Supplementary Material to: Replicating the Results in "A New Model of Trend Inflation" using Particle Markov Chain Monte Carlo

Nima Nonejad*

University of Rome "Tor Vergata" and CREATES

1 Summary of Contents

This online material contains several sections. We provide more technical details on the implementation of the PMCMC algorithms. We also provide examples that highlight some of the advantages of PMCMC compared to common Gibbs sampling approaches and intuitively justify applying PMCMC in practice. The last section provides additional figures and results regarding the AR-trend-bound model of Chan et al. (2013). Papers cited in this supplementary material are not necessarily listed in the references in the paper itself.

2 Stochastic Volatility in Mean Model

In order to provide an intuitive illustration of our PMCMC techniques, we consider the following stochastic volatility in mean (SVM) model, see also Chan (2015)

$$y_t = \mu + \lambda \exp(h_t) + \varepsilon_t, \quad \varepsilon_t \sim N(0, \exp(h_t)) \tag{2.1}$$

$$h_t = \mu_h + \phi_h (h_{t-1} - \mu_h) + \varepsilon_t^h, \quad \varepsilon_t^h \sim N\left(0, \sigma_h^2\right), \tag{2.2}$$

for $t = 1, ...T$, where y_t denotes the underlying series of interest, $E\left[\varepsilon_t \varepsilon_t^h\right] = 0$ and $|\phi_h| < 1$. The interesting feature of this model is that contrary to the plain SV model, Gibbs sampling estimation of (2.1)-(2.2) is nontrivial since h_t appears in both the conditional mean and the conditional variance. Thus, well-known methods such as Kim et al. (1998) cannot be used to draw $h_{1:T} \sim p(h_{1:T} \mid \theta, y_{1:T})$. Within a Gibbs sampling

*Department of Economics and Finance, University of Rome "Tor Vergata", Via Columbia 2, 00133 Rome, Italy. E-mail: nimanonejad@gmail.com.

scheme, one can typically draw $h_{1:T} \sim p(h_{1:T} \mid \theta, y_{1:T})$ using an AR-MH procedure, see for instance Chan (2015). Conditional on $h_{1:T}$, sampling the volatility parameters, μ_h, ϕ_h and σ_h^2 is relatively easy, see Kim et al. (1998). Finally, we can draw $\mu, \lambda \sim p(\mu, \lambda \mid h_{1:T}, y_{1:T})$ in one-block as regression parameters of a linear Gaussian model. One of the main advantages of PMCMC is that it requires limited design effort from the practitioner's part, especially if one is interested in making changes in a particular model. For instance, compared to the plain SV model, modifying the model to incorporate SVM effects requires minor changes. Specifically, we only need to augment the measurement equation in the particle filter (PF) to include $\lambda \exp(h_t)$. Otherwise, we maintain the same sampling and propagation steps as for the plain SV model[1]. We refer the reader to Creal (2012) for a very good introduction to particle filtering. On the other hand, making the same modification using common Gibbs sampling requires more programming effort, see Chan (2015). In Section 2.1, we show that handling bounded processes is also straightforward using PMCMC. Furthermore, for all of the models that we consider, the integrated likelihood at iteration i, $p\left(y_{1:T} \mid \theta^{(i)}\right)$, is automatically available as byproduct of our estimation procedure. Therefore, we can use it to compute the marginal likelihoods, ML, and compare models using Bayes factors, BF. Below, we briefly outline how to estimate (2.1)-(2.2) using PG-AS and PMMH. We then compare our posterior output with output from the AR-MH procedure of Chan (2015) using simulated data[2].

We collect the model parameters, μ, λ, μ_h, ϕ_h, σ_h^2, in θ, and let i denote PMCMC iteration number i out of the total N iterations. Our PMCMC sampling steps of (2.1)-(2.2) are as follows:

- Particle Gibbs (with ancestor sampling, PG-AS)

In PG, we act as if we are operating within a Gibbs sampling scheme. However, there is one major difference, namely, that we draw $h_{1:T} \sim p(h_{1:T} \mid \theta, y_{1:T})$ using the conditional particle filter. Conditional on $h_{1:T}$, drawing $\theta \sim p(\theta \mid h_{1:T}, y_{1:T})$ can be performed using standard Gibbs techniques. However, PG can suffer from a serious drawback, which is that the underlying mixing can be very poor when there is path degeneracy in the SMC sampler[3]. In this paper, in order to avoid any such issues, we choose to use the particle Gibbs with ancestor sampling (PG-AS) approach of Lindsten et al. (2014). PG-AS is able to alleviate the path degeneracy problem in a very computationally easy way. Specifically, the original PG kernel is modified

[1] Furthermore, we can also easily replace the Gaussian density with a Student-t density in order to account for heavy tails.

[2] In order to estimate (2.1)-(2.2) using AR-MH, we modify the codes available at Joshua Chan's website `http://people.anu.edu.au/joshua.chan/`.

[3] As mentioned in the main text, this problem can also be addressed by adding a backward simulation step to the PG sampler, yielding a method denoted as PG with backward simulation, see Whiteley et al. (2010).

using a so-called "ancestor sampling" step. This way the same effect as backward sampling is achieved, but without the need to run an explicit backward pass. The steps of our PG-AS sampler are as follows:

1. At iteration i, let $h_{1:T}^{(i-1)}$ be a fixed reference trajectory of $h_{1:T}$ sampled at iteration $i-1$ of PG-AS. Sample $h_{1:T}^{(i)} \mid \theta^{(i-1)}, y_{1:T}$ using the conditional particle filter with ancestor sampling, CPF-AS, see Lindsten et al. (2014) for more details[4].

2. Sample $\theta^{(i)} \mid h_{1:T}^{(i)}, y_{1:T}$ using standard Gibbs techniques.

3. Set $i = i+1$ and goto 1.

- Particle marginal Metropolis-Hastings, PMMH

Contrary to PG-AS, the PMMH algorithm directly targets $p(\theta, h_{1:T} \mid y_{1:T})$. Thus, we can sample θ and $h_{1:T}$ all-at-once[5]. Let $\theta^{(i-1)}$ and $p\left(y_{1:T} \mid \theta^{(i-1)}\right)$ denote the parameter vector and likelihood value at iteration $i-1$. Our PMMH algorithm is as follows:

1. At iteration i, sample a candidate, $\theta^* \sim q\left(\theta^* \mid \theta^{(i-1)}\right)$. Run a PF using θ^* and the data, $y_{1:T}$.

2. Sample a candidate, $h_{1:T}^*$ given a realization of the weighted samples, $\{w_T^{(j)}, h_{1:T}^{(j)}\}_{j=1}^M$, where M is the number of particles, see page 276 of Andrieu et al. (2010) for more details. Let $p(y_{1:T} \mid \theta^*)$ denote the corresponding integrated likelihood estimate using θ^*.

