs to be large enough to

give sufficient flexibility for reflecting the structure of the conditional distribution f_ε , but once this threshold is passed, increasing the number of basis elements further does no longer change the fit to the data much due to the impact of the penalty. For moderately smooth regression functions, Ruppert (2002) recommends to use a default of about 35 (or 40) basis functions. To capture the pdf of ε_t in an SV model, we expect such a choice to easily provide sufficient flexibility, and hence have chosen K accordingly in our analyses (see below). To select the smoothing parameter in a data-driven way, we will consider cross-validation (see Section 2.3.2).

In preliminary simulation experiments, we found that with an equidistant spacing of the knots our approach tended to produce estimated densities that were overly smooth around the peak of the true distribution and too wiggly in the tails. This is related to the fact that the basis coefficients systematically decay towards the tails of the estimated distribution which would require an adaptive amount of smoothing instead of a global smoothing parameter. As a simple yet effective strategy to achieve such adaptiveness, we consider increasingly wider distances between the B-spline basis densities towards the tails instead of the common equidistant specification. Since we still rely on the unweighted difference penalty in (8), this effectively increases the penalty for the tails of the distribution.

2.3 Inference

2.3.1 Parameter estimation

The use of the forward algorithm leads to a very fast evaluation of the penalized log-likelihood given in (8). A numerical maximization of the penalized log-likelihood is therefore feasible in typical cases, even for high m and hence very close approximations to the likelihood in (3); some computing times are given in Section 4. Since the first part of expression (8) is susceptible to numerical overflow, it is required to compute its logarithm, which involves a minor difficulty since we are dealing with a matrix product. However, techniques to address this issue are standard: Zucchini and MacDonald (2009) describe a straightforward scaling strategy for calculating the logarithm of an HMM-type matrix product likelihood (see their Chapter 3).

In practice, one also has to select the value of m , the number of intervals used in the discretization of the log-volatility process, and the range of possible g_t -values considered in the

numerical integration. In our experience, estimates usually stabilize for values of m around 50 (cf. Langrock et al., 2012; Langrock and King, 2013). The minimum and maximum values for g_t have to be chosen sufficiently large to cover the essential domain of the log-volatility process, but not too large, in order to maintain sufficient fineness of the grid. More guidance on this issue is provided in Langrock et al. (2012). Another technical issue in the numerical maximization is that of local maxima: it may sometimes happen that the numerical search fails to find the MLE, and return a local maxima instead. The best way to address this issue seems to be to use a number of different sets of initial values in order to find and verify the global maximum. Uncertainty quantification, for both the parameters of the underlying log-volatility process and the density of ε_t , can be conducted using a parametric bootstrap, but the computational burden is relatively high.

2.3.2 Choice of the smoothing parameter

Cross-validation techniques can be used to choose the smoothing parameter. For a given series of log-returns, we suggest to generate C random partitions such that in each partition a suitable percentage of the observations form the calibration sample, while the remaining observations constitute the validation sample. For each of the C partitions and any given λ , the model is then calibrated by estimating the parameters using only the calibration sample (treating the data points from the validation sample as missing data). Subsequently, proper scoring rules (Gneiting and Raftery, 2007) can be used on the validation sample to assess the calibrated model for the given λ . For computational convenience, we consider the log-likelihood of the validation sample (now treating the data points from the calibration sample as missing data), under the model fitted in the calibration stage, as the score of interest. From some pre-specified set of possible smoothing parameters, e.g., $\{2^n | n = r, r + 1, \dots, s\}$ where r and s are integers, we then select the λ with the highest mean score, over all C cross-validation samples. The number of samples C needs to be high enough to give meaningful scores, but must not be too high to allow for the approach to be computationally feasible. In our real data analyses (see below), the consideration of $C = 40$ samples led to stable results yet was computationally feasible.

2.3.3 Model checking

We have already seen that the use of the HMM forward algorithm provides an efficient and convenient way to evaluate the SV model likelihood. HMM machinery can also be exploited in order to check the goodness of fit of a given model. Following Zucchini and MacDonald (2009), we consider one-step-ahead forecast pseudo-residuals, which are given by

$$r_t = \Phi^{-1}\left(F(y_t \mid y_{t-1}, y_{t-2}, \dots, y_1)\right).$$

Here Φ denotes the cumulative distribution function of the standard normal distribution, and $F(y_t \mid y_{t-1}, y_{t-2}, \dots, y_1)$ is the cumulative distribution function of y_t given all observations up to time $t - 1$. Using the HMM-type approximation, this can be written as

$$\begin{aligned} F(y_t \mid y_{t-1}, y_{t-2}, \dots, y_1) &\approx \sum_{i=1}^m \zeta_i F(y_t \mid g_t = b_i^*) \\ &= \sum_{i=1}^m \zeta_i \int_{-\infty}^{y_t} \exp(-b_i^*/2) f_\varepsilon(x \exp(-b_i^*/2)) dx, \end{aligned} \quad (9)$$

where ζ_i is the i th entry of the vector $\tilde{\alpha}_{t-1}\mathbf{\Omega}/(\tilde{\alpha}_{t-1}\mathbf{1}')$, which is defined as

$$\tilde{\alpha}_{t-1} = \delta\mathbf{P}(y_1)\mathbf{\Omega}\mathbf{P}(y_2)\mathbf{\Omega}\cdots\mathbf{\Omega}\mathbf{P}(y_{t-1}),$$

