

Technical Appendix to: A New Model of Trend Inflation*

Joshua C.C. Chan Gary Koop
Australian National University University of Strathclyde

Simon M. Potter
Federal Reserve Bank of New York

September 2012

Appendix A: MCMC Algorithm

This appendix develops a posterior simulation algorithm for **AR-trend-bound**: the bounded inflation model given in (5). The other models are restricted special cases of this model and, thus, the MCMC algorithm is restricted in the obvious manner in each case. The one exception of this is the UC-SV model of Stock and Watson (2007), which we label **Trend-SV** in the paper. This involves one extra state equation for the stochastic volatility in the inflation defining trend inflation. This is drawn using the stochastic volatility described in this appendix.

Except for the parameters a and b , the prior is described in Section 3.2. The priors for a and b are assumed to be uniform on the intervals (\underline{a}, \bar{a}) and (\underline{b}, \bar{b}) respectively, where $\underline{a} = 0$, $\bar{a} = 1.5$, $\underline{b} = 3.5$ and $\bar{b} = 5$.

The MCMC algorithm sequentially draw from (we suppress the dependence on y_0):

1. $p(\tau | y, \rho, h, \theta)$;
2. $p(\rho | y, \tau, h, \theta)$;

*The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York or the Federal Reserve System. Gary Koop is a Fellow of the Rimini Center for Economic Analysis. This research was supported by the ESRC under grant RES-062-23-2646.

3. $p(h | y, \tau, \rho, \theta)$;
4. $p(a | y, \tau, \rho, h, \sigma_\tau^2, \sigma_\rho^2, \sigma_h^2, b)$;
5. $p(b | y, \tau, \rho, h, \sigma_\tau^2, \sigma_\rho^2, \sigma_h^2, a)$;
6. $p(\sigma_\tau^2, \sigma_\rho^2, \sigma_h^2 | y, \tau, \rho, h, a, b) =$
 $p(\sigma_\tau^2 | y, \tau, \rho, h, a, b)p(\sigma_\rho^2 | y, \tau, \rho, h, a, b)p(\sigma_h^2 | y, \tau, \rho, h, a, b)$.

In order to derive the above conditional densities, we first rewrite the measurement equation in (5) as

$$Ky = \mu_0 + K\tau + \varepsilon, \quad \varepsilon \sim N(0, \Omega_y),$$

where $\Omega_y = \text{diag}(e^{h_1}, \dots, e^{h_T})$ and

$$\mu_0 = \begin{pmatrix} \rho_1 y_0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad K = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -\rho_2 & 1 & 0 & \cdots & 0 \\ 0 & -\rho_3 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \cdots & -\rho_T & 1 \end{pmatrix}.$$

Since $|K| = 1$ for any ρ , K is invertible. Therefore, we have

$$(y | \rho, h, y_0) \sim N(K^{-1}\mu_0 + \tau, (K'\Omega_y^{-1}K)^{-1}),$$

i.e.,

$$\log p(y | \rho, h, y_0) \propto -\frac{1}{2}\iota_T' h - \frac{1}{2}(y - K^{-1}\mu_0 - \tau)' K'\Omega_y^{-1}K(y - K^{-1}\mu_0 - \tau), \quad (1)$$

where ι_T is a $T \times 1$ column of ones. It is important to note that both K and Ω_y^{-1} are sparse, i.e., they contain only a small number of non-zero elements. As such, computational benefits of working with sparse matrix algorithms can be exploited in this setting. In fact, (1) can be evaluated quickly without the need of inverting any large matrices (i.e., one needs not compute K^{-1}). We refer the readers to Chan and Jeliazkov (2009) for details. Similarly, we can write

$$H\tau = \varepsilon^\tau,$$

where

$$H = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \cdots & -1 & 1 \end{pmatrix}.$$

That is, the prior density for τ is given by

$$\log p(\tau | \sigma_\tau^2, a, b) \propto -\frac{1}{2}\tau' H' \Omega_\tau^{-1} H \tau + g_\tau(\tau, a, b, \sigma_\tau^2), \quad (2)$$

where $a < \tau_t < b$ for $t = 1, \dots, T$, $\Omega_\tau = \text{diag}(\omega_\tau^2, \sigma_\tau^2, \dots, \sigma_\tau^2)$ and

$$\begin{aligned} g_\tau(\tau, a, b, \sigma_\tau^2) = & -\log \left(\Phi \left(\frac{b}{\omega_\tau} \right) - \Phi \left(\frac{a}{\omega_\tau} \right) \right) \\ & - \sum_{t=2}^T \log \left(\Phi \left(\frac{b - \tau_{t-1}}{\sigma_\tau} \right) - \Phi \left(\frac{a - \tau_{t-1}}{\sigma_\tau} \right) \right). \end{aligned}$$

