

Unobserved Components with Stochastic Volatility: Simulation-based Estimation and Signal Extraction

Mengheng Li¹ Siem Jan Koopman^{2,3,4*}

¹ *UTS Business School, University of Technology Sydney, Australia*

² *School of Business and Economics, Vrije Universiteit Amsterdam, The Netherlands*

³ *CREATES, Aarhus University, Denmark*

⁴ *Tinbergen Institute Amsterdam, The Netherlands*

Abstract

The unobserved components time series model with stochastic volatility has gained much interest in econometrics, especially for the purpose of modelling and forecasting inflation. We present a feasible simulated maximum likelihood method for parameter estimation from a classical perspective. The method can also be used for evaluating the marginal likelihood function in a Bayesian analysis. We show that our simulation-based method is computationally feasible, for both univariate and multivariate models. We assess the performance of the method in a Monte Carlo study. In an empirical study, we analyse U.S. headline inflation using different univariate and multivariate model specifications.

Keywords: Importance Sampling, Kalman Filter, Monte Carlo Simulation, Stochastic Volatility, Unobserved Components Time Series Model, Inflation Forecasting.

JEL Classification: C32, C53, E31, E37.

*Corresponding author: S J Koopman, Department of Econometrics and Data Science, School of Business and Economics, Vrije Universiteit Amsterdam, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands. Email s.j.koopman@vu.nl; Telephone +31205986019; Fax +31205986020. The Online Appendix provides more material and results; it is made available at <http://sjkoopman.net/papers/LK2019Appendix.pdf>. SJK would like to acknowledge financial support by the Center for Research in Econometric Analysis of Time Series (DNRF78), CREATES, funded by the Danish National Research Foundation. We would like to thank Co-Editor Marco Del Negro and the referees for their valuable comments and support.
This version: 8 November 2020.

1 Introduction

The unobserved components (UC) time series model decomposes a time series into a trend component driven by permanent shocks and a cycle component driven by transitory shocks. For macroeconomic time series, the former typically relies on non-stationary dynamics while the latter is usually modelled through stationary dynamics. Hence, both the trend of the time series and the deviation from it are treated as latent dynamic stochastic processes. The most basic example is the local level model with a random walk process for the time-varying trend and a white noise process for the cycle component; see [Harvey \(1989\)](#) for further details.

The unobserved components model extended with stochastic volatility (UCSV) for both the permanent and transitory shocks, proposed and developed by [Stock and Watson \(2007, 2008\)](#), has gained much interest in studies on the modelling and forecasting of inflation. [Chan \(2013\)](#) compares UCSV models with different specifications for the cycle component. Other UCSV model specifications are considered by [Harvey \(2011\)](#), [Kim et al. \(2014\)](#), [Mertens and Nason \(2017\)](#) and [Cecchetti et al. \(2017\)](#). These studies show that forecast functions implied by UCSV models are flexible and can deliver accurate forecasts for inflation.

The multivariate extension of the UCSV model can accommodate survey data of inflation expectations to anchor the trend component and can embed structural relations implied by the Phillips curve and Okun's law in the cycle component; see, for example, [Harvey \(2011\)](#), [Kim et al. \(2014\)](#) and [Hasenzagl et al. \(2020\)](#). The multivariate UCSV model with a multivariate random walk for the trend and a vector autoregressive (VAR) process for the cycle is considered by [Mertens \(2016\)](#) and [Chan et al. \(2018\)](#). Finally, [Stock and Watson \(2016\)](#) consider a factor model with SV for modelling sectoral inflation data.

The tasks of parameter estimation and forecasting for univariate and multivariate UCSV models are typically based on Bayesian Markov chain Monte Carlo methods, as in [Stock and Watson \(2007, 2008\)](#) and [Chan \(2017\)](#), and possibly in combination with particle filtering methods, as in [Shephard \(2015\)](#). Given the lack of analytical expressions for the likelihood function of UCSV models, the classical method of maximum likelihood does also rely on simulation-based methods. For both Bayesian and classical approaches, the major challenge for estimation is to integrate out the unobserved components and the stochastic volatility in a numerically efficient manner for the purpose of evaluating the (marginal) likelihood function.

For this purpose, we develop a simulation-based importance sampling procedure for UCSV models. We adopt this procedure to enable maximum likelihood estimation, which is a novel development in the context of (multivariate) UCSV models.

The importance sampling method is developed and discussed in earlier work for a general class of nonlinear non-Gaussian state space models; see, for example, [Shephard and Pitt \(1997\)](#), [Durbin and Koopman \(1997\)](#) and [Richard and Zhang \(2007\)](#). The method is regularly used and illustrated for the maximum likelihood estimation of parameters in the stochastic volatility model and its extensions; see [Sandmann and Koopman \(1998\)](#), [Koopman and Hol Uspensky \(2002\)](#) and [Liesenfeld and Richard \(2003\)](#), among others. A modification of the importance sampling method for the purpose of increasing computational efficiency in its application to state space models is developed by [Koopman et al. \(2015\)](#). In our current study, we further improve this modification in two directions, benefiting its application to both univariate and multivariate UCSV models. Our proposed simulation-based maximum likelihood treatment of the UCSV model is illustrated in a Monte Carlo study where we provide evidence of its accuracy and its computational efficiency in comparison with particle filtering. The UCSV model and its multivariate extensions are further illustrated in an empirical study for the modelling and forecasting of U.S. headline inflation. We show that our simulation-based method is capable of handling 18 time series jointly within a multivariate UCSV model.

The remainder of the paper is organised as follows. [Section 2](#) formulates the UCSV model. [Section 3](#) develops the simulated maximum likelihood method. [Section 4](#) conducts the Monte Carlo study. [Section 5](#) presents an empirical illustration of modelling and forecasting U.S. inflation. We conclude with a discussion in [Section 6](#).

2 Unobserved components with stochastic volatility

We introduce the univariate and multivariate unobserved components time series models with stochastic volatility (UCSV) in general terms; further details are provided in [Sections 4](#) and [5](#) where the results of our Monte Carlo and empirical studies are presented, respectively.

2.1 Univariate model specification

The UCSV model decomposes the observed univariate time series y_1, \dots, y_T into a trend component π_t and a cycle component ψ_t , as given by

$$\begin{aligned} y_t &= \pi_t + \psi_t, & \pi_t &= \pi_{t-1} + \eta_t, & \eta_t &\sim \text{NID}(0, \exp h_{\eta,t}), \\ \psi_t &= \varphi\psi_{t-1} + \varepsilon_t + \vartheta\varepsilon_{t-1}, & \varepsilon_t &\sim \text{NID}(0, \exp h_{\varepsilon,t}), \end{aligned} \quad (1)$$

for $t = 1, \dots, T$, where π_t is specified as a random walk process driven by permanent shocks η_t , and where ψ_t is specified as a stationary autoregressive moving average (ARMA) process driven by transitory shocks ε_t ; the coefficients φ and ϑ are treated as unknown parameters. The shocks η_t and ε_t are normally and independently (serially and mutually) distributed (NID) with means equal to zero and variances equal to $\exp h_{\eta,t}$ and $\exp h_{\varepsilon,t}$, respectively. The stochastically time-varying log-variances $h_{\eta,t}$ and $h_{\varepsilon,t}$ are modelled jointly by the bivariate random walk processes, as given by

$$\begin{aligned} h_{\eta,t} &= h_{\eta,t-1} + \sigma_\eta \zeta_{\eta,t}, & \begin{pmatrix} \zeta_{\eta,t} \\ \zeta_{\varepsilon,t} \end{pmatrix} &\sim \text{NID} \left(0, \begin{bmatrix} 1 & \rho_{\eta\varepsilon} \\ \rho_{\eta\varepsilon} & 1 \end{bmatrix} \right), \\ h_{\varepsilon,t} &= h_{\varepsilon,t-1} + \sigma_\varepsilon \zeta_{\varepsilon,t}, \end{aligned} \quad (2)$$

for $t = 1, \dots, T$, where the disturbances $\zeta_{\eta,t}$ and $\zeta_{\varepsilon,t}$ are normally distributed with zero means and unity variances, are scaled by the strictly positive volatility-of-volatility coefficients σ_η and σ_ε , respectively, are serially independent, and are mutually dependent with correlation $-1 \leq \rho_{\eta\varepsilon} \leq 1$. The coefficients σ_η , σ_ε and $\rho_{\eta\varepsilon}$ are treated as unknown parameters.