$t = 2, \dots, T$, with δ , $\mathbf{P}(y_k)$ and $\mathbf{\Omega}$ defined as above. These $\tilde{\alpha}_t$'s constitute the SV model analogue to the HMM forward probabilities; see the appendix for more details on the latter. The representation given in (9) is only approximate due to the discretization of the log-volatility process, but as for the likelihood the accuracy also of this approximation can be made arbitrarily accurate by increasing m . In the context of SV models such residuals were first used by Kim et al. (1998). If the fitted model is correct, then the pseudo-residuals are distributed standard normal. Thus, forecast pseudo-residuals can be used to identify extreme values, and the general suitability of the model can be checked by using, e.g., quantile-quantile-plots or formal tests for normality.

2.3.4 Decoding

Again building on existing HMM machinery, estimates of the underlying log-volatility can easily be obtained using the Viterbi algorithm, which is an efficient dynamic programming

algorithm for computing the most likely Markov chain state sequence to have given rise to observations stemming from an HMM (see Langrock et al., 2012, or Chapter 5 of Zucchini and MacDonald, 2009).

3 Simulation experiments

To generate artificial data which adhere to many of the stylized facts discussed above, we use an SV model as in (2), with $\phi = 0.98$, $\sigma = 0.1$ and ε_t specified as $\varepsilon_t = 0.02(\zeta_{1,t} - 1)^{\alpha_t}(\zeta_{2,t} + 1)^{1-\alpha_t} + 0.006$, where $\zeta_{1,t}$ and $\zeta_{2,t}$ are mutually independent iid sequences of Student-t random variables with 6 and 8 degrees of freedom, respectively, and α_t are iid Bernoulli variables, each taking on the value 1 with probability 0.35. This specification results in a skewed and leptokurtic distribution (skewness ≈ -0.22 ; kurtosis ≈ 3.58) with zero mean. For an illustration of the shape of this distribution, see Figure 1.

For this model, we conducted 200 simulation runs, with $T = 4000$ observations being generated in each run. In each run, the final 1000 observations of the generated series were used only to assess the predictive capacity of various models, which were previously fitted to the first 3000 observations. To make a fairly extensive simulation study feasible, we did not conduct a cross-validation for the smoothing parameter λ within each simulation run. Instead, we ran cross-validations only in 10 preliminary simulation runs, trying the values 256, 512, 1024, 2048, 4096 and 8192, and then fixed λ for the main 200 simulation runs at the value which was selected most often by cross-validation in the preliminary runs (namely $\lambda = 1024$). This procedure resulted in a good performance (see below), but in fact the results could potentially be further improved by conducting a cross-validation within each simulation run. We set $K = 15$, resulting in 31 B-spline basis densities that were used in the estimation. To obtain a benchmark for the semiparametric model SV_{sp} , we further fitted the basic models SV_0 and SV_t to each generated series, also using the HMM-based discretization approach described in Section 2.1.

For the SV_{sp} model, the sample mean estimates of the parameters ϕ and σ were obtained as 0.978 (sample standard error: 0.007) and 0.103 (0.017), respectively. For the SV_0/SV_t model, the sample mean estimates of the parameters ϕ and σ were obtained as 0.972/0.977 (sample standard errors: 0.010/0.007) and 0.117/0.102 (0.023/0.017), respectively. The SV_0 model underestimated the persistence parameter and overestimated the variance of the log-

volatility process. The latter is due to the model’s inability to capture the slight excess kurtosis of the true conditional distribution (see further comments on this issue in Section 4.2). In contrast, both the SV_{sp} model and the SV_t model yielded approximately unbiased estimates of the parameters related to the log-volatility process. Concerning the conditional process, y_t , Figure 1 displays the true pdf of ε_t and the corresponding pdfs that were estimated using the nonparametric approach. From the graphic we can see that all 200 fits seem fairly reasonable.




Figure 1: True density of ε_t considered in the simulation experiments (red line) and its 200 estimates obtained using the nonparametric approach (grey lines).