Combining the previous prior and likelihood for τ , we obtain the log conditional density $\log p(\tau | y, \rho, h, \theta)$ as follows:

$$\begin{aligned} \log p(\tau | y, \rho, h, \theta) \propto & -\frac{1}{2}(y - K^{-1}\mu_0 - \tau)' K' \Omega_y^{-1} K (y - K^{-1}\mu_0 - \tau) \\ & - \frac{1}{2}\tau' H' \Omega_\tau^{-1} H \tau + g_\tau(\tau, a, b, \sigma_\tau^2) \\ \propto & -\frac{1}{2}(\tau - \hat{\tau})' D_\tau^{-1} (\tau - \hat{\tau}) + g_\tau(\tau, a, b, \sigma_\tau^2), \end{aligned}$$

where $a < \tau_t < b$ for $t = 1, \dots, T$, and

$$D_\tau = (H' \Omega_\tau^{-1} H + K' \Omega_y^{-1} K)^{-1}, \quad \hat{\tau} = D_\tau K' \Omega_y^{-1} K (y - K^{-1}\mu_0).$$

Since the conditional density is non-standard, we sample τ via an independence-chain Metropolis-Hasting (MH) step. Specifically, candidate draws are first obtained from the $N(\hat{\tau}, D_\tau)$ distribution with the precision-based algorithm in Chan and Jeliazkov (2009), and they are accepted or rejected via an acceptance-rejection Metropolis-Hasting (ARMH) step.

To draw from $p(\rho | y, \tau, h, \theta)$, rewrite the measurement equation as

$$y^* = X\rho + \varepsilon,$$

where $X = \text{diag}(y_0^*, \dots, y_{T-1}^*)$, $y^* = (y_1^*, \dots, y_T^*)'$ and $y_t^* = y_t - \tau_t$. From the state equation we also have

$$H\rho = \varepsilon^\rho.$$

Therefore, the log conditional density $\log p(\rho | y, \tau, h, \theta)$ is given by:

$$\log p(\rho | y, \tau, h, \theta) \propto -\frac{1}{2}(\rho - \hat{\rho})' D_\rho^{-1} (\rho - \hat{\rho}) + g_\rho(\rho, \sigma_\rho^2),$$

where $0 < \rho_t < 1$ for $t = 1, \dots, T$, and

$$g_\rho(\rho, \sigma_\rho^2) = - \sum_{t=2}^T \log \left(\Phi \left(\frac{1 - \rho_{t-1}}{\sigma_\rho} \right) - \Phi \left(\frac{-\rho_{t-1}}{\sigma_\rho} \right) \right),$$

$$D_\rho = (H' \Omega_\rho^{-1} H + X' \Omega_y^{-1} X)^{-1}, \quad \hat{\rho} = D_\rho X' \Omega_y^{-1} y^*, \quad \Omega_\rho = \text{diag}(\omega_\rho^2, \sigma_\rho^2, \dots, \sigma_\rho^2).$$

As before, we implement an ARMH step with approximating density $N(\hat{\rho}, D_\rho)$. For $p(h | y, \tau, \rho, \theta)$, we directly use the algorithm in Chan and Strachan (2012).

To draw from the bounds of the interval trend inflation is restricted to lie in, note that the log conditional densities for a and b are given by

$$\begin{aligned} \log p(a | y, \tau, \rho, h, \sigma_\tau^2, \sigma_\rho^2, \sigma_h^2, b) &\propto g_\tau(\tau, a, b, \sigma_\tau^2) \\ \log p(b | y, \tau, \rho, h, \sigma_\tau^2, \sigma_\rho^2, \sigma_h^2, a) &\propto g_\tau(\tau, a, b, \sigma_\tau^2) \end{aligned}$$

with supports $\underline{a} < a < \min\{\bar{a}, \min\{\tau_t\}\}$ and $\max\{b, \max\{\tau_t\}\} < b < \bar{b}$, respectively. Since each conditional density is one-dimensional with bounded support, draws from each density can be obtained via a Griddy-Gibbs step (using a uniform grid with 300 grid points, accurate up to at least 2 decimal places).