3. Let $p(\theta)$ denote the prior density of θ. With probability

$$1 \quad \wedge \quad \frac{p(y_{1:T} \mid \theta^*)\, p(\theta^*)\, q\left(\theta^{(i-1)} \mid \theta^*\right)}{p\left(y_{1:T} \mid \theta^{(i-1)}\right) p\left(\theta^{(i-1)}\right) q\left(\theta^* \mid \theta^{(i-1)}\right)}, \tag{2.3}$$

   set $\theta^{(i)} = \theta^*$, $h_{1:T}^{(i)} = h_{1:T}^*$, and $p\left(y_{1:T} \mid \theta^{(i)}\right) = p(y_{1:T} \mid \theta^*)$; otherwise set $\theta^{(i)} = \theta^{(i-1)}$, $h_{1:T}^{(i)} = h_{1:T}^{(i-1)}$ and $p\left(y_{1:T} \mid \theta^{(i)}\right) = p\left(y_{1:T} \mid \theta^{(i-1)}\right)$.

4. Set $i = i+1$ and goto 1.

[4]Specifically, in PG-AS, we sample a new value for the Mth index variable, $a_t^{(M)}$, where $a_t^{(j)}$ is the jth index variable and M is the number of particles. On the other hand, in the PG algorithm of Andrieu et al. (2010), we set $a_t^{(M)} = M$. Even though this is a small modification, improvements in mixing can be quite considerable, see Lindsten et al. (2014) for more details. Furthermore, it is important to note that we are actually drawing $h_{1:T} \sim p\left(h_{1:T}, a_{1:T}^{(M)} \mid \theta, y_{1:T}\right)$. Thus, from a technical point of view, we are not drawing from the true conditional posterior, $p(h_{1:T} \mid \theta, y_{1:T})$, but from a very close approximation, $p\left(h_{1:T}, a_{1:T}^{(M)} \mid \theta, y_{1:T}\right)$. However, in order to ease the notation burden and avoid unnecessary confusions for the reader, we use the notation $h_{1:T} \sim p(h_{1:T} \mid \theta, y_{1:T})$.

[5]We can also sample θ element-by-element within the PMMH framework, see for instance Flury and Shephard (2011). However, this will increase the computation time drastically as we need to run the particle filter twice for each block.

2.1 Simulation example

We simulate $T = 500$ observations from (2.1)-(2.2). We first generate $h_{1:T}$ through the volatility parameters, and then generate $y_{1:T}$ using μ, λ and $h_{1:T}$. The true value of θ is set to $(0.1, -0.05, 0.5, 0.98, 0.02)'$[6]. We set the number of particles, M, to 100 for PG-AS, 1000 for PMMH and take 20000 draws from the joint posterior, $p(h_{1:T}, \theta \mid y_{1:T})$, after a burn-in of 5000[7].

For the PMMH algorithm, we initially specify the covariance matrix for the MH increments as $\triangle\mu^{(i)} = 0.1\xi_1^{(i)}$, $\triangle\lambda^{(i)} = 0.1\xi_2^{(i)}$, $\triangle\mu_h^{(i)} = 0.1\xi_3^{(i)}$, $\triangle\phi_h^{(i)} = 0.01\xi_4^{(i)}$ and $\triangle\sigma_h^{2(i)} = 0.01\xi_5^{(i)}$, $i = 1, ..., 50$, where $\xi_k \sim N(0,1)$, $k = 1, ..., 5$. Thereafter, we update $\Sigma^{(i)}$ using the covariance of $\{\theta^{(n-1)}\}_{n=1}^{i-1}$, i.e. $\Sigma^{(i)} = d_s\Sigma^{(i-1)} + I\xi$, where we set $d_s = 0.40$. Throughout this paper, we achieve MH acceptance ratios around 30% to 40%, which are also in accordance with Flury and Shephard (2011). We suggest choosing M as follows: A good indication of when we reach a sufficient number of particles (in the sense of achieving a likelihood estimate that is not too jittery) is when the rate with which the MH acceptance ratio increases with M starts to slow down and improvements become only marginal. For instance, in Figure 1, we estimate (2.1)-(2.2) with $M = 10$, 100, 500, 1000 and compare results. Clearly $M = 10$ is not sufficient as the MH acceptance ratio is almost zero. Furthermore, the inefficiency factors of $h_{1:T}$, the autocorrelation functions (ACF) of the posterior draws of $\psi = \left(\mu_h, \phi_h, \sigma_h^2\right)'$ show very high autocorrelation and thus poor mixing. For $M = 100$, we obtain a MH acceptance ratio of 10% and the ACFs of ψ decrease. However, the likelihood function can still be too jittery in some periods, which means that there is a chance that the algorithm gets stuck at a particular point. For instance, in panel (b) of Figure 1, we see that the chain gets stuck for a while towards $N = 8000$. For $M = 500$ and $M = 1000$, we obtain almost identical results. For instance, we obtain a MH acceptance ratio of 34% (38%) for $M = 500$ ($M = 1000$). The inefficiency factors of $h_{1:T}$ are only marginally lower, and the ACFs of ψ show that the Markov chain is very well-mixing[8].

We plot posterior estimates of $\exp(h_t/2)$, $t = 1, ..., T$, using the aforementioned methods in panels (a), (c) and (e) of Figure 2. Obviously, we obtain almost identical posterior results for AR-MH and PMCMC. We also obtain very similar posterior parameter estimates, see Table 1. In panel (e) of Figure 2, we report

[6] We choose the same starting values, θ_0, for all of the samplers. Furthermore, we use Matlab's `setGlobalStream` feature to ensure that we are basically using the same seeding when we take draws for different samplers.

[7] In general, we find that PG-AS works very well for $M \geq 10$. We simply set $M = 100$ to have a decent number of particles, well-mixing posterior draws and reasonable computation time. For instance, compared to $M = 100$, PG-AS estimation of (2.1)-(2.2) using $M = 1000$ takes almost three times as long.

[8] Ideally, we would prefer higher MH acceptance ratios. However, we refer the reader to Table 3 of Flury and Shephard (2011), where even for a linear model in which MH can be performed using exact likelihood through the Kalman filter, the authors obtain MH acceptance ratios of around 40%.

the inefficiency factors of $h_{1:T}$ for each algorithm. Evidently, AR-MH and PG-AS produce similar results, whereas the inefficiency factors of $h_{1:T}$ associated with PMMH are marginally higher. However, this is understandable as in PMMH, we update $(h_{1:T}, \theta)$ all-at-once, whereas for the other methods $h_{1:T}$ and θ are sampled sequentially. In panels (d) and (f) of Figure 2, we report the ACFs of the posterior draws of ϕ_h and σ_h^2 after the burn-in period. Overall, AR-MH and PMCMC are very capable of producing posterior draws of these parameters that are not highly autocorrelated.