In the given scenario, we also assessed the predictive capacity of the three different modelling approaches, represented by the models SV_0 , SV_t and SV_{sp} . We did this by calculating, in each simulation run and under each of the three types of models fitted to the first 3000 observations, the log-likelihood score for the final 1000 observations, denoted by $llk_i(SV_0)$, $llk_i(SV_t)$ and $llk_i(SV_{sp})$, with i indicating the simulation run. These scores were compared to the corresponding score obtained when using the true model, i.e., the one that was actually used to generate the artificial data; we denote this score by $llk_i(SV_{true})$. In this simulation experiment, with skewed and leptokurtic conditional distribution, the averaged differences between the scores obtained for SV_0 , SV_t and SV_{sp} on the one hand, and the score obtained

for the true model on the other hand, were obtained as

$$\begin{aligned}\frac{1}{200} \sum_{i=1}^{200} (\text{llk}_i(\text{SV}_0) - \text{llk}_i(\text{SV}_{\text{true}})) &= -9.48, \\ \frac{1}{200} \sum_{i=1}^{200} (\text{llk}_i(\text{SV}_t) - \text{llk}_i(\text{SV}_{\text{true}})) &= -9.14 \quad \text{and} \\ \frac{1}{200} \sum_{i=1}^{200} (\text{llk}_i(\text{SV}_{\text{sp}}) - \text{llk}_i(\text{SV}_{\text{true}})) &= -3.52.\end{aligned}$$

On average, the SV_{sp} model hence fitted the out-of-sample data substantially better than its parametric counterparts. Considering the individual simulation runs, the SV_{sp} model had a better predictive performance than the two parametric models in 176 out of 200 cases. These results are hardly surprising given the skewness of the distribution chosen for ε_t . Nevertheless, they do demonstrate both the practical feasibility and the potential benefits of our approach.

In order to have a benchmark for these results, we ran a second simulation study in which we generated data from the SV_t model, specifying $\varepsilon_t^{(0)}$ in (1) to be a Student-t distribution with 10 degrees of freedom and $\mu = 0.02$. Except of the true distribution of the conditional distribution, this second simulation experiment was configured exactly as the first. Again considering the three models SV_0 , SV_t and SV_{sp} , we obtained

$$\begin{aligned}\frac{1}{200} \sum_{i=1}^{200} (\text{llk}_i(\text{SV}_0) - \text{llk}_i(\text{SV}_{\text{true}})) &= -5.44, \\ \frac{1}{200} \sum_{i=1}^{200} (\text{llk}_i(\text{SV}_t) - \text{llk}_i(\text{SV}_{\text{true}})) &= -0.62 \quad \text{and} \\ \frac{1}{200} \sum_{i=1}^{200} (\text{llk}_i(\text{SV}_{\text{sp}}) - \text{llk}_i(\text{SV}_{\text{true}})) &= -2.30.\end{aligned}$$

In this scenario, the (correct) SV_t model had a better predictive performance than the two other models in 152 out of 200 cases, while the SV_{sp} model still showed a better predictive performance than the (incorrect) SV_0 model in 164 cases. Thus, overall, the results demonstrate a) the potential of the nonparametric approach to considerably improve the predictive capacity in scenarios where the true conditional distribution deviates from the functional form imposed by either the normal or the Student-t distribution (e.g., if it is skewed), and b) that it can perform almost as well as parametric modelling approaches in scenarios where those are adequate.

4 Application to stock returns

4.1 The data

The SV_{sp} model was fitted to series of daily log-returns for three stocks, namely Sony Corporation, Merck & Co. and Microsoft Corporation, and for the stock index S&P 500. The adjusted closing prices, p_t , for the period 03.01.2000–01.08.2013, were downloaded from “finance.yahoo.com”, and the daily log-returns were computed as $y_t = \log(p_t/p_{t-1})$. To assess the out-of-sample predictive performance of various models, we divided each of the four series into two parts:

- In-sample period: 03.01.2000–31.12.2007,
- Out-of-sample period: 02.01.2008–01.08.2013.

The dividing date was chosen to lie before the outburst of the recent financial crisis, which culminated in the collapse of Lehman Brothers Holdings Inc. in September 2008. The four time series that were analyzed are displayed in Figure 2.




Figure 2: Time series of log-returns on shares of three companies (Sony, Merck and Microsoft) and on the stock index S&P 500; the observations from the in-sample and out-of-sample periods are displayed in black and gray, respectively.

For comparison purposes, we also fitted the two basic models SV_0 and SV_t to each of the four series, using the same partition into in-sample period and out-of-sample period. All models were fitted using the data from the in-sample period only, and the data from the out-of-sample period were used to assess the predictive performance of the three different models (as detailed below).

4.2 Results

To each of the four series, the SV_{sp} model was fitted using $K = 20$ and hence 41 B-spline basis densities to represent the density of ε_t , $m = 100$ intervals in the discretization of the log-volatility process, numerically integrating over the log-volatility values from the interval $[b_0, b_{100}] = [-5, 5]$. Smoothing parameters were selected via cross-validation as described in Section 2.3.2, in each of $C = 40$ cross-validation partitions using 90% of the data points in the calibration stage. Fitting the SV_{sp} model took about 10 minutes per series on an i7 CPU, at 2.7 GHz and with 4 GB RAM.

Table 1: Parameter estimates that were obtained when fitting the three different models to the four series of log-returns considered.