To draw from the error variances, note that $p(\sigma_\tau^2, \sigma_\rho^2, \sigma_h^2 | y, \tau, \rho, h, a, b)$ is the product of three densities. Hence, we can sample σ_τ^2 , σ_ρ^2 and σ_h^2 sequentially without affecting the efficiency of the sampler. The log conditional density $\log p(\sigma_\tau^2 | y, \tau, \rho, h, a, b)$ is given by

$$\begin{aligned} \log p(\sigma_\tau^2 | y, \tau, \rho, h, a, b) &\propto - (\underline{\nu}_\tau + 1) \log \sigma_\tau^2 - \frac{\underline{S}_\tau}{\sigma_\tau^2} - \frac{T-1}{2} \log \sigma_\tau^2 \\ &\quad - \frac{1}{2\sigma_\tau^2} \sum_{t=2}^T (\tau_t - \tau_{t-1})^2 + g_\tau(\tau, a, b, \sigma_\tau^2), \end{aligned}$$

which is a non-standard density. We therefore implement an MH step with the proposal density

$$IG \left(\underline{\nu}_\tau + \frac{T-1}{2}, \underline{S}_\tau + \frac{1}{2} \sum_{t=2}^T (\tau_t - \tau_{t-1})^2 \right).$$

Similarly, the log conditional density $\log p(\sigma_\rho^2 | y, \tau, \rho, h, a, b)$ is given by

$$\begin{aligned} \log p(\sigma_\rho^2 | y, \tau, \rho, h, a, b) &\propto - (\underline{\nu}_\rho + 1) \log \sigma_\rho^2 - \frac{\underline{S}_\rho}{\sigma_\rho^2} - \frac{T-1}{2} \log \sigma_\rho^2 \\ &\quad - \frac{1}{2\sigma_\rho^2} \sum_{t=2}^T (\rho_t - \rho_{t-1})^2 + g_\rho(\rho, \sigma_\rho^2). \end{aligned}$$

Again, a draw from $p(\sigma_\rho^2 | y, \tau, \rho, h, a, b)$ is obtained via an MH step with the proposal density

$$IG \left(\underline{\nu}_\rho + \frac{T-1}{2}, \underline{S}_\rho + \frac{1}{2} \sum_{t=2}^T (\rho_t - \rho_{t-1})^2 \right).$$

Finally, $p(\sigma_h^2 | y, \tau, \rho, h, a, b)$ is a standard inverse-Gamma density

$$IG \left(\underline{\nu}_h + \frac{T-1}{2}, \underline{S}_h + \frac{1}{2} \sum_{t=2}^T (h_t - h_{t-1})^2 \right).$$

Appendix B: Prior Sensitivity Analysis

This appendix presents results from a small prior sensitivity analysis where we estimate the **AR-trend-bound** model with the following three sets of priors:

- P1: $a \sim U(0, 1.5)$ and $b \sim U(3.5, 5)$;
- P2: $a \sim U(0, 2.5)$ and $b \sim U(2.5, 10)$;
- P3: $a \sim TN(0, 2.5; 1, 1)$ and $b \sim TN(2.5, 10; 3.5, 1)$.

The first of these is the benchmark prior used in the paper. The second is a much more noninformative priors which allow for the bounds on trend inflation to be very wide. The third investigates sensitivity to choice of functional form. It can be seen that estimates of trend inflation change only slightly, and estimates of the other latent states are virtually the same for the three priors.




Figure 1: Estimates for τ_t , ρ_t and h_t under various priors.

References

Chan, J. and Jeliazkov, I. (2009). “Efficient simulation and integrated likelihood estimation in state space models,” *International Journal of Mathematical Modelling and Numerical Optimisation*, 1, 101-120.

Chan, J. and Strachan, R. (2012). “Estimation in non-linear non-Gaussian state-space models with precision-based methods,” CAMA Working Paper Series 2012-08, Australian National University, Canberra.

Stock, J. and Watson, M. (2007). “Why has U.S. inflation become harder to forecast?,” *Journal of Money, Credit and Banking* 39, 3-33.