

STRATHCLYDE

DISCUSSION PAPERS IN ECONOMICS



A NEW MODEL OF TREND INFLATION

BY

**JOSHUA CHAN, GARY KOOP AND
SIMON POTTER**

No. 12-02

**DEPARTMENT OF ECONOMICS
UNIVERSITY OF STRATHCLYDE
GLASGOW**

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

February 13, 2012

Abstract: This paper introduces a new model of trend (or underlying) inflation. In contrast to many earlier approaches, which allow for trend inflation to evolve according to a random walk, ours is a bounded model which ensures that trend inflation is constrained to lie in an interval. The bounds of this interval can either be fixed or estimated from the data. Our model also allows for a time-varying degree of persistence in the transitory component of inflation. The bounds placed on trend inflation mean that standard econometric methods for estimating linear Gaussian state space models cannot be used and we develop a posterior simulation algorithm for estimating the bounded trend inflation model. In an empirical exercise with CPI inflation we find the model to work well, yielding more sensible measures of trend inflation and forecasting better than popular alternatives such as the unobserved components stochastic volatility model.

Keywords: Constrained inflation, non-linear state space model, underlying inflation, inflation targeting, inflation forecasting, Bayesian

*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.

1 Introduction

There are numerous studies analyzing the behavior of inflation over the last few decades. There seems to be agreement that the persistence and volatility of inflation have changed over time (see, among many others, Cogley and Sargent, 2005, Stock and Watson, 2007, 2010, Koop and Potter, 2007 and Cogley, Primiceri and Sargent, 2010). This literature has highlighted the importance of the appropriate econometric modeling of the inflation process.

One key finding of many researchers is that the persistence of inflation was reduced after around 1990. This finding may be partially explained by the fact that many central banks, reacting to the great inflation of the 1970s, introduced some type of inflation targeting regime. In such regimes the central bank either announces a point target or a target range for the rate of inflation so as to fix long-term inflation expectations. Other central banks do not have official inflation targeting regimes but behave in similar ways to inflation targeting central banks, particularly in their focus on long term inflation expectations. Successful monetary policy that stabilizes long run inflation expectations is implicitly affecting trend inflation and the persistence of inflation.

This has led to a great policy interest in inflation expectations and measures of trend (or underlying) inflation constructed using time series methods.¹ A large literature has emerged on estimating trend inflation (see Clark and Doh, 2011 as an example of a recent paper which discusses various approaches to modeling trend inflation and surveys much of the related literature). Much of this literature models trend inflation as a driftless random walk either in a univariate times series model or as an assumption embedded in a multivariate time series model (e.g., among many others, Smets and Wouters, 2003, Cogley and Sargent, 2005, Ireland, 2007, Stock and Watson, 2007, Cogley and Sbordone, 2008 and Cogley, Primiceri and Sargent, 2010).² The use of a random walk specification has the counter-intuitive implication

¹There is also a literature using direct estimates of inflation expectations (e.g. using the Survey of Professional Forecasters) as proxies for trend inflation (e.g. Clark and Davig, 2008 or Williams, 2009). In this paper, we do not consider such sources of information about trend inflation.

²In this paper, we use univariate time series methods, estimating trend inflation using only information in observed inflation. As noted in Stock and Watson (2010), it is often hard to improve upon univariate forecasting procedures with inflation. However, the basic issues discussed, involving imposing constraints on trend inflation, will hold in multivariate models such as TVP-VARs.

that trend inflation (and long term inflation expectations) can grow in an unbounded fashion. There are few models of the inflation process that restrict the variation of trend inflation and long term inflation expectations. The purpose of the present paper is to fill this gap; to develop methods for bounding trend inflation in a manner consistent with the implicit inflation target ranges of central banks.

We develop a new model for inflation which restricts trend inflation to lie within bounds. These bounds can either be fixed or estimated from the data and we investigate both approaches as well as the sensitivity to choice of bounds. In our empirical work, involving quarterly CPI inflation, we find that inclusion of bounds is important in developing models with sensible properties and obtaining reasonable estimates of trend inflation. It also leads to improved forecast performance, particularly at longer horizons.