	SV_0			SV_t				SV_{sp}	
	$\hat{\phi}$	$\hat{\sigma}$	$\hat{\beta}$	$\hat{\phi}$	$\hat{\sigma}$	$\hat{\beta}$	$\hat{\nu}$	$\hat{\phi}$	$\hat{\sigma}$
Sony	0.957	0.249	0.019	0.992	0.092	0.017	6.698	0.994	0.087
Merck	0.825	0.545	0.014	0.992	0.086	0.012	4.670	0.993	0.078
Microsoft	0.979	0.239	0.015	0.994	0.116	0.014	6.305	0.994	0.122
S&P 500	0.991	0.114	0.010	0.992	0.104	0.009	25.724	0.998	0.088

The estimates of the parameters ϕ , σ , β (only for the SV_0 and for the SV_t model) and ν (only in case of the SV_t model), for the four series considered, are given in Table 1. The results obtained for the SV_0 model illustrate the problems of the conditional normal distribution to capture the extreme returns. Especially for the Merck stock, where on September 30, 2004, the withdrawal of the drug Rofecoxib from the market caused heavy losses, SV_0 performs badly. Indeed, the only way for the SV_0 model to cope with the associated extreme negative return of -0.31 , which occurs in a period of calm market, is to assign a very high uncertainty to the log-volatility process (as expressed by a high $\hat{\sigma}$ and a small $\hat{\phi}$). This results in an undersmoothing of the volatility. By contrast, the SV_t model's leptokurtic conditional distribution leads to much more plausible estimates for ϕ and σ , with the results being sim-

ilar to those obtained for the SV_{sp} model. In these two models, extreme returns are assigned to the tail of the return distribution, rather than to big jumps in the log-volatility process, as in the SV_0 model. The same pattern is found for the other two stock return series — although to a lesser extent. Only for the stock index S&P 500, the estimate of σ obtained using the SV_0 model is of the same magnitude as the corresponding estimates obtained when applying SV_t and SV_{sp} . This is not surprising since in the index the extreme returns of individual companies play a smaller role, which is also reflected by the much lighter tail of the conditional distribution in the fitted SV_t model.

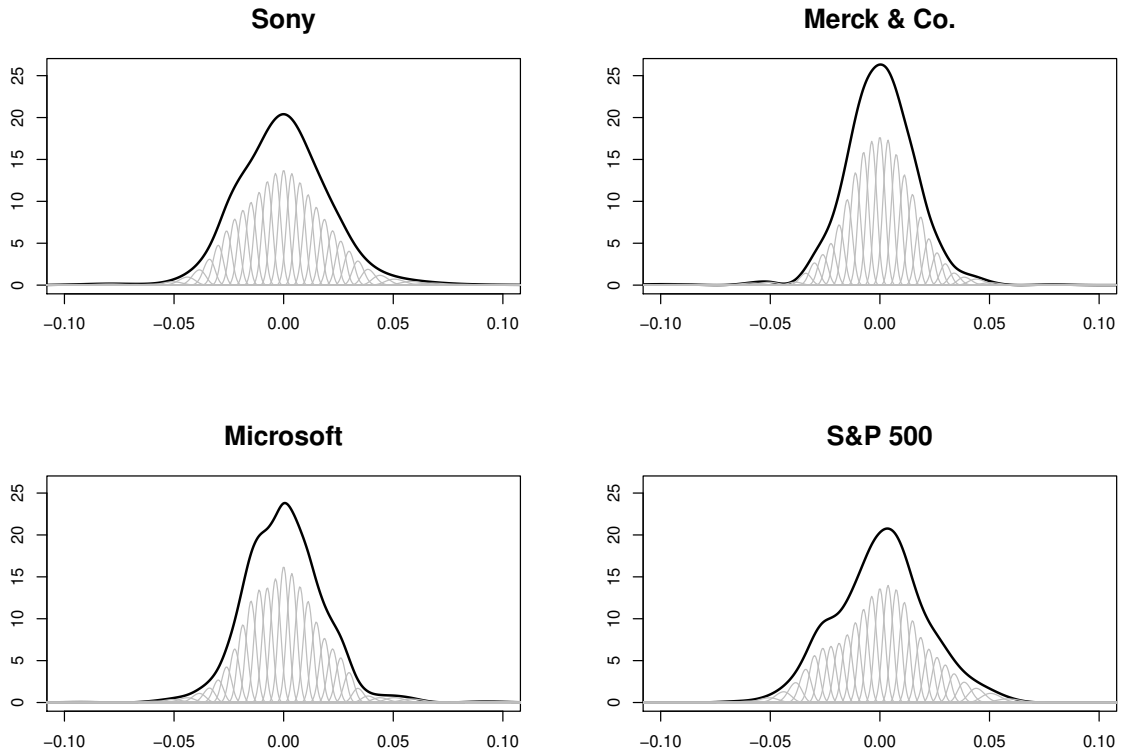


Figure 3: Conditional densities of ε_t estimated using the nonparametric approach, for the four series of log-returns considered, and underlying weighted B-splines that generate these densities via a linear combination (in grey).