The UCSV model is originally introduced by [Stock and Watson \(2007\)](#) for the modelling and forecasting of U.S. inflation; they specify the model as given by (1)-(2) with $\varphi = \vartheta = 0$, that is $\psi_t = \varepsilon_t$. This model specification is also analysed in detail by [Shephard \(2015\)](#). The UCSV model with ψ_t modelled as an ARMA process is studied by [Chan \(2013\)](#). Other dynamic specifications for the components can also be considered. For example, [Stock and Watson \(1998\)](#) formulate an integrated random walk process for π_t , while [Harvey \(2011\)](#) adopts a stochastically evolving cyclical process for ψ_t . Finally, [Gordon \(1990\)](#), [Mertens \(2016\)](#) and [Chan et al. \(2018\)](#) consider model extensions by including exogenous effects and other stochastic variables in (1).

2.2 Multivariate model specification

The multivariate version of the UCSV model is obtained by considering y_t as an $N \times 1$ observation vector, that is $y_t = (y_{1,t}, \dots, y_{N,t})'$ where $y_{i,t}$ is the i th time series variable of y_t , for $i = 1, \dots, N$. We also treat all other variables in (1) as $N \times 1$ vectors, while the ARMA scalar coefficients φ and ϑ become $N \times N$ matrices Φ and Θ , respectively, with possible parsimonious specifications $\Phi = \varphi \cdot I_N$ and $\Theta = \vartheta \cdot I_N$, where I_N is the $N \times N$ identity matrix. The vector disturbances η_t and ε_t in (1) are NID with zero mean vectors and variance matrices $\Sigma_{\eta,t}$ and $\Sigma_{\varepsilon,t}$, respectively. Different specifications for the variance matrices can be considered for the multivariate UCSV (MUCSV) model. For example, in their analysis of sectoral inflation, [Stock and Watson \(2016\)](#) adopt the MUCSV model with both the trend and cycle decomposed into common and idiosyncratic factors. Their model can be represented as in (1) with $\varphi = \vartheta = 0$, that is $y_{i,t} = \pi_{i,t} + \varepsilon_{i,t}$, and with the decompositions

$$\pi_{i,t} = \omega_{\eta,i} \pi_t^c + \pi_{i,t}^*, \quad \varepsilon_{i,t} = \omega_{\varepsilon,i} \varepsilon_t^c + \varepsilon_{i,t}^*, \quad (3)$$

where π_t^c and ε_t^c are the common factors while $\pi_{i,t}^*$ and $\varepsilon_{i,t}^*$ are the idiosyncratic factors, and with time-invariant loadings $\omega_{\eta,i}$ and $\omega_{\varepsilon,i}$ for the common trend and cycle, respectively, for $i = 1, \dots, N$. The loadings are subject to the identification restrictions $\omega_{\eta,1} = \omega_{\varepsilon,1} = 1$. We do not consider time-varying loadings as in the analysis of [Stock and Watson \(2016\)](#). The common and idiosyncratic trend factors are modelled as random walks, we have $\pi_t^c = \pi_{t-1}^c + \eta_t^c$ and $\pi_{i,t}^* = \pi_{i,t-1}^* + \eta_{i,t}^*$, implying that $\pi_{i,t}$ is also modelled as a random walk since from (3) we have

$$\pi_{i,t} = \omega_{\eta,i} (\pi_{t-1}^c + \eta_t^c) + \pi_{i,t-1}^* + \eta_{i,t}^* = \omega_{\eta,i} \pi_{t-1}^c + \pi_{i,t-1}^* + \omega_{\eta,i} \eta_t^c + \eta_{i,t}^* = \pi_{i,t-1} + \eta_{i,t},$$

where $\eta_{i,t} = \omega_{\eta,i} \eta_t^c + \eta_{i,t}^*$, for $i = 1, \dots, N$. Hence, the shocks in the model $y_{i,t} = \pi_{i,t} + \varepsilon_{i,t}$ of [Stock and Watson \(2016\)](#) are subject to the decomposition

$$\xi_{i,t} = \omega_{\xi,i} \xi_t^c + \xi_{i,t}^*, \quad \xi_t^c \sim \text{NID}(0, \exp h_{\xi,t}^c), \quad \xi_{i,t}^* \sim \text{NID}(0, \exp h_{\xi,i,t}^*), \quad \xi = \eta, \varepsilon, \quad (4)$$

where $h_{\xi,t}^c$ and $h_{\xi,i,t}^*$ are specified as $h_{\xi,t}$ in (2) and scaled by σ_ξ^c and $\sigma_{\xi,i}^*$, respectively, for $\xi = \eta, \varepsilon$ and $i = 1, \dots, N$. Similar to specification (2), we allow for correlation between the two common shocks given by $-1 \leq \rho_{\eta\varepsilon}^c \leq 1$. Hence, the disturbance variance matrices $\Sigma_{\xi,t} = \text{Var}(\xi_t)$, for $\xi = \eta, \varepsilon$, are given by

$$\Sigma_{\xi,t} = \exp(h_{\xi,t}^c) [\omega_\xi \omega_\xi'] + \Omega_{\xi,t}, \quad \Omega_{\xi,t} = \text{diag} [\exp(h_{\xi,1,t}^*), \dots, \exp(h_{\xi,N,t}^*)], \quad \xi = \eta, \varepsilon, \quad (5)$$

for $t = 1, \dots, T$, with $N \times 1$ vector $\omega_\xi = (\omega_{\xi,1}, \dots, \omega_{\xi,N})'$ and where $\text{diag}[a_1, \dots, a_N]$ is the $N \times N$ diagonal matrix with diagonal elements a_1, \dots, a_N . The number of unknown parameters for each variance matrix is $2 \times N$ (we have $N - 1$ loadings $\omega_{\xi,i}$ and $N + 1$ scalings σ_ξ^c and $\sigma_{\xi,i}^*$). The total number of parameters is $4N + 1$, including the correlation coefficient $\rho_{\eta\varepsilon}^c$. Alternative MUCSV model specifications are presented in Section 5.

3 Simulated maximum likelihood estimation

We develop a parameter estimation method based on simulated maximum likelihood (SML). Conditional on these estimates, similar simulation-based methods enable the signal extraction of unobserved variables such as π_t , ψ_t , $h_{\eta,t}$ and $h_{\varepsilon,t}$. The details are in the Online Appendix.

3.1 Simulated likelihood function

Importance sampling methods are generally used for the evaluation of intractable densities. Their implementation for the maximum likelihood estimation of parameters in standard SV models have been highly effective; for illustrations of importance sampling in this context, see Shephard and Pitt (1997), Durbin and Koopman (1997) and Richard and Zhang (2007), among others. Next we consider importance sampling for the UCSV model (1)-(2).

The likelihood function of the UCSV model is denoted by $L(\theta)$ where θ is the vector of model parameters. We collect all stochastically time-varying processes related to the (co)variances at time t in the $k \times 1$ vector h_t , for $t = 1, \dots, T$; for example, for the univariate UCSV model (1)-(2) we have $k = 2$ and $h_t = (h_{\eta,t}, h_{\varepsilon,t})'$. The complete set of (co)variance

paths is represented by the set $H = \{h_1, \dots, h_T\}$. The joint density for y is given by

$$p(y; \theta) = \int p(y, H; \theta) dH = \int p(y|H; \theta) p(H; \theta) dH,$$

where $y = (y'_1, \dots, y'_T)'$. The conditional density $p(y|H; \theta)$ can be evaluated by the Kalman filter since the UCSV model is linear and Gaussian for given trajectories in H . Hence, the Kalman filter acts as a Rao-Blackwellisation step by analytically integrating out latent trend and cycle components in a highly computationally efficient manner; see [Chen and Liu \(2000\)](#). To obtain $p(y; \theta)$, or the likelihood $L(\theta)$, the remaining task is to numerically integrate out H which is of a prohibitively high dimension. For this purpose, we extend the *numerically accelerated importance sampling* (NAIS) method of [Koopman et al. \(2015\)](#) towards (M)UCSV models, as NAIS has originally been designed for univariate SV models. The NAIS method constructs a Gaussian importance density $g(H|y; \theta)$ such that simulations of trajectories of H can mimic simulations from $p(H|y; \theta)$ accurately and computationally efficient. It follows from the developments in [Durbin and Koopman \(1997\)](#) that

$$L(\theta) = L_g(\theta) \omega(\theta), \quad \omega(\theta) = \int \frac{p(y|H; \theta)}{g(y|H; \theta)} g(H|y; \theta) dH, \quad (6)$$

where $L_g(\theta)$ is the Gaussian likelihood function with respect to the importance density model $g(y, H; \theta)$ and $\omega(\theta)$ is the importance weight function. An unbiased and consistent Monte Carlo estimator of $\omega(\theta)$ is given by

$$\widehat{\omega}(\theta) = \frac{1}{M} \sum_{i=1}^M \frac{p(y|H^i; \theta)}{g(y|H^i; \theta)}, \quad H^i \sim g(H|y; \theta), \quad (7)$$

where M is the number of simulations, and where the simulations for H^i , for $i = 1, \dots, M$, are obtained from a simulation smoother applied to the importance density model $g(y, H; \theta)$; see, for example, [Durbin and Koopman \(2002\)](#) and [Chan and Jeliazkov \(2009\)](#). The Monte Carlo estimate of the log-likelihood function is subject to the standard log-normal bias correction and is given by

$$\log \widehat{L}(\theta) = \log L_g(\theta) + \log \widehat{\omega}(\theta) + \widehat{s}(\theta)_\omega^2 / (2M\widehat{\omega}(\theta)^2), \quad (8)$$

where $\widehat{s}(\theta)_\omega^2$ is the sample variance corresponding to the sample average $\widehat{\omega}(\theta)$. In a classical

analysis, the maximum likelihood estimate of θ is obtained via the numerical maximisation of (8). In a Bayesian analysis, the marginal likelihood is obtained by evaluating $\widehat{L}(\theta)$. The Online Appendix provides more details on the importance sampling estimation method applied to the UCSV model.