A further contribution of this paper is to introduce a computational algorithm, based on Chan and Strachan (2012), which allows for the efficient estimation of state space models involving inequality restrictions such as the ones in our model. Many models used for estimating trend inflation (e.g. the unobserved components stochastic volatility, UC-SV, model of Stock and Watson, 2007 or the time-varying parameter vector autoregressive model, TVP-VAR) are state space models and trend inflation is based on a vector of states. An advantage of this is that standard methods for statistical inference, involving the Kalman filter and state smoother, exist. However, when the states are subject to inequality constraints (such as occurs when trend inflation is bounded to lie in an interval), these methods are no longer valid and the obvious extensions of these methods to deal with inequality constraints can be computationally inefficient or even infeasible. These points are discussed in Koop and Potter (2011) who find the extension of the multi-move sampler proposed by Cogley and Sargent (2005) to be computationally infeasible unless the inequality constraints rarely bind. Koop and Potter (2011) propose an extension of a single-move sampler which is found to work when constraints are often binding, but this algorithm can be computationally demanding. The algorithm used in the present paper is not based on Kalman filtering and state smoothing, but rather uses a fast and simple Gaussian approximation to the posterior of the state vector. This approximation is then used to derive a proposal density for an accept-reject Metropolis-Hastings algorithm. We find this algorithm to work very well, providing accurate estimates of trend inflation in computationally efficient fashion.

This paper is organized as follows. In the next section, we discuss models

for trend inflation and introduce our bounded model. The third section carries out a prior predictive analysis using these models in order to illustrate the properties of these models and show the importance of bounding trend inflation. The fourth section describes a posterior simulation algorithm. The fifth section presents empirical results and the sixth section concludes.

2 Models for Trend Inflation

A wide variety of models for inflation can be placed into the following unobserved components framework:

$$\pi_t = \tau_t + c_t,$$

where π_t is an observed measure of inflation, τ_t is the inflation trend with the property that

$$\lim_{j \rightarrow \infty} E_t[\pi_{t+j}] = E_t[\tau_{t+j}] \text{ with probability 1} \quad (1)$$

and c_t is the inflation gap with the property that:

$$\lim_{j \rightarrow \infty} E_t[c_{t+j}] = 0 \text{ with probability 1.} \quad (2)$$

A simple constant parameter model in this framework would be

$$\begin{aligned} \tau_t &= \tau \\ c_t &= \rho_1 c_{t-1} - \dots - \rho_p c_{t-p} + \varepsilon_t, \end{aligned}$$

with $\varepsilon_t \sim N(0, \sigma_c^2)$. Indeed much of the early literature on inflation forecasting used such a framework where the information set was extended to include a measure of resource utilization. In this paper, following much of the recent literature, we focus on univariate modeling and impose that any unit root behavior is in the trend inflation component.

The simplest case we examine is a local trend model with stochastic volatility in the inflation gap:

$$\begin{aligned} \tau_t &= \tau_{t-1} + \varepsilon_t^\tau \\ c_t &= \varepsilon_t \exp\left(\frac{h_t}{2}\right), \\ h_t &= h_{t-1} + \varepsilon_t^h \end{aligned} \quad (3)$$

where $\varepsilon_t^\tau \sim N(0, \sigma_\tau^2)$, $\varepsilon_t \sim N(0, 1)$ and $\varepsilon_t^h \sim N(0, \sigma_h^2)$. These errors are assumed to be independent of one another and at all leads and lags.