3 Unobserved Components Model with SV Effects, US Inflation

PMCMC techniques are also very flexible if we want to simultaneously generate several, possibly bounded latent states from their respective conditional posteriors. We consider the following well-known unobserved components model with SV effects for US quarterly inflation

$$y_t = \tau_t + \varepsilon_t, \quad \varepsilon_t \sim N(0, \exp(h_t)) \tag{3.1}$$

$$\tau_t = \tau_{t-1} + \varepsilon_t^{\tau}, \quad \varepsilon_t^{\tau} \sim N\left(0, \sigma_{\tau}^2\right) \tag{3.2}$$

$$h_t = \mu_h + \phi_h (h_{t-1} - \mu_h) + \varepsilon_t^{h}, \quad \varepsilon_t^{h} \sim N\left(0, \sigma_h^2\right). \tag{3.3}$$

Estimating this specification is very easy using PMCMC. We can simply let $\gamma_t = (\tau_t, h_t)'$, modify the propagation and resampling steps of the particle filter such that we simulate particles for γ_t, and then generate $\gamma_{1:T} \sim p(\gamma_{1:T} \mid \theta, y_{1:T})$. Below, we report results for (3.1)-(3.3) using PG-AS, PMMH and compare our results with a Gibbs sampling approach. Specifically, for the latter approach, we generate $\tau_{1:T} \sim p(\tau_{1:T} \mid \theta, h_{1:T}, y_{1:T})$ and $h_{1:T} \sim p(h_{1:T} \mid \theta, \tau_{1:T}, y_{1:T})$ sequentially using the precision sampler of Chan (2013).

Results are reported in Figure 3 and Table 2. Similar to the SVM model, we obtain almost identical posterior estimates of the latent processes and the model parameters. However, contrary to typical Gibbs sampling procedures such as Chan et al. (2013), bounding one or all of the processes requires minor coding effort. For instance, assume that we want to bound $\tau_{1:T}$ between a and b, i.e. $\tau_t \in (a,b)$, $t = 1,...,T$. For this version, all we need to do is to change the propagation step of the particle filter such that we simulate particles from a truncated Normal distribution, i.e. $\tau_t^{(j)} \sim TN(a, b, \tau_{t-1}^{\left(a_t^{(j)}\right)}, \sigma_{\tau}^2)$, $j = 1,...,M$. We illustrate this point graphically in panels (c) and (d) of Figure 4. In panel (c), we plot $M = 20$ particles for $t = 1,...,20$ time periods for a model where the unobserved process is unbounded. In panel (d), we consider the same

model, however, we bound the unobserved process between $a = 0$ and $b = 4$. Obviously, in the latter case, all of the generated particles are within the prespecified interval, which makes generating the unobserved process very easy.

Finally, as previously mentioned, one major advantage of PMCMC is that $p\left(y_{1:T} \mid \theta^{(i)}\right)$, $i = 1,...,N$, is easily available as a byproduct of the particle filter. More importantly, Del Moral (2004) shows that this output is an unbiased estimate of $p(y_{1:T} \mid \theta)$, see also Flury and Shephard (2011). Therefore, we can use these quantities to compute different statistics. In this paper. we use $p\left(y_{1:T} \mid \theta^{(i)}\right)$ and calculate the marginal likelihood (ML) using the method of Gelfand-Dey (1994), see also Koop (2004). We can also compute ML using the the conditional likelihood, $p(y_{1:T} \mid \tau_{1:T}, h_{1:T}, \theta)$. However, we refer the reader to Chan and Grant (2015) for more details on the pitfalls of estimating ML based on the conditional likelihood.

For (3.1)-(3.3), we obtain a ML value of -551.71. We then compare this model with versions of (3.1)-(3.3), where $a = 0$ and $b = 5, 6, ..., 10$, i.e. we increase b by 1. For each specification, we estimate the model and compute ML. We report posterior means of $\tau_{1:T}$ for these versions in panels (e) and (f) Figure 4. According to ML, we must increase b to 9 before we obtain the highest ML value. For instance, for the version with $a = 0$, $b = 5$, we obtain a ML of -563.64. As we increase b, we obtain higher ML values. For the version with $a = 0$, $b = 9$, we obtain a ML of -551.06, which slightly outperforms (3.1)-(3.3). For $a = 0$, $b = 10$, we obtain a ML of -551.15.

4 AR-trend-bound with a and b Estimated

Estimating a and b in the AR-trend-bound model is also easy using PMCMC. For PG-AS, we can simply implement the Griddy Gibbs procedure to draw $a \sim p\left(a \mid \tau_{1:T}, \sigma_\tau^2\right)$ and $b \sim p\left(b \mid \tau_{1:T}, \sigma_\tau^2\right)$ from their respective conditional posteriors. For PMMH, we follow Chan et al. (2013), let $p(a) \sim Uniform(0, 1.5)$, $p(b) \sim Uniform(2.5, 5)$, and augment the parameter vector to include a and b. We then proceed to sample θ all-at-once. At each iteration, we simply use Matlab's `while` command to ensure that the restrictions in θ are satisfied. We estimate the posterior mean of a at 1.43 and b at 4.20, with posterior standard deviations of 0.52 and 0.38, respectively, see panels (e) and (f) of Figure 7. Compared to AR-trend-bound, $a = 0$ and $b = 5$, we estimate $\tau_{1:T}$ at a lower rate, whereas posterior estimates of $\rho_{1:T}$ and $h_{1:T}$ are similar to the version where $a = 0$ and $b = 5$. Finally, by comparing the marginal likelihoods, we see that this version marginally outperforms AR-trend-bound, $a = 0$ and $b = 5$, see Table 1 in the main text.

References

[1] Chan, J. 2015. "The Stochastic Volatility in Mean Model with Time-Varying Parameters: An Application to Inflation Modeling". *Journal of Business and Economic Statistics* forthcoming.

[2] Chan, J., and A. L. Grant. 2015. "Pitfalls of estimating the marginal likelihood using the modified harmonic mean". *Economic Letters 131(1)*: 29-33.

[3] Creal, D. 2012. "A Survey of Sequential Monte Carlo Methods for Economics and Finance". *Econometric Reviews 31(3)*: 245-296.

[4] Del Moral, P. 2004. "Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications". Springer, New York.

[5] Flury, T., and N. Shephard. 2011. "Bayesian inference based only on simulated likelihood: particle filter analysis of dynamic economic models". *Econometric Theory 27(5)*: 933-956.

[6] Kim, S., N. Shephard, and S. Chib. 1998. "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models" . *Review of Economic Studies 65(3)*: 361-393.

[7] Koop, G. 2003. "Bayesian Econometrics". John Wiley & Sons Ltd.

[8] Whiteley, N., C. Andrieu, and A. Doucet. 2010. "Efficient Bayesian Inference for Switching State-Space Models using Particle Markov chain Monte Carlo methods". Bristol Statistics Research Report 10:04.