3.2 Importance density for UCSV: a twofold NAIS modification

The modification of NAIS for UCSV models consists of: (i) the joint treatment of the UC part, consisting of π_t and ψ_t in (1), and the SV part, consisting of $h_{\eta,t}$ and $h_{\varepsilon,t}$ in (2), for constructing the joint importance density $g(y, H; \theta)$; (ii) the introduction of the pruning method.

3.2.1 Modification (i): linear dynamic process in mean

The joint importance density $g(y, H; \theta) = g(y|H; \theta)g(H; \theta)$ has two parts with $g(H; \theta) \equiv p(H; \theta)$ representing the dynamic Gaussian model (2), or one of its multivariate extensions, and with Gaussian observation density $g(y|H; \theta)$ representing an approximation of $p(y|H; \theta)$ as implied by model (1). In addition to the latent SV processes collected in H , the observation density also depends on the latent trend π_t and cycle ψ_t as in (1), but both in a linear way. Conditional on H , the linear dependence allows us to integrate out π_t and ψ_t recursively via the prediction error decomposition applied to the observation model (1). We obtain

$$p(y|H; \theta) = \prod_{t=1}^T p(y_t|y_1, \dots, y_{t-1}, H; \theta) = \prod_{t=1}^T p(v_t|H; \theta), \quad (9)$$

where $v_t \equiv v_t(H; \theta) \sim \text{NID}(0, F_t)$ is the prediction error with its variance $F_t \equiv F_t(H; \theta)$, both are obtained from the Kalman filter; see Durbin and Koopman (2012, Chapter 7). For the construction of the Gaussian importance density $g(H|v; \theta) = \prod_{t=1}^T g(h_t|v; \theta)$ in (7), where $v = (v'_1, \dots, v'_T)'$, we consider the linear Gaussian model $g(v|H; \theta)g(H; \theta) \equiv p(H; \theta) \prod_{t=1}^T g(v_t|H; \theta)$ where h_t is the mean for the observation density $g(v_t|H; \theta)$, for all t . Given the linear and Gaussian structures, we can apply a simulation smoother to generate draws from $g(H|v; \theta)$.

The main step for the construction of an importance density is to find an expression for $g(v|H; \theta) = \prod_{t=1}^T g(v_t|H; \theta)$ such that density $g(H|v; \theta)$ fits $p(H|v; \theta)$ as close as possible, in

terms of some distance measure. For this purpose, we formulate the observation density as

$$g(v_t|H; \theta) = \exp\left(r_t + b_t' h_t - \frac{1}{2} h_t' C_t h_t\right), \quad t = 1, \dots, T, \quad (10)$$

where r_t is an integrating constant and $k \times 1$ vector b_t and $k \times k$ matrix C_t are functions of v and θ . The importance sampling parameters are $b_t \equiv b_t(v; \theta)$ and $C_t \equiv C_t(v; \theta)$ as they interact with h_t in (10), for all t . It can be shown that the Gaussian importance observation density (10) is implied by the model

$$v_t^+ = h_t + \epsilon_t^+, \quad \epsilon_t^+ \sim N(0, C_t^{-1}), \quad t = 1, \dots, T, \quad (11)$$

where $v_t^+ = C_t^{-1} b_t$ such that $v_t^+ \equiv v_t^+(v; \theta)$; it follows that $g(v_t^+|h_t; \theta) \equiv g(v_t|H; \theta)$, for all t .

We obtain optimal importance sampling parameters by solving the same minimum squared error criterion as adopted by [Richard and Zhang \(2007\)](#). We solve

$$\min_{b_t, C_t} \int \lambda_t^2(v_t, h_t; \theta) \omega_t(v_t, h_t; \theta) g(h_t|v; \theta) dh_t, \quad \omega_t(v_t, h_t; \theta) = \frac{p(v_t|H; \theta)}{g(v_t^+|h_t; \theta)}, \quad (12)$$

where $\lambda_t(v_t, h_t; \theta) = \log \omega_t(v_t, H; \theta) = \log p(v_t|H; \theta) - r_t - b_t' h_t + \frac{1}{2} h_t' C_t h_t$, for $t = 1, \dots, T$. Instead of using Monte Carlo simulation to evaluate the integral as done by [Richard and Zhang \(2007\)](#), we adopt the NAIS method of [Koopman et al. \(2015\)](#) that utilises fast numerical integration by means of high precision Gauss-Hermite quadrature. When h_t is univariate, or $k = 1$, the minimisation of (12) reduces to

$$\min_{b_t, C_t} \sum_{j=1}^K w_{tj} \lambda_t^2(v_t, \tilde{h}_{tj}; \theta), \quad w_{tj} = w^*(z_j) \exp(z_j^2 / 2) \omega_t(v_t, \tilde{h}_{tj}; \theta) / \sqrt{2\pi}, \quad (13)$$

with K tabulated abscissae or Gauss-Hermite nodes z_j and associated weights $w^*(z_j)$, and $\tilde{h}_{tj} = \hat{h}_t + \sqrt{V_t} z_j$ for $j = 1, \dots, K$, where \hat{h}_t and V_t are the state estimate and its corresponding variance obtained from Kalman smoother applied to importance model (11). Given that v_t^+ in (11) relies on variables b_t and C_t which need to be optimised, the minimisation (13) is achieved through iteration. Starting with $b_t = 1$ and $C_t = 1$ for all t , the minimisation (13) is solved for each t by a weighted least square (WLS) with K ‘‘dependent’’ variables $\log p(v_t^+|h_{tj}; \theta)$, ‘‘explanatory’’ variables \tilde{h}_{tj} and \tilde{h}_{tj}^2 , and weights w_{tj} , for $j = 1, \dots, K$. The WLS estimates

update b_t and C_t which are used for the next iteration. Convergence of this process is usually achieved within a small number of iterations. Moreover, due to common Gauss-Hermite nodes throughout, only one matrix inversion is needed for all WLS calculations. We typically use $K = 10$ nodes. We discuss more computational details in the Online Appendix.

3.2.2 Modification (ii): multiple latent effects in variance

The multiplicity of the integration in (12) is determined by k , the dimensionality of time-varying processes related to (co)variances of the MUCSV model. For example in (2), integration takes the form

$$\int x(h_t)dh_t = \int \int x(h_{\eta,t}, h_{\varepsilon,t})dh_{\eta,t}dh_{\varepsilon,t},$$

where $x(h_t)$ can be any function of h_t ; for example, $x(h_t) = \lambda_t^2(y_t, h_t; \theta)$ in (12). The numerical burden of integration increases exponentially with k . To manage the curse of dimensionality in the computations, we modify the NAIS method by using a basic pruning device. The top left panel of Figure 1 presents the Gauss-Hermite grid that results in a rectangular set of function evaluations at the nodes (z_{j_1}, z_{j_2}) for $j_1, j_2 = 1, \dots, K$. The pruning cuts off many node combinations that lead to near-zero likelihood values as they are located far in the tails of the bivariate normal distribution. We set the threshold

$$\tau_K := w^*(z_1) \cdot w^*(z_{\lfloor \frac{K+1}{2} \rfloor})/K,$$

where $\lfloor a \rfloor$ is the largest integer smaller than a . By dropping all grid points that correspond to weights smaller than τ_K , we obtain the approximation

$$\int x(h_t)dh_t \approx \sum_{j_1=1}^K \sum_{j_2=1}^K \mathbf{1}_{\{w^*(z_{j_1})w^*(z_{j_2}) \geq \tau_K\}} w_{tj_1} w_{tj_2} x(\tilde{h}_{tj_1}, \tilde{h}_{tj_2}),$$

with indicator function $\mathbf{1}_{\{b\}}$ that takes the value 1 when b is true, and 0 otherwise. The effect of pruning is presented in the top right panel of Figure 1. Since H is a draw from $g(H|v_1^+, \dots, v_T^+; \theta)$ with respect to the model (11), elements in \hat{h}_t are correlated. When the function inputs are correlated, the rectangular Gauss-Hermite grid becomes a parallelogram as obtained from a rotated plane as in the bottom graphs of Figure 1. It is shown in Jäckel (2005) that pruning remains valid after such rotations. Alternatively, a grid can be constructed