Figure 3 displays the nonparametrically estimated densities of the conditional distribution in the SV_{sp} model, for the four series considered. The skewness of the nonparametrically estimated distribution f_ε is -0.73 , -2.61 , -1.23 and -0.77 , for Sony, Merck, Microsoft and S&P 500, respectively. This is in line with the stylized facts attributed to financial assets which propagate a gain/loss asymmetry as large drawdowns generally exceed large upward movements (Cont, 2001), resulting in a left tail in the conditional distribution that is more

extreme than the right tail (Durham, 2006). However, to some extent the skewness also stems from an asymmetric density close to the center, a phenomenon that could be related to the lack of a leverage effect in our model. Indeed, Delatola and Griffin (2013) gave evidence that if data stem from an SV model with a strong leverage effect, then a nonparametric SV model not including this effect may infer a multimodal conditional distribution.

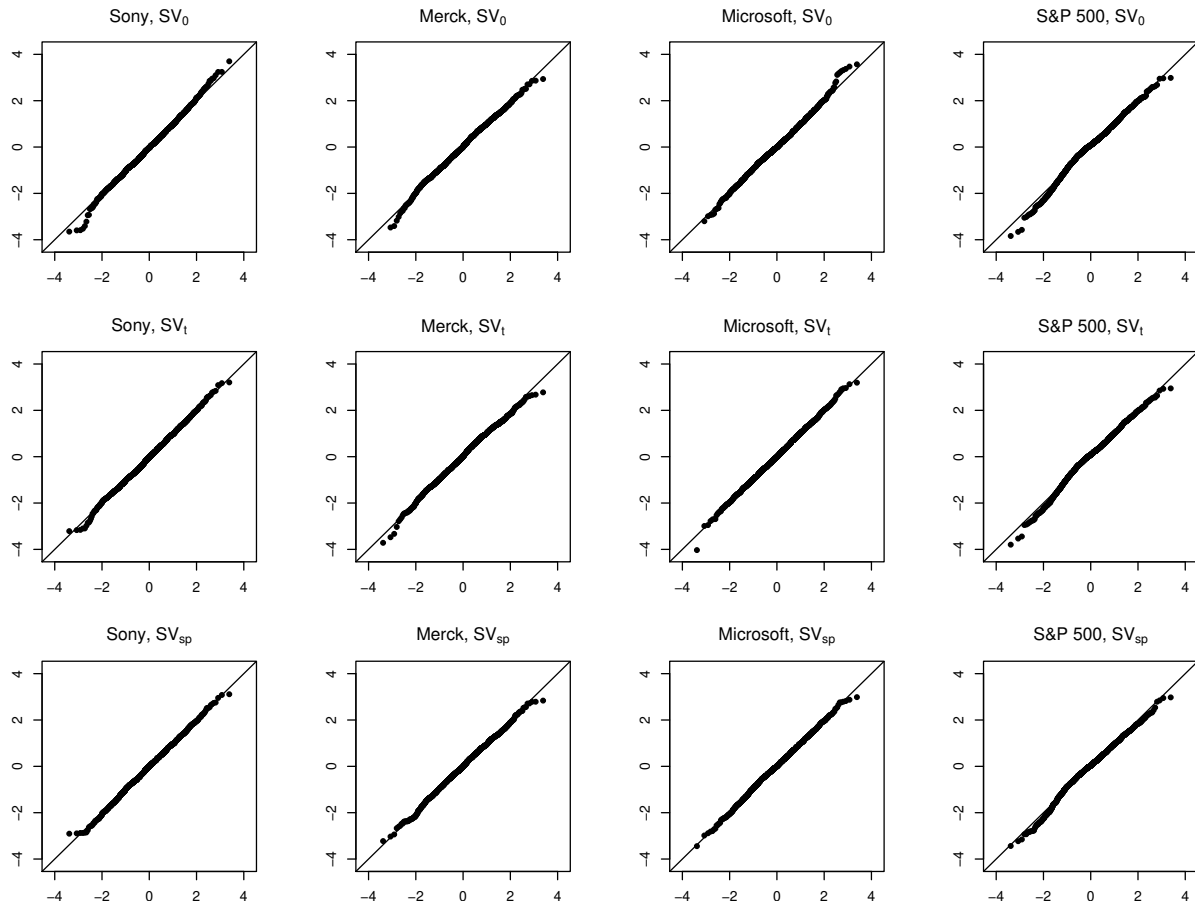


Figure 4: Assessment of the absolute fit of the models in terms of their predictive performance: quantile-quantile plots of the out-of-sample forecast pseudo-residuals obtained from the fitted models SV_0 (top row), SV_t (second row) and SV_{sp} (bottom row), for Sony (first column), Merck (second column), Microsoft (third column) and S&P 500 (fourth column). Sample quantiles are given on the vertical axes, and quantiles of the standard normal are given on the horizontal axis.