Table 1: Estimation results, stochastic volatility in mean model

Parameter	true	mean	std. dev	16%-tile	84%-tile	R_B	Geweke
Method: AR-MH							
μ	0.10	0.1297	0.0782	0.0528	0.2076	1.92	-1.86
λ	-0.05	-0.0953	0.0686	-0.1631	-0.0277	1.64	1.48
μ_h	0.50	-0.0810	0.6038	-0.4724	0.3238	1.24	0.89
ϕ_h	0.98	0.9819	0.0112	0.9715	0.9929	13.75	0.37
σ_h^2	0.02	0.0219	0.0077	0.0149	0.0286	39.78	-0.82
Method: PG-AS							
μ	0.10	0.1299	0.0777	0.0529	0.2068	1.80	-0.91
λ	-0.05	-0.0950	0.0679	-0.1620	-0.0279	1.60	1.63
μ_h	0.50	-0.0680	0.5701	-0.4516	0.3214	1.32	-0.88
ϕ_h	0.98	0.9812	0.0113	0.9703	0.9924	13.84	0.19
σ_h^2	0.02	0.0227	0.0080	0.0157	0.0296	39.64	0.71
Method: PMMH							
μ	0.10	0.1269	0.0776	0.0499	0.2056	22.27	-1.53
λ	-0.05	-0.0945	0.0687	-0.1629	-0.0263	23.18	1.48
μ_h	0.50	-0.0871	0.5365	-0.4817	0.3123	33.40	1.44
ϕ_h	0.98	0.9815	0.0106	0.9712	0.9920	27.74	-0.32
σ_h^2	0.02	0.0222	0.0071	0.0155	0.0288	25.55	0.22

This table reports estimation results of (2.1)-(2.2) using the mentioned methods. R_B: Inefficiency factors using a bandwidth, B, of 100. Geweke: Geweke's convergence statistics, see Koop (2004).

Table 2: Estimation results, unobserved components model with SV effects, US inflation

Parameter	mean	std. dev	16%-tile	84%-tile	R_B	Geweke
Method: AR-MH						
μ_h	1.0435	1.1360	0.1758	1.9496	1.98	0.02
ϕ_h	0.9766	0.0161	0.9614	0.9920	14.60	0.83
σ_h^2	0.1089	0.0424	0.0706	0.1467	32.93	-1.33
σ_τ^2	0.2329	0.1183	0.1191	0.3539	49.13	1.57
Method: PG-AS						
μ_h	1.0594	1.1370	0.2009	1.9649	1.76	-1.09
ϕ_h	0.9768	0.0160	0.9617	0.9920	13.70	-0.61
σ_h^2	0.1055	0.0411	0.0687	0.1415	31.23	0.97
σ_τ^2	0.2285	0.1119	0.1200	0.3410	45.14	1.10
Method: PMMH						
μ_h	1.1054	1.2023	0.2416	2.0265	41.28	1.55
ϕ_h	0.9768	0.0151	0.9628	0.9911	37.78	1.19
σ_h^2	0.1028	0.0384	0.0673	0.1369	34.02	-1.13
σ_τ^2	0.2711	0.1237	0.1464	0.4014	33.91	0.10

This table reports estimation results of (3.1)-(3.3) using the mentioned methods. R_B: Inefficiency factors using a bandwidth, B, of 100. Geweke: Geweke's convergence statistics, see Koop (2004).