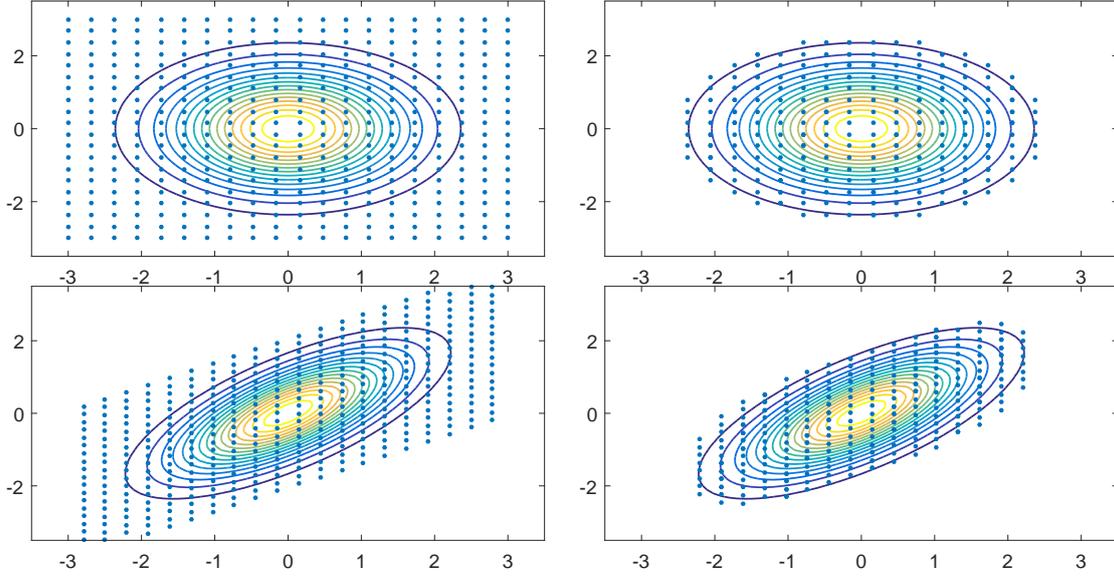


Figure 1
Two-dimensional Gauss-Hermite grid

Top left: Gauss-Hermite grid leads to a rectangular set of function evaluations.

Top right: Effect of pruning. **Bottom left:** Set of function evaluations due to correlation.

Bottom right: Effect of pruning when correlation is present.

using sparse techniques as in Gerstner and Griebel (1998). These efficiency efforts may come with the cost of $g(\cdot)$ being a less accurate approximation of $p(\cdot)$, leading to a higher sampling variance for the importance weight function $\omega(\cdot)$ in (6). However, despite this variance increase and efficiency loss, the importance sampling method itself remains valid.

These NAIS modifications are specifically applicable for MUCSV models, also when many (co)variance processes are present. In a broad range of model specifications, the importance density can be constructed in two consecutive steps. In the first step, we treat the MUCSV model as a set of N separate equations and establish for each series its importance density, and the corresponding trend and cycle estimates. In the second step, we adopt the univariate treatment of a multivariate state space model as described in Durbin and Koopman (2012, Chapter 6), to obtain the joint importance density that accounts of the covariance structures. Such treatments for MUCSV models are described in the Online Appendix.

4 Monte Carlo Evidence

We present the results of two Monte Carlo studies for the univariate UCSV model. First, we assess the accuracy of parameter estimation based on SML for different finite sample sizes. Second, we compare the computational efficiency of our proposed importance sampling method with three different particle filter methods.

4.1 Performance of Simulated Maximum Likelihood Estimation

We first study asymptotic and finite sample properties of estimates obtained via the proposed simulated maximum likelihood (SML) method. We consider the UCSV model (1) - (2) with $\varphi = \vartheta = 0$, such that $\psi_t = \varepsilon_t$. The data generation process (DGP) reduces to the [Stock and Watson \(2007\)](#) model $y_t = \pi_t + \exp(h_{\varepsilon,t}/2)\varepsilon_t$ with random walk state $\pi_{t+1} = \pi_t + \exp(h_{\eta,t}/2)\eta_t$. For $z = \varepsilon, \eta$ and $t = 2, \dots, T$, log-volatility processes are given by $h_{z,t+1} = \alpha_z + \phi_z h_{z,t} + \sigma_z \zeta_{z,t}$ and for $t = 1$, $h_{z,t} \sim N(\frac{\alpha_z}{1-\phi_z}, \frac{\sigma_z^2}{1-\phi_z^2})$. All four disturbances are mutually and serially independent standard normal variates. The true parameter values in the simulations are presented in [Table 1](#). We simulate 1000 time series for each sample size $T = 100, 300, 1000$. In our SML estimation procedure, we use $K = 10$ Gauss-Hermite nodes for obtaining the importance density and $M = 200$ importance samples to compute the log-likelihood function. The key estimation results are presented in [Table 1](#). The results for other DGPs can be delivered upon request.

Overall the estimation accuracy increases as T increases, as is seen by the decreasing sample standard deviation of the 1000 estimates for each DGP. The downward estimation bias for σ_η for $T = 100$ is related to the “pile-up” problem that [Stock and Watson \(1998\)](#) attribute to small sample size. This issue is mitigated for $T = 300$ where the σ_η estimate is close to its true value. Under the column *Asymptotics* in [Table 1](#), the sample average of standard errors of simulated time series with $T = 1000$ is shown. This provides a heuristic diagnostic that supports the validity of asymptotic properties of SML estimates as it is close to the sample standard deviation of estimates for $T = 300$ or $T = 1000$. Furthermore, the negligible Monte Carlo numerical standard deviations suggest that the SML estimates are stable and insensitive to random seeds used in the estimation. Given the SML estimates, the residual diagnostic test statistics for normality, serial correlation and heteroskedasticity ([Durbin and Koopman,](#)

2012, Section 2.12.1) can be computed using standardised prediction errors which are obtained from the particle filter; see Doucet et al. (2001). We conclude that UCSV models with their parameter estimates from SML, capture the dynamic properties in the simulated data well. We also present low fractions of rejections for the diagnostic residual tests among the 1000 replications. It is not surprising that the strong evidence of correct model specification is somewhat weaker for $T = 100$.

Table 1
Results for Monte Carlo Study UCSV Model

	True	Estimated parameters by SML			
		$T = 100$	$T = 300$	$T = 1000$	<i>Asymptotics</i>
α_ε	-0.1	0.23 (0.16)	-0.06 (0.04)	-0.12 (0.04)	0.01 [0.0010]
α_η	-0.2	-6.21 (4.77)	-0.16 (0.07)	-0.17 (0.04)	0.03 [0.0021]
ϕ_ε	0.9	0.67 (0.28)	0.93 (0.10)	0.91 (0.07)	0.08 [0.0000]
ϕ_η	0.9	0.99 (0.03)	0.84 (0.13)	0.90 (0.11)	0.06 [0.0003]
σ_ε	0.3	0.47 (0.34)	0.41 (0.18)	0.32 (0.14)	0.16 [0.0008]
σ_η	0.2	0.08 (0.06)	0.18 (0.03)	0.21 (0.03)	0.04 [0.0016]
Normality		0.18	0.04	0.04	–
Box-Ljung		0.08	0.11	0.06	–
Heteroskedasticity		0.10	0.00	0.00	–
Tail index		4.81 (2.47)	8.74 (1.61)	8.63 (1.21)	2.69 [0.1610]
Variance existence		0.37	0.02	0.00	–

We report sample average of estimates with sample standard deviation in parentheses for σ_z , $z = \pi, y$ across 1000 replications, for three sample sizes. In the last column under *Asymptotics*, the sample average of standard errors is reported, and within the square bracket are the Monte Carlo numerical standard deviations based on 30 different random seeds. The fraction of rejections of residual tests (normality, autocorrelation and, heteroskedasticity) at the 5% level are also reported, using one-step ahead standardised prediction errors produced by the particle filter. Estimated tail index and fraction of rejection for the existence of variance of importance weights are also reported.