Quantile-quantile plots of the out-of-sample one-step-ahead forecast pseudo-residuals associated with the returns observed in the out-of-sample period 02.01.2008–01.08.2013, under the models fitted to the data from the in-sample period, are given in Figure 4. The p-values for the Jarque–Bera tests applied to the pseudo-residuals are listed in Table 2. The results show that the SV_0 model provides poor out-of-sample forecasts for all three stock price time series, as the corresponding pseudo-residuals show large deviations from normality. The SV_t

model is able to adequately forecast the Sony and Microsoft stocks, but exhibits some problems in the forecasts for the Merck series. Overall, the SV_{sp} shows a slightly higher accuracy in the forecasts of the stock returns, especially with respect to the extreme negative returns. For the S&P 500 index, all three models perform badly, with the Jarque-Bera test rejecting the null hypothesis of normally distributed pseudo-residuals in each case. While for the models SV_0 and SV_t at least part of the reason for the poor performance is the inaccurate forecast of extreme negative returns, for the SV_{sp} model the reason for the poor performance lies solely in the inaccurate forecast of moderate losses, while more extreme losses are again captured well. We believe the reason for the bad performance in this particular case to lie in the long persistent decline of stock prices during the financial and economic crisis which ensued the collapse of Lehman Brothers.

Table 2: Assessment of the absolute fit of the models in terms of their predictive performance: p-values of Jarque–Bera tests applied to out-of-sample one-step-ahead ahead forecast pseudo-residuals.

	Sony	Merck	Microsoft	S&P 500
SV_0	0.048**	0.002**	< 0.001**	< 0.001**
SV_t	0.896	0.055*	0.795	< 0.001**
SV_{sp}	0.818	0.333	0.407	< 0.001**

As in the simulation study, we also calculated the log-likelihood scores for the observations from the out-of-sample period, for the three different models and each of the four series analyzed. The results of this comparative assessment of the out-of-sample predictive performance of the models are displayed in Table 3. The log-likelihood scores portray a similar picture as above, with the SV_0 model performing worse than the two models SV_t and SV_{sp} (at least for all stock returns), which again perform similarly well. While the SV_{sp} model performs slightly better for the Sony and Merck time series, SV_t has the edge for Microsoft and S&P 500.

With regard to the index S&P 500, as well as other indices, it should be noted that these exhibit less extreme dynamics due to the averaging over several stocks. As a consequence, the simple model SV_0 in this case performs about as well as do the more flexible models SV_t and SV_{sp} . However, the quantile-quantile plots do indicate that even for the index S&P 500 our semiparametric model has a slightly improved forecast accuracy at the extreme end of the lower tail.

Table 3: Assessment of the relative fit of the models in terms of their predictive performance: log-likelihood scores for the out-of-sample period. For each company, the highest score is underlined.

	Sony	Merck	Microsoft	S&P 500
$\text{llk}(\text{SV}_0)$	3265.31	3891.75	3778.64	4228.95
$\text{llk}(\text{SV}_t)$	3269.58	3913.12	<u>3799.58</u>	<u>4230.53</u>
$\text{llk}(\text{SV}_{\text{sp}})$	<u>3271.38</u>	<u>3914.65</u>	3798.64	4228.74

5 Discussion

The stylized facts of asset returns indicate that simple parametric distributions, such as the normal or the Student’s t-distribution, may not be well-suited to describe the shape of the conditional distribution in SV models. Thus, a nonparametric modelling of the conditional distribution, which allows for heavy tails, gain/loss asymmetry and other unusual features, may bear considerable advantages. In this manuscript, we developed a powerful and flexible frequentist framework for a nonparametric estimation of the conditional distribution in a discrete-time SV model. The approach exploits the strengths of the HMM machinery, in particular allowing for model checking, forecasting and volatility estimation.

The computational burden for estimating a model of the proposed type is low, namely in the order of a couple of minutes for the considered series and fixed smoothing parameter. Applying cross-validation techniques to choose a data-driven smoothing parameter is, however, computationally demanding: computing times for this part of the analysis were about 10-15 hours per series we analyzed. These computing times can be substantially reduced by employing parallel computing, allowing for computing times below one hour. Although the model specifications are not directly comparable, Delatola and Griffin (2011) report computing times of up to a day for the model fitting with their Bayesian approach, for series that were slightly longer than the ones we considered and when fixing the concentration parameter for the Dirichlet process mixture (which represents an analogue to the smoothing parameter in our setting).

A technical issue with the presented method which calls for further research concerns the configuration of the B-spline basis densities used in the estimation. We employed an ad hoc approach to account for the fact that in the tails of the conditional distribution only few observations are available to infer the shape of the density. Our approach effectively increases

the penalty for non-smoothness in the tails of the distribution. The use of equally spaced sample quantiles, as suggested by Ruppert (2002), seems a promising avenue to explore in this regard. An alternative would be to follow the literature on adaptive smoothing parameter selection, e.g., Krivobokova et al. (2008), where the smoothing parameter would be specified as another spline function on the log-volatility domain. However, this would considerably increase the complexity of the likelihood-based analysis and may therefore easily lead to numerical or identifiability problems.