Figure 1: PMMH estimation results, stochastic volatility in mean model

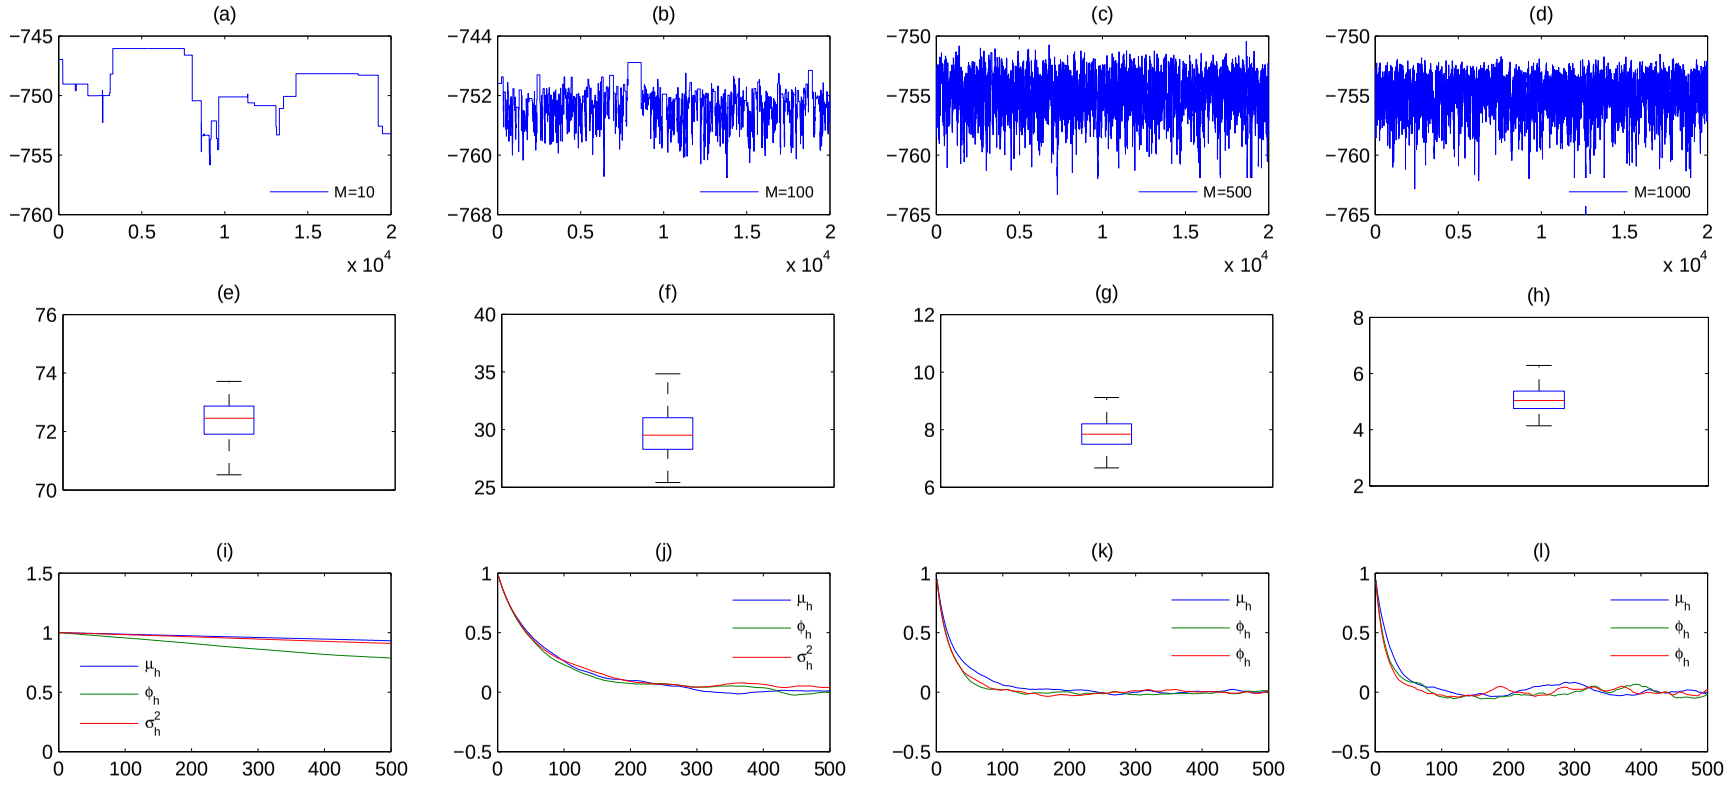

Panels (a)-(d): Likelihood against iteration for $M = 10$, 100, 500 and 1000. Panels (e)-(h): Inefficiency factors of $h_{1:T}$ for $M = 10$, 100, 500 and 1000. Panels (i)-(l): Autocorrelation functions of the posterior draws of the volatility parameters for $M = 10$, 100, 500 and 1000.

Figure 2: Estimation results, stochastic volatility in mean model

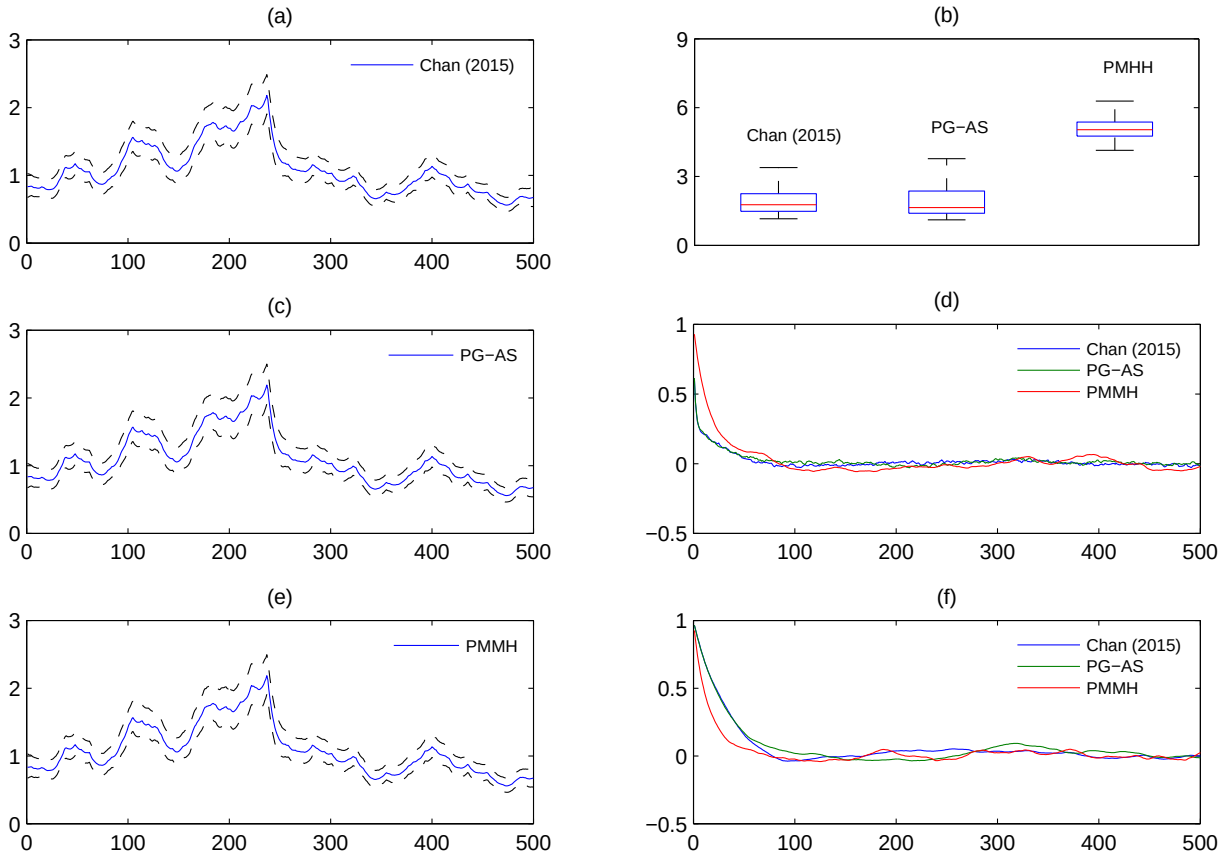

Panels (a), (c) and (e): Posterior estimates of $\exp(h_{1:T}/2)$, $t = 1,...,T$, using the mentioned methods. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles. Panel (b): Inefficiency factors of $h_{1:T}$ using the mentioned methods. Panel (d): Autocorrelation functions of the posterior draws of ϕ_h using the mentioned methods. Panel (f): Autocorrelation functions of the posterior draws of σ_h^2 using the mentioned methods.

Figure 3: Estimation results, unobserved components model with SV effects, US inflation