To ensure the properties of asymptotic normality for the SML estimates, the distribution of the importance weights must have bounded second moments; see Geweke (1989). The tail index of a distribution indicates the lower bound of bounded moments. We consider the distribution of the importance weights, evaluated at the SML estimate of θ , and estimate its tail index using a method of moments as proposed by Dekkers et al. (1989). We report the sample average and standard deviation of these tail index estimates over the 1000 Monte Carlo

simulations. We also report the rejection rate of a likelihood-based test for the existence of the second moment as proposed by [Koopman et al. \(2009\)](#). We can conclude from the reported results that our proposed importance sampling procedure ensures the existence of a variance for the importance weights, when the sample size is reasonably large, say $T \geq 300$. For cases where this evidence is less strong, and the variance of weights is possibly unbounded, [Chan et al. \(2017\)](#) have developed methods for the bias correction of importance sampling.

4.2 Comparisons with particle filters

Sequential Monte Carlo methods such as the particle filter (PF) are used for solving signal extraction problems related to general Markov processes of which the UCSV model (1) - (2) is an example; see [Doucet et al. \(2001\)](#). By applying the prediction error decomposition to the UCSV model, the PF is also capable of computing an unbiased estimate of its likelihood function $\hat{L}(\theta)$ in (8). The PF can therefore be regarded as an alternative to the importance sampling algorithm based on (6) in our proposed SML method. Given the construction of the PF, and its randomness through the recursive updating, it is not designed for the numerical maximization of $\hat{L}(\theta)$ with respect to θ ; see [Malik and Pitt \(2011\)](#) and [Kantas et al. \(2015\)](#). However, the evaluation of the likelihood $\hat{L}(\theta)$ has also other purposes. For example, it can be used within the particle Markov chain Monte Carlo method where parameter inference is based on $p(\theta|Y) \propto p(Y|\theta)p_0(\theta)$, see [Andrieu et al. \(2010\)](#) and [Shephard \(2015\)](#), and within a Bayesian model comparison where the marginal likelihood $p(Y) = \int p(Y|\theta)p_0(\theta)d\theta$ is computed, see [Shephard and Pitt \(1997\)](#), [Chen and Liu \(2000\)](#), and [Chan \(2017\)](#); in both examples, $p_0(\theta)$ denotes the prior density of θ . The PF replaces $p(Y|\theta)$ with its Monte Carlo estimate $\hat{L}(\theta)$ by sequentially integrating out the latent time-varying processes related to the (co)variances in H via $p(Y|\theta) = \int p(Y|H;\theta)p(H;\theta)dH$. Our importance sampling algorithm can also be implemented in these cases with the notable distinction of integrating out H simultaneously instead of sequentially over time.

The accuracy and efficiency of the methods used in a Bayesian analysis depend on the estimation method for $\hat{L}(\theta)$. Hence, it is of interest to compare the PF with our importance sampling algorithm used in the proposed SML method for evaluating likelihood function. Here we consider the bootstrap filter (BF), the auxiliary PF (APF) of [Pitt and Shephard \(1999\)](#) and

the tempered PF (TPF) of [Herbst and Schorfheide \(2019\)](#). We use the same 1000 generated time series with sample size $T = 300$ as above. The number of particles is set to $M = 1000$ and $M = 5000$, a much larger number than the number in SML where we set $M = 50$ and $M = 100$. All computations are carried out by the object-oriented matrix language `0x` of [Doornik \(2007\)](#) and the library of state space functions `SsfPack` of [Koopman et al. \(2008\)](#) on a quad-core desktop. [Table 2](#) summarises a selection of our main results. We can conclude that the performance of SML is competitive in comparison to PF, even when the number of simulations is as small as $M = 50$: *(i)* the sample variance of the 1000 values of $\log \widehat{L}(\hat{\theta})$ is small; and *(ii)* the tail indices of the importance weight distributions suggest that the simulation methods are well-behaved. More evidence is presented in the Online Appendix.

The BF method shows the worst performance with a large log-likelihood variance and a possibly unbounded second moment for the importance weight distribution. The APF and TPF methods mitigate this by reducing the variance and ensuring the existence of a second moment. However, their computational costs have increased due to more function evaluations taking place in the forward weight resampling for APF, and in the tempering iterations for TPF. The computing time for SML is much smaller than the three competing PF methods. Although PF methods are more general, and we only consider small-scale cases in the Monte Carlo study, we can conclude that the importance sampling algorithm can provide a fast and accurate alternative to the particle filter, whether it is used for SML or in a Bayesian analysis.

Table 2
Likelihood Estimation Comparisons for UCSV Models

Method	SML		BF		APF		TPF	
no. simulations or particles	50	200	1000	5000	1000	5000	1000	5000
Var($\log \widehat{L}(\hat{\theta})$)	14.86	14.84	73.94	52.2	48.07	29.49	44.65	18.76
Tail index	10.48	10.74	1.34	2.61	3.65	5.94	5.06	11.65
Average run time (in seconds)	0.08	0.09	1.14	3.18	4.84	15.00	1.78	6.68

We report the sample average of variances and tail indexes of log-likelihood evaluated at the SML estimates $\hat{\theta}$ across 1000 simulated series with $T = 300$. Sample average of run time for estimating log-likelihood relative to the proposed SML method with $M = 50$ simulations used in estimation is also reported. The acronyms BF, APF and TPF denote bootstrap filter, auxiliary particle filter and tempered particle filter, respectively.

5 Empirical Illustration: U.S. Inflation Forecasting

To illustrate our importance sampling method in an empirical study, we consider UCSV and MUCSV models for the analysis of quarterly U.S. headline inflation. In case of the MUCSV models, we adopt an additional set of 17 U.S. sectoral inflation variables which can be grouped into the sectors of durable goods, non-durable goods, and services; see the Online Appendix for details. All time series range from 1960Q1 up to 2017Q2, with $T = 230$. We consider the univariate UCSV model (1) - (2), plus the restricted cases of $\vartheta = 0$ (UCSV-AR), $\varphi = 0$ (UCSV-MA), and $\varphi = \vartheta = 0$ (UCSV-SW). The UCSV-SW specification reduces to the local level model with SV as considered by [Stock and Watson \(2007\)](#) for forecasting inflation.

For the MUCSV models, we analyse the observation vector $y_t = (y_{1,t}, y_{2,t}, \dots, y_{18,t})$ with headline inflation $y_{1,t}$ and sectoral inflation variables $y_{2,t}, \dots, y_{18,t}$, that is $N = 18$. The details of the MUCSV model are provided in Section 2.2 which imply that $y_{1,t} = \pi_t^c + \pi_{1,t}^* + \varepsilon_{i,t}$ and $y_{i,t} = \omega_{\eta,i} \pi_t^c + \pi_{i,t}^* + \varepsilon_{i,t}$, for $i = 2, \dots, 18$. Hence, we define U.S. trend inflation as $\pi_t^c + \pi_{1,t}^*$. We consider three different specifications for the variance matrices $\Sigma_{\xi,t}$, with $\xi = \eta, \epsilon$, in (5). The specification in Section 2.2 is labelled as MUCSV-SW. The equicorrelation specification MUCSV-EC is given by $\Sigma_{\xi,t} = \Omega_{\xi,t}^{\frac{1}{2}} R_{\xi,t} \Omega_{\xi,t}^{\frac{1}{2}}$ where diagonal variance matrix $\Omega_{\xi,t}$ is defined in (5) and correlation matrix $R_{\xi,t}$ has ones on the diagonal while all off-diagonal elements are set equal to $\rho_{\xi,t} = (\exp(-h_{\xi,t}^\rho) + 1)^{-1}$, and where $h_{\xi,t}^\rho$ is specified as $h_{\xi,t}$ in (2) with scaling parameter σ_ξ^ρ and correlation coefficient $\rho_{\eta\epsilon}^\rho$.

The dynamic factor specification MUCSV-DF allows the permanent or the transitory shocks, or both, to be a linear combination of a small set of $1 \leq r < N$ common factors. In case of the permanent shock, we have $\eta_t = \Lambda_\eta \eta_t^f$ with $N \times r$ factor loading matrix Λ_η , $r \times 1$ disturbance vector $\eta_t^f \sim N(0, \Omega_{\eta,t}^f)$, and $r \times r$ variance matrix $\Omega_{\eta,t}^f = \text{diag}[\exp(h_{\eta,1,t}^f), \dots, \exp(h_{\eta,r,t}^f)]$. It follows that $\Sigma_{\eta,t} = \Lambda_\eta \Omega_{\eta,t}^f \Lambda_\eta'$. In effect, the observation and trend components in (1) can be expressed as a dynamic factor model given by $y_t = \Lambda_\eta \pi_t^f + \psi_t$ with the $r \times 1$ vector of dynamic factors modelled as $\pi_t^f = \pi_{t-1}^f + \eta_t^f$, for $t = 1, \dots, T$. The identification of the parameters in the MUCSV-DF specification is enforced by imposing a unity lower triangular structure for the loading matrix Λ_η . We finally allow for a single correlation coefficient $\rho_{\eta\epsilon}^1$ between the log-volatility innovations for the transitory shock driving the first cycle component (variable of interest) and for the permanent shock driving the first factor.