While we modelled the conditional distribution in the SV model in a nonparametric way, we still assumed a parametric distribution form of the innovations in the log-volatility process, which is not necessary. Furthermore, the possible incorporation of leverage effects into the model — i.e., the explicit modelling of a (negative) correlation between returns and subsequent log-volatilities, as often done in parametric SV modelling (e.g., Harvey and Shephard, 1996, Jacquier et al., 2004) — was not discussed in the present manuscript, since we felt that in this first step towards a frequentist framework for semiparametric SV modelling it is advisable to focus on the inferential machinery, rather than on exploring the various possible variations in the model structure. Corresponding extensions are to be explored in future research.

Overall, the approach shows promise as a useful novel tool for analyzing time series of daily log-returns. We have illustrated that the approach *can* lead to an improved predictive capacity compared to basic parametric SV models, and in the real data analyses we found some notable distributional shapes. In particular, our model revealed negative skewness and heavy tails in the conditional distribution of the returns we analyzed, while still identifying the behavior of the log-volatilities that is typical of SV models. In out-of-sample comparisons the parametric model with Student-t conditional distribution performed similarly compared to our semiparametric model (and both performed much better than the model with Gaussian conditional distribution, at least for stock returns). It should be noted here that all validation samples that we considered involve rather extreme dynamics as they comprise the recent financial crisis. Further research needs to be done to investigate the performance of our approach in different scenarios, including calmer markets. The present manuscript is intended mainly to introduce the frequentist estimation framework and to outline its potential. While much work remains to be done, we strongly believe that the model’s flexibility with regard to describing asymmetries and extreme events, and the relative accessibility of the maximum

likelihood framework, will render our approach and potential future extensions a useful tool in portfolio management.

Supplementary Material

R and C++ code to 1) generate artificial data as in the simulation study and to 2) fit, to the generated data, the SV_0 , SV_t and SV_{sp} models (format: main R code in .R file and additional C++ code in .cpp file).

Appendix – HMM essentials

This appendix reviews some HMM basics. A standard m -state HMM has the same two process structure as SV models and SSMS, only that the unobserved process is a Markov chain and hence discrete- rather than continuous-valued. Consider an HMM with observable process $\{X_t\}_{t=1}^T$ and underlying Markov chain $\{S_t\}_{t=1}^T$. Given the current state of S_t , the variable X_t is usually assumed to be conditionally independent from previous and future observations and states. The Markov chain is typically considered to be of first order, and the probabilities of transitions between the different states are summarized in the $m \times m$ transition probability matrix $\mathbf{\Gamma} = (\gamma_{ij})$, where $\gamma_{ij} = \Pr(S_{t+1} = j | S_t = i)$, $i, j = 1, \dots, m$. The initial state probabilities are summarized in the vector $\boldsymbol{\pi}$, where $\pi_i = \Pr(S_1 = i)$, $i = 1, \dots, m$. It is usually convenient and appropriate to assume $\boldsymbol{\pi}$ to be the stationary distribution. For the described HMM, with observations given by x_1, \dots, x_T and underlying states denoted by s_1, \dots, s_T , the likelihood is given by

$$\begin{aligned} \mathcal{L}^{\text{HMM}} = f(x_1, \dots, x_T) &= \sum_{s_1=1}^m \dots \sum_{s_T=1}^m f(x_1, \dots, x_T | s_1, \dots, s_T) f(s_1, \dots, s_T) \\ &= \sum_{s_1=1}^m \dots \sum_{s_T=1}^m \pi_{s_1} \prod_{t=1}^T f(x_t | s_t) \prod_{t=2}^T \gamma_{s_{t-1}, s_t}. \end{aligned}$$

In this form the likelihood involves m^T summands, which would make a numerical maximization infeasible in most cases. However, there is a much more efficient way of calculating the likelihood \mathcal{L}^{HMM} , given by a recursive scheme called the *forward algorithm*. To see this, we consider the vectors of forward variables, defined as $\boldsymbol{\alpha}_t = (\alpha_t(1), \dots, \alpha_t(m))$, $t = 1, \dots, T$,

where $\alpha_t(j) = f(x_1, \dots, x_t, S_t = j)$, $j = 1, \dots, m$. We then have the recursion:

$$\boldsymbol{\alpha}_1 = \boldsymbol{\pi}\mathbf{Q}(x_1), \quad \boldsymbol{\alpha}_{t+1} = \boldsymbol{\alpha}_t\boldsymbol{\Gamma}\mathbf{Q}(x_{t+1}), \quad (10)$$

where $\mathbf{Q}(x_t) = \text{diag}(f_1(x_t), \dots, f_m(x_t))$, with $f_i(x_t) = f(x_t|S_t = i)$. The recursion (10) can be derived in a straightforward manner using the HMM dependence structure. The likelihood can then be written as a matrix product:

$$\mathcal{L}^{\text{HMM}} = \sum_{i=1}^m \alpha_T(i) = \boldsymbol{\pi}\mathbf{Q}(x_1)\boldsymbol{\Gamma}\mathbf{Q}(x_2)\dots\boldsymbol{\Gamma}\mathbf{Q}(x_T)\mathbf{1},$$

where $\mathbf{1} \in \mathbb{R}^m$ is a column vector of ones. For more details on HMMs, see, e.g., Zucchini and MacDonald (2009).