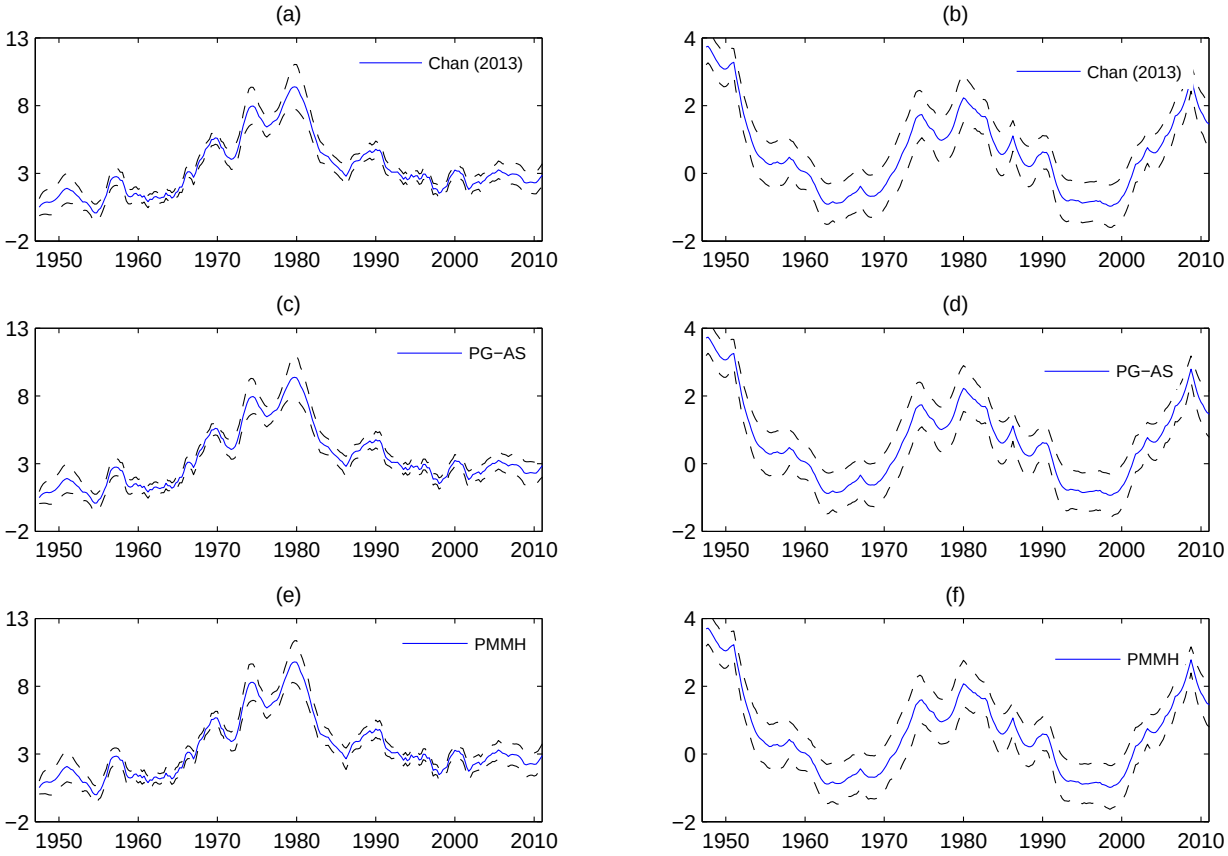

Panels (a), (c) and (e): Posterior estimates of $\tau_{1:T}$, $t = 1, ..., T$, using the mentioned methods. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles. Panels (b), (d) and (f): Posterior estimates of $h_{1:T}$, $t = 1, ..., T$, using the mentioned methods. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles.

Figure 4: Estimation results, unobserved components model with SV effects, US inflation

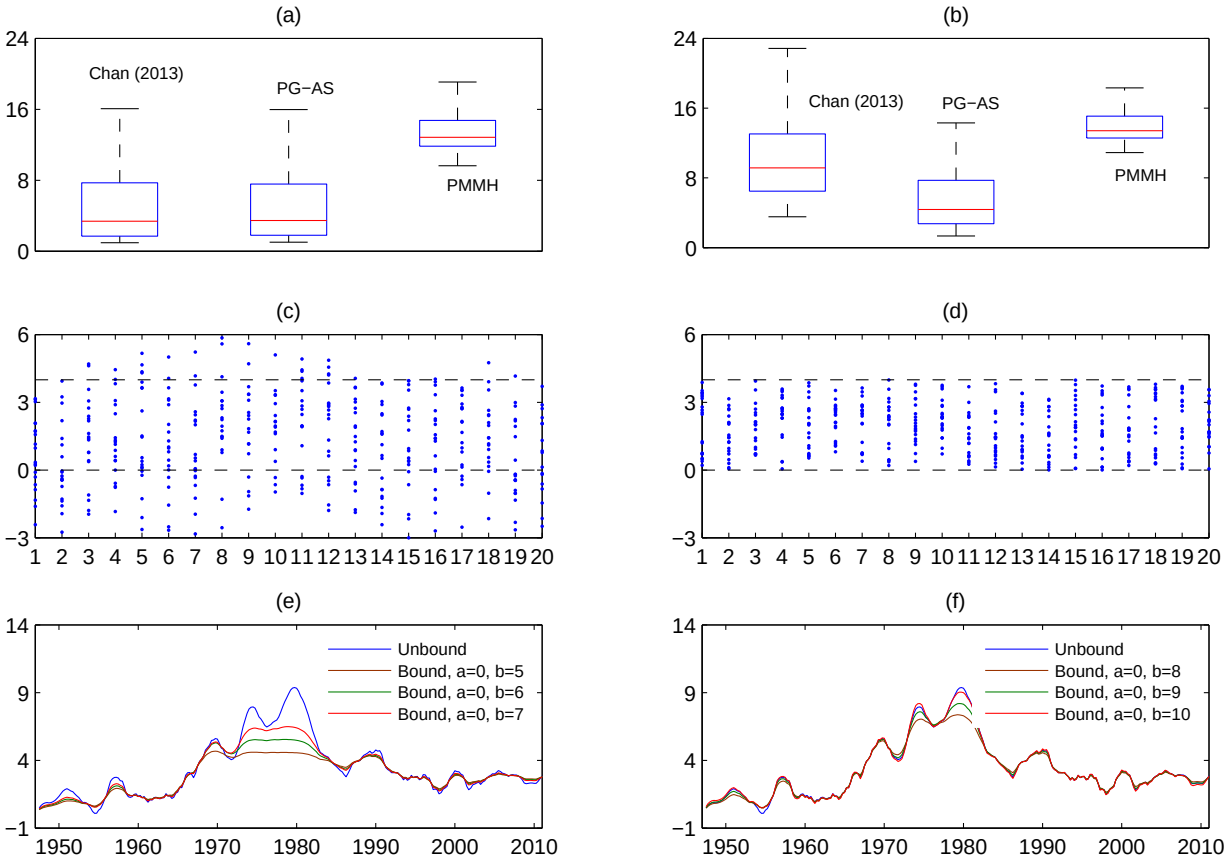

Panels (a) and (b): Inefficiency factors of $\tau_{1:T}$ and $h_{1:T}$ using the corresponding methods. Panels (c) and (d): Illustration of the particle filter, unbounded situation (left), bounded situation (right). The dotted lines denote the respective lower and upper bounds. Panels (e) and (f): Posterior mean estimates of $\tau_{1:T}$, $t = 1, ..., T$, for the unobserved components model with SV effects with and without bounds on $\tau_{1:T}$.