The UCSV and MUCSV model specifications are reviewed in Table 3 with their properties provided in the Online Appendix. The number of unknown parameters are given as a function of N and r in the last column of Table 3.

Table 3
Summary of UCSV and MUCSV models

Acronym	N	$\Sigma_{\eta,t}$	$\Sigma_{\varepsilon,t}$	Parameters	$\#(N, r)$
UCSV	$= 1$	$\exp h_{\eta,t}$	$\exp h_{\varepsilon,t}$	$\sigma_{\eta}, \sigma_{\varepsilon}, \rho_{\eta\varepsilon}$	3
MUCSV					
-SW	> 1	$\exp(h_{\eta,t}^c) [\omega_{\eta} \omega'_{\eta}] + \Omega_{\eta,t}$	$\exp(h_{\varepsilon,t}^c) [\omega_{\varepsilon} \omega'_{\varepsilon}] + \Omega_{\varepsilon,t}$	$\omega_{\eta}, \omega_{\varepsilon}$	$2(N - 1)$
				$\sigma_{\eta}^c, \sigma_{\varepsilon}^c, \rho_{\eta\varepsilon}^c$	3
				$\sigma_{\eta}^*, \sigma_{\varepsilon}^*$	$2N$
-EC	> 1	$\Omega_{\eta,t}^{\frac{1}{2}} R_{\eta,t} \Omega_{\eta,t}^{\frac{1}{2}}$	$\Omega_{\varepsilon,t}^{\frac{1}{2}} R_{\varepsilon,t} \Omega_{\varepsilon,t}^{\frac{1}{2}}$	$\sigma_{\eta}^*, \sigma_{\varepsilon}^*$	$2N$
				$\sigma_{\eta}^{\rho}, \sigma_{\varepsilon}^{\rho}, \rho_{\eta\varepsilon}^{\rho}$	3
-DF	> 1	$\Lambda_{\eta} \Omega_{\eta,t}^f \Lambda'_{\eta}$	$\Omega_{\varepsilon,t}$	Λ_{η}	$r(N - r)$
					$r(r - 1)/2$
				$\sigma_{\eta}^f, \sigma_{\varepsilon}^*, \rho_{\eta\varepsilon}^{\mathbf{1}}$	$r + N + 1$

We consider the univariate UCSV model (1) - (2) and its multivariate MUCSV model of Section 2.2 for $N \times 1$ observation vector y_t and the shocks $\eta_t \sim \text{NID}(0, \Sigma_{\eta,t})$ and $\varepsilon_t \sim \text{NID}(0, \Sigma_{\varepsilon,t})$. We report the acronym of the model, the specifications of the two variance matrices, the parameter vector and its dimension. For the MUCSV-SW model, the variance matrices correspond to equation (5), the parameter σ_{ξ}^c scales the volatility of the common shock and σ_{ξ}^* does it for the idiosyncratic shocks, with $\xi = \eta, \varepsilon$. For the MUCSV-EC model, the diagonal variance matrix $\Omega_{\xi,t}$ relies on the volatility scale parameters in σ_{ξ}^* and, similarly, the correlation matrix $R_{\xi,t}$ relies on σ_{ξ}^{ρ} , for $\xi = \eta, \varepsilon$. For the MUCSV-DF model, the $r \times r$ diagonal variance matrix $\Omega_{\eta,t}^f$ relies on the volatility scale parameters in σ_{η}^f and the $N \times N$ diagonal variance matrix $\Omega_{\varepsilon,t}$ relies on the volatility scale parameters in σ_{ε}^* . The scalar log-volatility correlation coefficients in the MUCSV models are represented by $\rho_{\eta\varepsilon}^x$, for $x = c, \rho, \mathbf{1}$. Finally, $\#(N, r)$ denotes the number of parameters as a function of N and r , where r is the number of factors in MUCSV-DF. In case $N = 18$ and $r = 3$, we have 73, 39 and 70 parameters for the MUCSV specifications SW, EC and DF, respectively. In addition, we restrict coefficient matrices with unknown parameters as $\varphi = \varphi^* I_k$ and $\vartheta = \vartheta^* I_k$, with scalar coefficients φ^* and ϑ^* , where $k = N$ for SW and EC, and $k = r$ for DF.

5.1 In-sample results

Table 4 presents a selection of our estimation results for the full sample. The estimates of the scaling parameters σ_{η} and σ_{ε} for the permanent and transitory volatilities in (2), respectively, and of the correlation coefficient $\rho_{\eta\varepsilon}$ are provided. For all models, the estimated permanent scaling is smaller compared to its transitory counterpart, while the estimated correlation

is positive and significant for the univariate models and for the MUCSV-EC specification. Hence, the shock in headline inflation has both an immediate and a more lasting impact on the uncertainty in measuring and forecasting inflation. We further report the in-sample root mean squared error (RMSE) and mean predictive likelihood score (MPLS) statistics which reflect in-sample point- and density-fits of Geweke and Amisano (2011), respectively. These goodness-of-fit statistics are based on the one-step ahead prediction errors for U.S. headline inflation which are computed by the Kalman filter, conditional on the time-varying importance sampling estimates of the latent variables in the variances of UCSV models and in the variance matrices of the MUCSV models. The MPLS values are lowest for the multivariate models while the RMSE values are slightly lower for the univariate models. We may conclude that the *overall* fit for U.S. headline inflation is best for MUCSV models; similar findings are reported in Stock and Watson (2016). Interestingly, the MUCSV-EC specification shows a good performance delivering the highest MPLS among all models and lowest RMSE among multivariate models. Similar to Chan (2013), we also find that the RMSE amongst the univariate models is lowest for specifications with a moving average term in the inflation cycle component ψ_t in (1) with $\vartheta \neq 0$. The in-sample RMSE and MPLS statistics for a UC model without SV are equal to 0.31 and -0.71 (not reported), respectively. These statistics indicate that the SV part is important for modelling the dynamic properties of inflation; see Stock and Watson (2007). The estimated tail indices of the importance weights provide evidence that the SML method yields estimates with valid asymptotic properties. We can conclude that the reported indices imply the existence of a bounded second moment for the importance weight distribution; see Geweke (1989).

The importance sampling estimates of the trend component in U.S. headline inflation from our seven models, that is π_t from the four UCSV models and $\pi_t^c + \pi_{1,t}^*$ from the three MUCSV models, are presented in the two panels of Figure 2. In the first panel, we present the seven estimates of the trend, together with the quarterly observations for headline inflation. In the second panel, we present the seven estimates of the trend, together with the 95% confidence interval obtained from the UCSV-SW model. The models with persistence in the inflation cycle ψ_t , such as the UCSV model with a full dynamic specification for ψ_t , deliver trend inflation estimates that are smoother. Overall, the seven estimates of the trend component in headline inflation show a similar gradual decline of inflation since the 1980s.

Table 4
UCSV Model Parameter Estimation Results

	UCSV-SW	-AR	-MA	UCSV	MUCSV-SW	-EC	-DF
σ_η	0.09	0.17	0.21	0.14	0.11, 0.24	0.18	0.13
σ_ε	0.34	0.35	0.38	0.33	0.23, 0.27	0.35	0.17
$\rho_{\eta\varepsilon}$	0.46	0.37	0.41	0.38	0.14	0.57	0.02
RMSE	0.27	0.25	0.21	0.21	0.25	0.22	0.23
MPLS	-0.50	-0.48	-0.44	-0.43	-0.38	-0.35	-0.40
tail-index	12.76	14.24	10.78	11.26	5.46	7.93	3.20

We report scaling parameters σ_η and σ_ε , and correlation coefficient $\rho_{\eta\varepsilon}$ for the UCSV model (1) - (2). For the MUCSV model, the specifications are in provided in Table 3. In case of SW, we present the scale parameters for both the common shock and the idiosyncratic shock in U.S. headline inflation (correlation coefficient is only for the common shock). In case of EC, we present σ_η^ρ , σ_ε^ρ , and $\rho_{\eta\varepsilon}^\rho$. In case of DF, we present σ_η^f , σ_ε^* , and $\rho_{\eta\varepsilon}^1$. We further report (i) the in-sample RMSE of one-step ahead prediction errors for U.S. headline inflation, obtained from the Kalman filter conditional on SV components for each model; (ii) the mean predictive likelihood score (MPLS) based on the same prediction errors; (iii) the estimated tail-index for the importance sampling weights, evaluated at the SML estimates. The in-sample RMSE and MPLS can be regarded as goodness-of-fit statistics. The estimated tail-index is indicative of the number of moments that exists for the importance weight distribution, see Section 4.1.