References

- Abanto-Valle, C., Bandyopadhyay, D., Lachos, V. and Enriquez, I. (2010). Robust Bayesian analysis of heavy-tailed stochastic volatility models using scale mixtures of normal distributions. *Computational Statistics & Data Analysis*, **54**, 2883–2898.
- Abraham, B., Balakrishna, N. and Sivakumar, R. (2006). Gamma stochastic volatility models. *Journal of Forecasting*, **25**, 153–171.
- Bartolucci, F. and De Luca, G. (2001). Maximum likelihood estimation of a latent variable time-series model. *Applied Stochastic Models in Business and Industry*, **17**, 5–17.
- Bartolucci, F. and De Luca, G. (2003). Likelihood-based inference for asymmetric stochastic volatility models. *Computational Statistics & Data Analysis*, **42**, 445–449.
- de Boor, C. (1978). *A practical guide to splines*. Springer, Berlin.
- Chib, S., Nardari, F. and Shephard, N. (2002). Markov chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, **108**, 281–316.
- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, **1**, 223–236.

- Delatola, E.-I. and Griffin, J.E. (2011). Bayesian nonparametric modelling of the return distribution with stochastic volatility. *Bayesian Analysis*, **6**, 901–926.
- Delatola, E.-I. and Griffin, J.E. (2013). A Bayesian semiparametric model for volatility with a leverage effect. *Computational Statistics & Data Analysis*, **60**, 97–110.
- Durham, G.B. (2006). Monte Carlo methods for estimating, smoothing, and filtering one- and two-factor stochastic volatility models. *Journal of Econometrics*, **133**, 273–305.
- Eilers, P.H.C. and Marx, B.D. (1996). Flexible smoothing with B -splines and penalties. *Statistical Science*, **11**, 89–121.
- Fridman, M. and Harris, L. (1998). A maximum likelihood approach for non-Gaussian stochastic volatility models. *Journal of Business & Economic Statistics*, **16**, 284–291.
- Gallant, A.R., Hsieh, D. and Tauchen, G.E. (1997). Estimation of stochastic volatility models with diagnostics. *Journal of Econometrics*, **81**, 159–192.
- Gneiting, T. and Raftery, A.E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, **102**, 359–378.
- Harvey, A.C. and Shephard, N. (1996). Estimation of an asymmetric stochastic volatility model for asset returns. *Journal of Business & Economic Statistics*, **14**, 429–434.
- Harvey, C.R. and Siddique, A. (2000). Conditional skewness in asset pricing tests. *Journal of Finance*, **55** 1263–1295.
- Jacquier, E., Polson, N.G. and Rossi, P.E. (2004). Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. *Journal of Econometrics*, **122**, 185–212.
- Jondeau, E. and Rockinger, M. (2003). Conditional volatility, skewness and kurtosis: existence, persistence and co-movements. *Journal of Economic Dynamics and Control*, **27**, 1699–1737.
- Jensen, M.J. and Maheu, J.M. (2010). Bayesian semiparametric stochastic volatility modeling. *Journal of Econometrics*, **157**, 306–316.
- Kim, S., Shephard, N. and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *The Review of Economic Studies*, **65**, 361–393.

- Krivobokova, T., Crainiceanu, C.M. and Kauermann, G. (2008). Fast adaptive penalized splines. *Journal of Computational and Graphical Statistics*, **17**, 1–20.
- Langrock, R. (2011). Some applications of nonlinear and non-Gaussian state-space modelling by means of hidden Markov models. *Journal of Applied Statistics*, **38**, 2955–2970.
- Langrock, R., MacDonald, I.L. and Zucchini, W. (2012). Some nonstandard stochastic volatility models and their estimation using structured hidden Markov models. *Journal of Empirical Finance*, **19**, 147–161.
- Langrock, R. and King, R. (2013). Maximum likelihood estimation of mark-recapture-recovery models in the presence of continuous covariates. *Annals of Applied Statistics*, **7**, 1709–1732.
- Ruppert, D. (2002). Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics*, **11**, 735–757.
- Ruppert, D., Wand, M.P. and Carroll, R.J. (2003). *Semiparametric Regression*. Cambridge University Press, Cambridge.
- Schellhase, C. and Kauermann, G. (2012). Density estimation and comparison with a penalized mixture approach. *Computational Statistics*, **27**, 757–777.
- Shephard, N. (1996). Statistical aspects of ARCH and stochastic volatility. In *Time Series Models: In econometrics, finance and other fields*, Cox, D.R., Hinkley, D.V. and Barndorff-Nielsen, O.E. (eds.), 1–67. Chapman & Hall, London.
- Zucchini, W. and MacDonald, I.L. (2009). *Hidden Markov Models for Time Series: An Introduction Using R*. Chapman & Hall, London.