Figure 5: Simulation against iteration and ACF of the model parameters, AR-trend-bound model, $a = 0$ and $b = 5$, US inflation

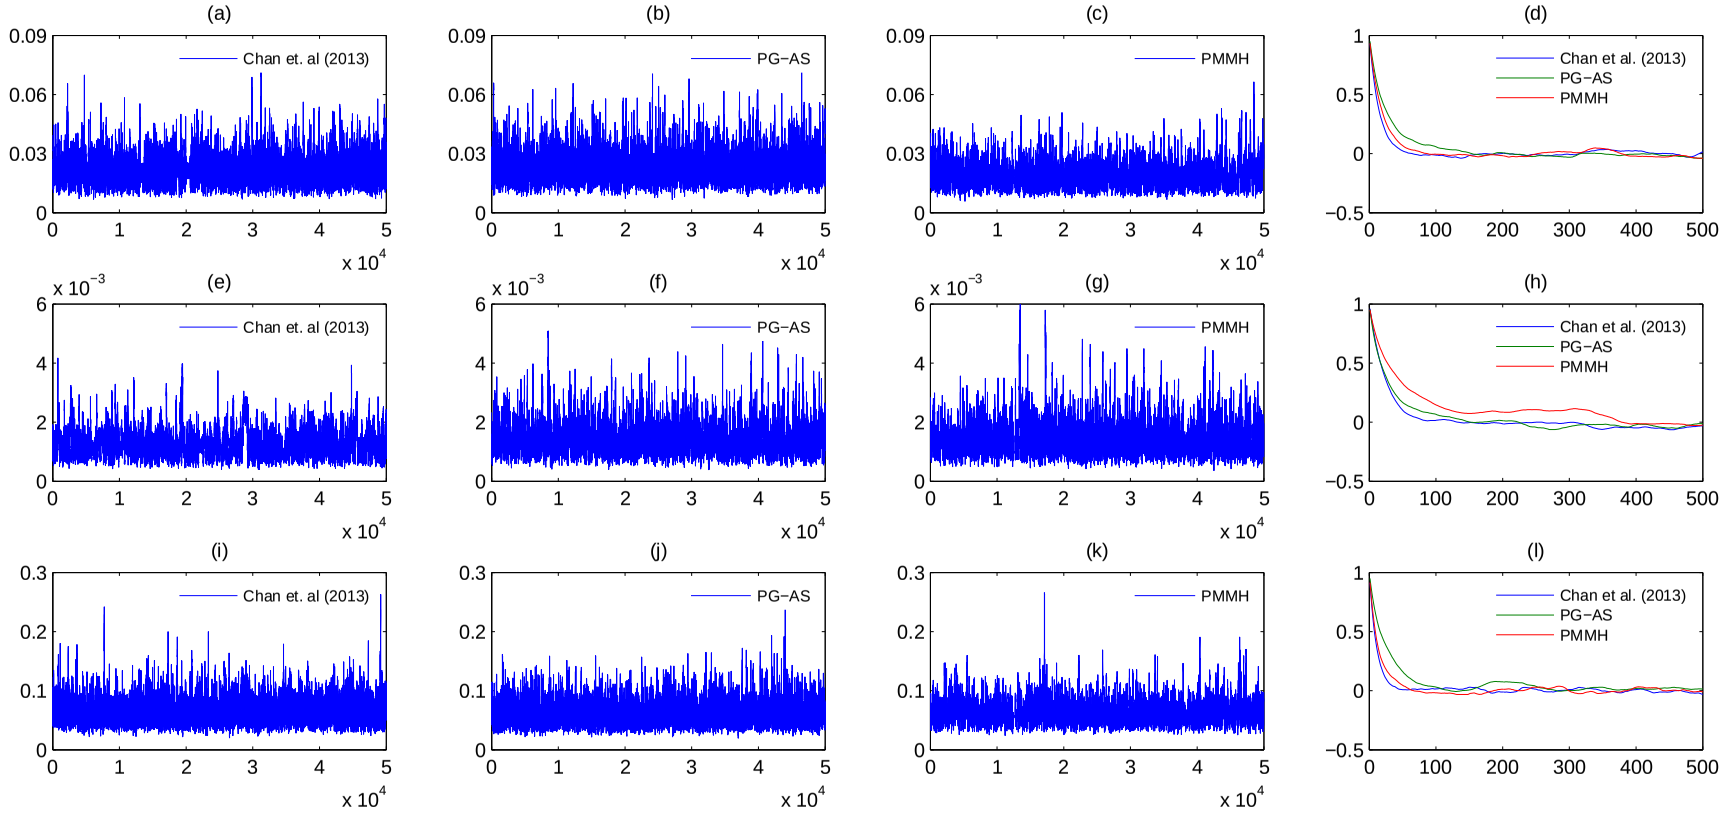

Panels (a), (b) and (c): Simulation against iteration, σ_τ^2, for each method. Panels (e), (f) and (g): Simulation against iteration, σ_ρ^2, for each method. Panels (i), (j) and (k): Simulation against iteration, σ_h^2, for each method. Panels (d), (h) and (l): Autocorrelation functions of of the posterior draws of σ_τ^2, σ_ρ^2 and σ_h^2 for each method.

Figure 6: Estimation results, AR-trend model, $\pi_t = \tau_t + \rho_t \pi_{t-1} + \varepsilon_t$, US inflation