The importance sampling estimates of permanent volatility $\exp(h_{\eta,t}/2)$ and transitory volatility $\exp(h_{\varepsilon,t}/2)$ are presented in Figure 3. For the MUCSV models, the two volatility processes are displayed for U.S. headline inflation only. We observe a persistent gradual decline in permanent volatility from the 1980s. This is mainly due to the *Great Moderation* that took place after the oil crises in the 1970s and the start of the Volcker-Greenspan monetary policy regime from the early 1980s. The transitory volatility estimates remain low until the financial crisis around 2008. During the financial crisis, the transitory volatility estimates show a two- to three-fold increase in a relatively short period. The estimates from the MUCSV-EC and -DF models present the highest increases. This high-volatility period is short-lived: the transitory volatility estimates have returned to their mid-1980 levels after 2010. The permanent volatility estimates are not affected by the financial crisis. Similar findings for U.S. headline inflation have been reported by Shephard (2015), Stock and Watson (2016) and Kleppe (2019).

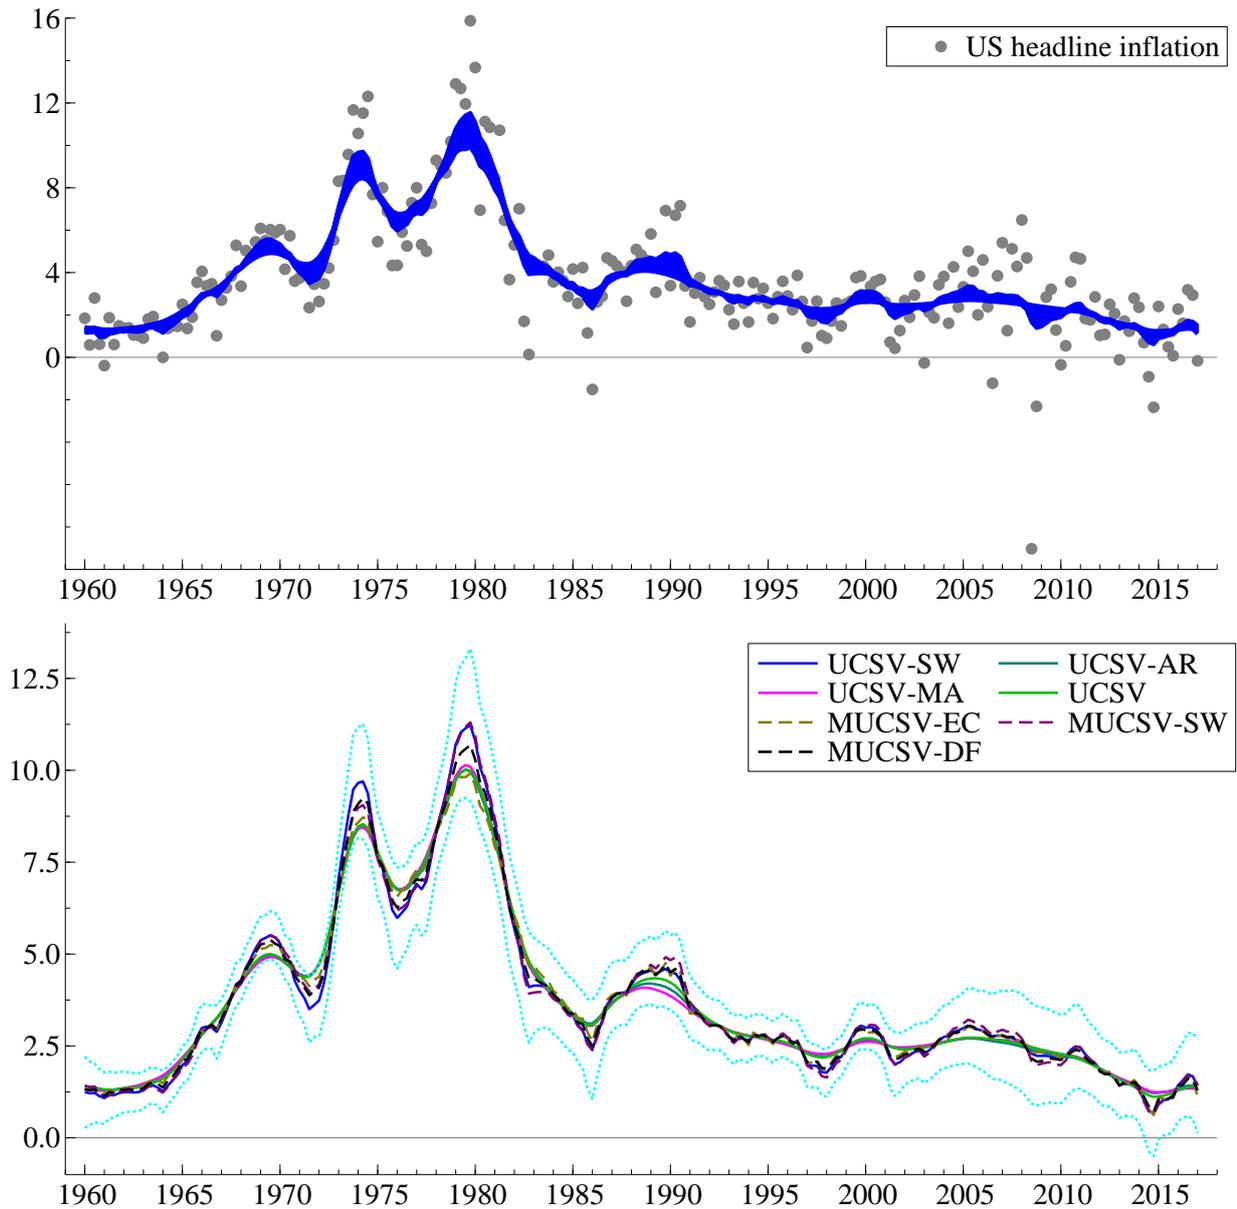


Figure 2
Estimates of U.S. trend inflation

We present importance sampling estimates of trend inflation obtained from UCSV (π_t) and MUCSV ($\pi_t^c + \pi_{1,t}^*$) models. The first panel shows the observed annualized inflation from 1960Q1 to 2017Q2, together with the range of the seven trend inflation estimates. The second panel presents the seven trend inflation estimates together with the 95% confidence interval from the UCSV-SW model (dotted lines), in annualized terms.

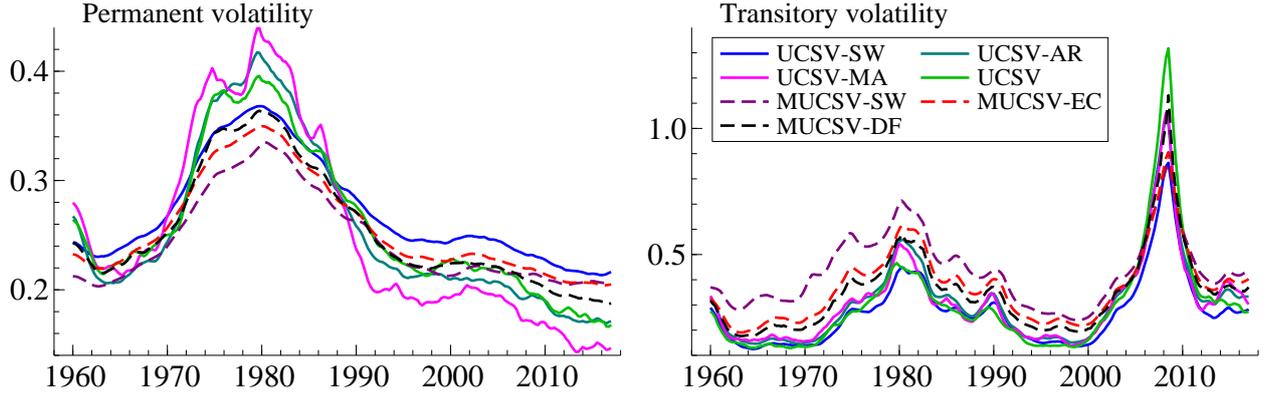


Figure 3
Estimates of stochastic volatility processes

We present the importance sampling estimates of permanent and transitory volatility for the U.S. headline inflation time series, from UCSV and MUCSV model.

5.2 Out-of-sample results

The results from our out-of-sample forecasting study are based on an expanding window exercise, starting from 1990Q1 until 2017Q2 minus h forecast horizons, for $h = 1, 4, 8, 12$ quarters. For each model and each horizon h , we re-estimate the parameters using the SML method and we obtain $106 - h$ forecast errors. When all sets of forecast errors are computed, we compute the out-of-sample RMSE and MPLS values. In Table 5 we report these RMSE and MPLS values relative to their corresponding values for the UCSV-SW model: RMSE as a ratio and MPLS as a difference. On the basis of both RMSE and MPLS, we can conclude that the UCSV models tend to perform better than MUCSV models for shorter forecast horizons. In particular, the forecast performance improves more by including moving average dynamics rather than including autoregressive dynamics. This finding confirms the results reported by Chan (2013). The MUCSV models show a better performance in forecast precision for longer horizons. The superior performance of MUCSV-EC model is convincing when we focus on the density forecasts in this study. We can conclude from the reported results that time-varying correlations across inflation sectors, for both trend and cycle components, are important for the modelling of predictive densities.

Table 5
Point and Density Forecast Performance of UCSV Models

		UCSV				MUCSV		
h		SW	AR	MA	ARMA	SW	EC	DF
RMSE	1	1	0.95	0.90	0.90	0.93	0.97	0.97
	4	1	0.96	0.93	0.93	0.95	0.95	1.06
	8	1	0.97	0.95	0.95	0.94	0.95	1.05
	12	1	0.98	0.96	0.96	0.94	0.95	1.02
MPLS	1	0	0.28	0.34	0.34	-0.06	0.27	0.14
	4	0	0.29	0.31	0.34	-0.02	0.26	-0.04
	8	0	0.21	0.20	0.22	0.00	0.24	0.10
	12	0	0.01	0.06	0.07	0.02	0.09	0.03

The point and density forecast precisions are measured by the out-of-sample root mean square error (RMSE) and mean predictive likelihood score (MPLS) statistics, respectively, for forecast horizon $h = 1, 4, 8, 12$ quarters. The values reported are relative to those from the UCSV-SW model, with RMSE as a ratio and MPLS as a difference. A small RMSE value and a large MPLS value suggest superior performance. The forecasts are obtained from an expanding window starting from 1990Q1; for each window, parameters are re-estimated. The values in bold indicate superior performance over all models.

6 Conclusion

A simulated maximum likelihood estimation method for multivariate unobserved components models with stochastic volatility is developed. The proposed method relies on the conditional linear and Gaussian specification of the model. Conditional on the stochastic volatility part, the Kalman filter can be adopted to evaluate the conditional likelihood function. To obtain the full likelihood function, we need to integrate out the stochastic volatility processes using a computationally efficient importance sampling method. The feasibility of the method is partly due to the use of pruned multivariate Gauss-Hermite quadrature. In a Monte Carlo study, we have provided evidence that our proposed method ensures the existence of variance for importance weights, implying a valid likelihood-based inference procedure. In the empirical study, we have shown the relevance of the stochastic volatility part of the model for the analysis and forecasting of U.S. headline inflation. We further have shown that the SML method for high-dimensional multivariate models is feasible. An interesting avenue for future research is to extend the model further with more time-varying features.

References

- Andrieu, C., A. Doucet, and R. Holenstein (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72(3), 269–342.
- Cecchetti, S. G., P. Hooper, A. K. Kashyap, and K. L. Schoenholtz (2017). Deflating inflation expectations: The implications of inflation’s simple dynamics. In *U.S. Monetary Policy Forum*.
- Chan, J. C. (2013). Moving average stochastic volatility models with application to inflation forecast. *Journal of Econometrics* 176(2), 162–172.
- Chan, J. C. (2017). The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling. *Journal of Business & Economic Statistics* 35(1), 17–28.
- Chan, J. C., T. E. Clark, and G. Koop (2018). A new model of inflation, trend inflation, and long-run inflation expectations. *Journal of Money, Credit and Banking* 50(1), 5–53.
- Chan, J. C., C. Hou, and T. T. Yang (2017). Asymptotic trimming for importance sampling estimators with infinite variance. Technical report.
- Chan, J. C. and I. Jeliazkov (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation* 1(1-2), 101–120.
- Chen, R. and J. S. Liu (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62(3), 493–508.
- Dekkers, A. L., J. H. Einmahl, and L. De Haan (1989). A moment estimator for the index of an extreme-value distribution. *The Annals of Statistics*, 1833–1855.
- Doornik, J. A. (2007). *Object-Oriented Matrix Programming Using Ox, 3rd ed.* London: Timberlake Consultants Press.

- Doucet, A., N. De Freitas, and N. Gordon (2001). An introduction to sequential Monte Carlo methods. In *Sequential Monte Carlo methods in practice*, pp. 3–14. Springer.
- Durbin, J. and S. J. Koopman (1997). Monte Carlo maximum likelihood estimation for non-Gaussian state space models. *Biometrika* 84(3), 669–684.
- Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89(3), 603–616.
- Durbin, J. and S. J. Koopman (2012). *Time series analysis by state space methods* (2nd ed.). Oxford University Press.
- Gerstner, T. and M. Griebel (1998). Numerical integration using sparse grids. *Numerical algorithms* 18(3-4), 209–232.
- Geweke, J. (1989). Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 1317–1339.
- Geweke, J. and G. Amisano (2011). Optimal prediction pools. *Journal of Econometrics* 164(1), 130–141.
- Gordon, R. J. (1990). The Phillips curve now and then. Technical report, National Bureau of Economic Research, Inc.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press.
- Harvey, A. C. (2011). Modelling the Phillips curve with unobserved components. *Applied Financial Economics* 21(1-2), 7–17.
- Hasenzagl, T., F. Pellegrino, L. Reichlin, and G. Ricco (2020). A model of the Fed’s view on inflation. *Review of Economics and Statistics* 102, forthcoming.
- Herbst, E. and F. Schorfheide (2019). Tempered particle filtering. *Journal of Econometrics* 210(1), 26–44.
- Jäckel, P. (2005). A note on multivariate Gauss-Hermite quadrature. *Technical report* (<https://jaeckel.org/ANoteOnMultivariateGaussHermiteQuadrature.pdf>).

- Kantas, N., A. Doucet, S. S. Singh, J. Maciejowski, and N. Chopin (2015). On particle methods for parameter estimation in state-space models. *Statistical Science* 30(3), 328–351.
- Kim, C., P. Manopimoke, and C. R. Nelson (2014). Trend inflation and the nature of structural breaks in the New Keynesian Phillips curve. *Journal of Money, Credit and Banking* 46(2-3), 253–266.
- Kleppe, T. S. (2019). Dynamically rescaled Hamiltonian Monte Carlo for Bayesian hierarchical models. *Journal of Computational and Graphical Statistics*, 1–32.
- Koopman, S. J. and E. Hol Uspensky (2002). The stochastic volatility in mean model: empirical evidence from international stock markets. *Journal of Applied Econometrics* 17(6), 667–689.
- Koopman, S. J., A. Lucas, and M. Scharth (2015). Numerically accelerated importance sampling for nonlinear non-Gaussian state-space models. *Journal of Business & Economic Statistics* 33(1), 114–127.
- Koopman, S. J., N. Shephard, and D. Creal (2009). Testing the assumptions behind importance sampling. *Journal of Econometrics* 149(1), 2–11.
- Koopman, S. J., N. Shephard, and J. A. Doornik (2008). *SsfPack 3.0: Statistical algorithms for models in state space form*. London: Timberlake Consultants Press.
- Liesenfeld, R. and J.-F. Richard (2003). Univariate and multivariate stochastic volatility models: estimation and diagnostics. *Journal of Empirical Finance* 10, 505–531.
- Malik, S. and M. K. Pitt (2011). Particle filters for continuous likelihood evaluation and maximisation. *Journal of Econometrics* 165, 190–209.
- Mertens, E. (2016). Measuring the level and uncertainty of trend inflation. *Review of Economics and Statistics* 98(5), 950–967.
- Mertens, E. and J. M. Nason (2017). Inflation and professional forecast dynamics: An evaluation of stickiness, persistence, and volatility. *CAMA Working Paper, Australian National University*.

- Pitt, M. K. and N. Shephard (1999). Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association* 94(446), 590–599.
- Richard, J.-F. and W. Zhang (2007). Efficient high-dimensional importance sampling. *Journal of Econometrics* 141(2), 1385–1411.
- Sandmann, G. and S. J. Koopman (1998). Estimation of stochastic volatility models via Monte Carlo maximum likelihood. *Journal of Econometrics* 87(2), 271–301.
- Shephard, N. (2015). Martingale unobserved component models. In S. J. Koopman and N. Shephard (Eds.), *Unobserved Components and Time Series Econometrics*, Chapter 10. Oxford University Press.
- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84(3), 653–667.
- Stock, J. H. and M. W. Watson (1998). Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association* 93(441), 349–358.
- Stock, J. H. and M. W. Watson (2007). Why has US inflation become harder to forecast? *Journal of Money, Credit and banking* 39(s1), 3–33.
- Stock, J. H. and M. W. Watson (2008). Phillips curve inflation forecasts. Technical report, National Bureau of Economic Research.
- Stock, J. H. and M. W. Watson (2016). Core inflation and trend inflation. *Review of Economics and Statistics* 98(4), 770–784.