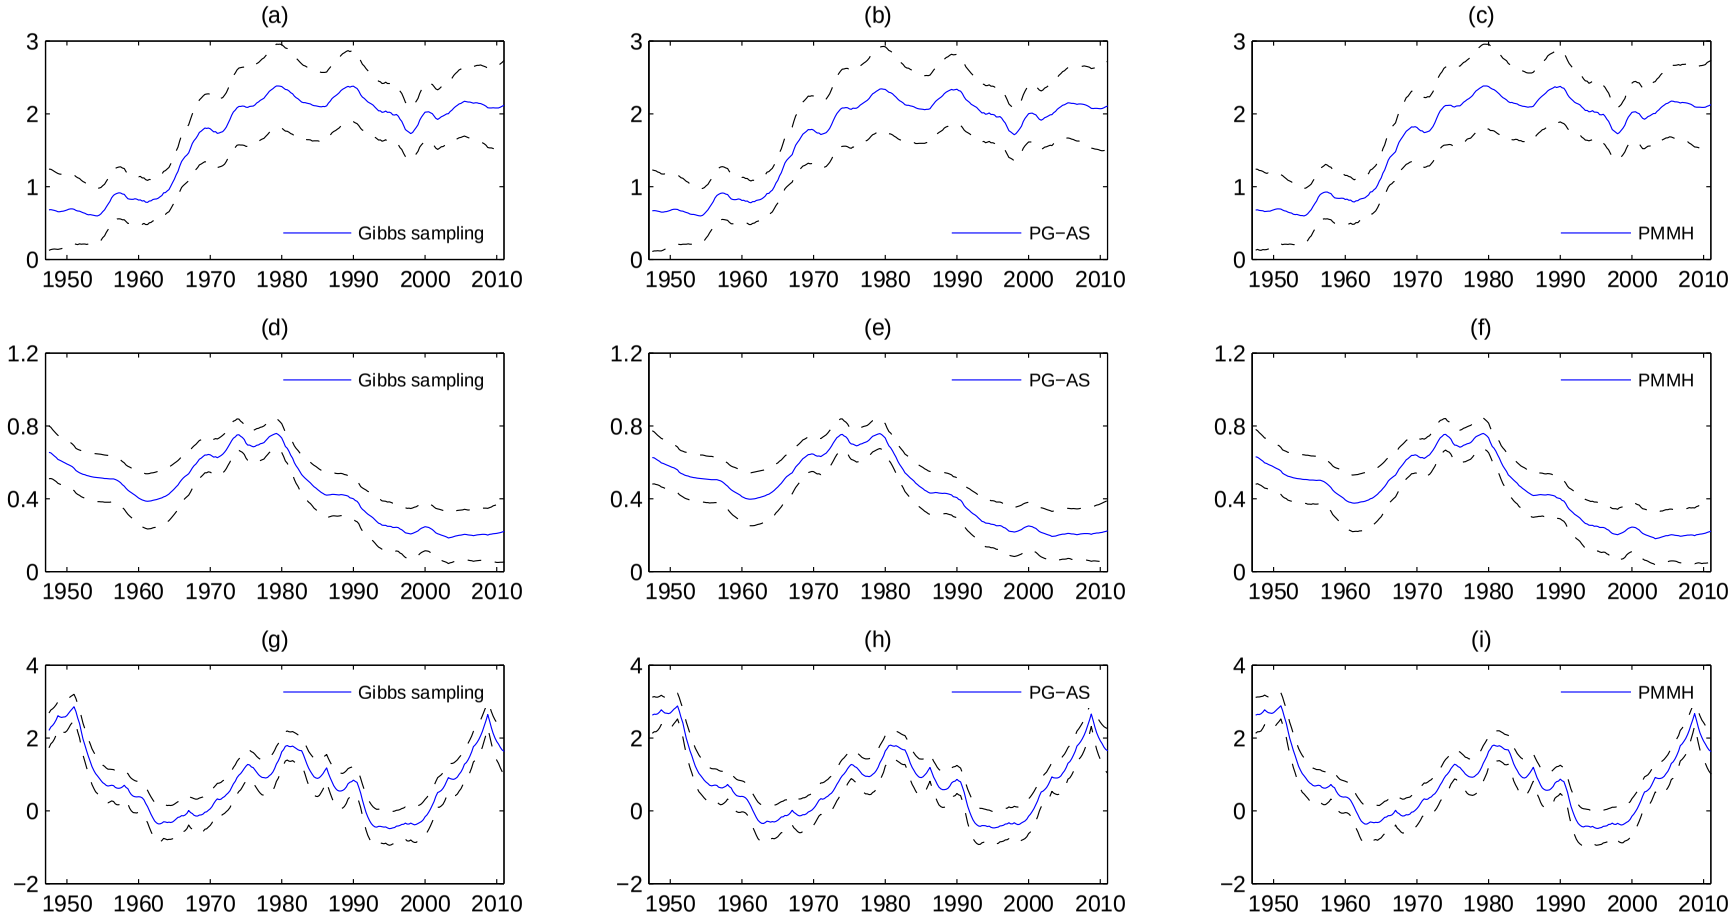

Panels (a), (b) and (c): Posterior estimates of $\tau_{1:T}$ using the mentioned methods. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles. Panels (d), (e) and (f): Posterior estimates of $\rho_{1:T}$ using the mentioned methods. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles. Panels (g), (h) and (i): Posterior estimates of $h_{1:T}$ using the mentioned methods. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles.

Figure 7: Estimation results, AR-trend model, a and b estimated, US inflation

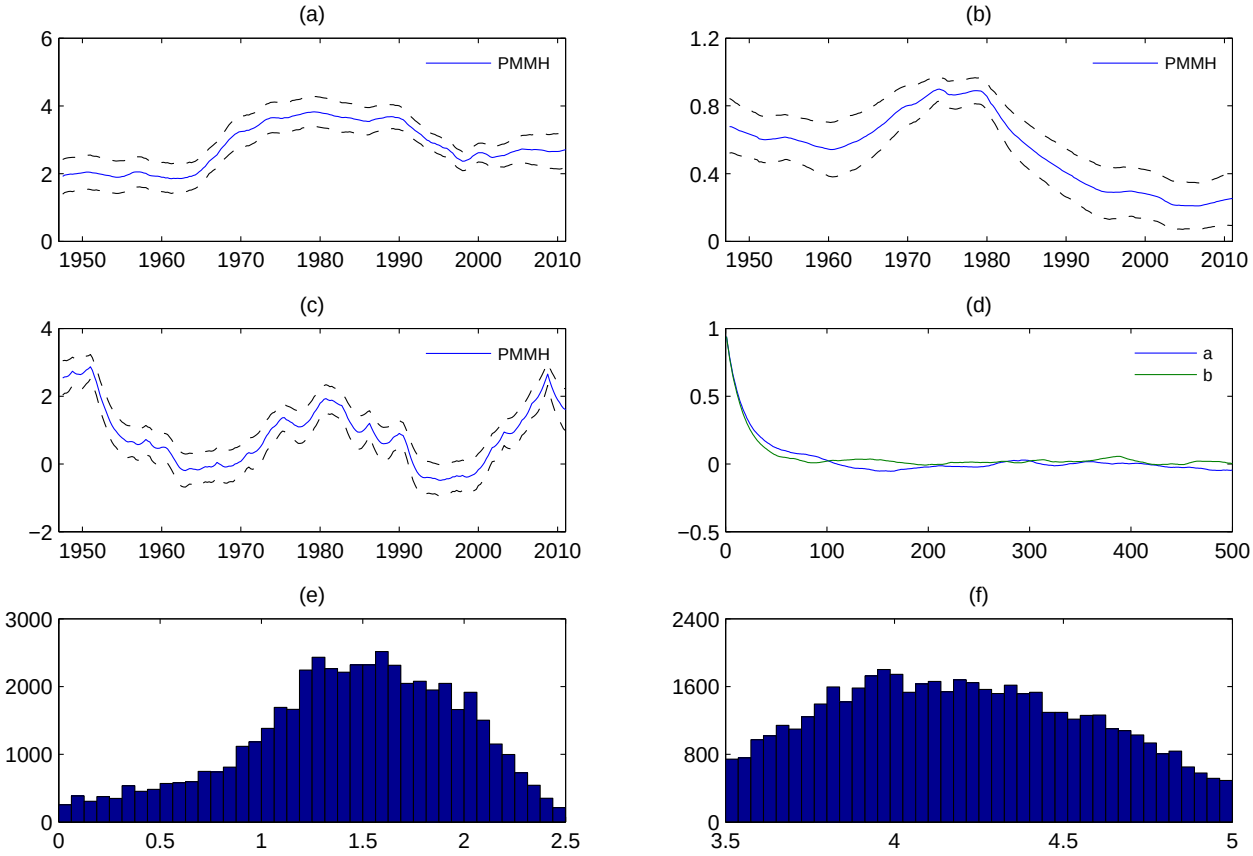

Panels (a), (b) and (c): Posterior estimates of $\tau_{1:T}$, $\rho_{1:T}$ and $h_{1:T}$ using PMMH. The blue lines indicate the posterior mean estimates. The black dotted lines are the 16 and 84 posterior percentiles. Panel (d): Autocorrelation functions of the posterior draws of a and b. Panels (e) and (f): Histograms of the posterior draws of a and b.