Under the assumption that the information set is $\{\tau_t, c_t, \tau_{t-1}, c_{t-1}, \dots\}$ then

$$\lim_{j \rightarrow \infty} E_t[\pi_{t+j}] = E_t[\tau_{t+j}] = \tau_t$$

In this specification trend inflation is a driftless random walk. Thus, trend inflation and long-term inflation expectations must evolve in an unbounded fashion. This is inconsistent with the idea that central banks may, implicitly or explicitly, be targeting inflation and acting decisively when inflation moves outside of a desirable range. Of course, if σ_τ is small and τ_0 is a “reasonable” number it might take a very long time for this model to produce unusual behavior. Furthermore, if the actual inflation trend is bounded then this unbounded model could still provide a good approximation to inflation dynamics. As discussed in Stock and Watson (2007), this particular unobserved components model is equivalent to an integrated moving average representation for the inflation process where the moving average coefficient and innovation variance are time-varying.

Stock and Watson (2007) further generalize this model by allowing for stochastic volatility in the innovation to the inflation trend

$$\begin{aligned} \varepsilon_t^\tau &\sim N(0, \exp(g_t)) \\ g_t &= g_{t-1} + \varepsilon_t^g \\ \varepsilon_t^g &\sim N(0, \sigma_g^2). \end{aligned} \tag{4}$$

This version allows for the inflation trend to change at varying rates at different points in time. Specifically, they find that g_t was high in the late 1970s but more recently it has been very low. Thus, if g_t is low and σ_g^2 is small, the unbounded nature of the martingale trend might not become apparent over long periods.

Cogley and Sargent, in a number of papers, take a different approach to time variation in the inflation process. They do not use an unobserved components framework, instead they model inflation as

$$\begin{aligned} \pi_t &= \phi_{0t-1} + \phi_{1t-1}\pi_{t-1} + \varepsilon_t \exp\left(\frac{h_t}{2}\right) \\ h_t &= h_{t-1} + \varepsilon_t^h \\ \phi_t &= \phi_{t-1} + \varepsilon_t^\phi, \end{aligned}$$

where in Cogley, Primiceri and Sargent (2010, CPS hereafter) they also allow for stochastic volatility in the state equation for the vector of coefficients ϕ_t .³ Their approach to defining trend inflation can be illustrated using a first order autoregression of order 1. They impose that the autoregressive parameter is inside the unit circle, thus the ratio

$$\frac{\phi_{0t}}{1 - \phi_{1t}} = \tau_t$$

is the long-run mean (trend) of the inflation process if future values of the vector ε_t^ϕ are equal to zero. Thus, they define the inflation gap as

$$c_t = \pi_{t+1} - \frac{\phi_{0t}}{1 - \phi_{1t}}$$

with associated model specification:

$$c_{t+1} = \phi_{1t} \left(\pi_t - \frac{\phi_{0t}}{1 - \phi_{1t}} \right) + \varepsilon_{t+1} \exp\left(\frac{h_{t+1}}{2}\right).$$

As CPS discuss, this formulation allows the modeler to separately investigate inflation gap persistence from variations in trend inflation.

An unobserved components model with an autoregression in the transitory component is a more direct way of producing this decomposition and we adopt such a framework in this paper. We will focus on the first order autoregression case:

$$\begin{aligned} (\pi_t - \tau_t) &= \rho_t (\pi_{t-1} - \tau_{t-1}) + \varepsilon_t \exp\left(\frac{h_t}{2}\right) \\ \tau_t &= \tau_{t-1} + \varepsilon_t^\tau \\ h_t &= h_{t-1} + \varepsilon_t^h \\ \rho_t &= \rho_{t-1} + \varepsilon_t^\rho \end{aligned}, \tag{5}$$

where $\varepsilon_t \sim N(0, 1)$ and $\varepsilon_t^h \sim N(0, \sigma_h^2)$. One primary goal of the present paper is to examine the implications of bounding the behavior of τ_t and ρ_t such that the inflation trend and gap satisfy the criteria given above without imposing unit root behavior on inflation.

³We follow the working paper version which is Cogley, Primiceri and Sargent (2008). While CPS focus mainly on the vector autoregressive model it is more direct to just consider the univariate case for the points we are making.

Consider first the question of bounding trend inflation. We assume that the innovation in the state equation has the following form

$$\varepsilon_t^\tau \sim TN(a - \tau_{t-1}, b - \tau_{t-1}; 0, \sigma_\tau^2),$$

where $TN(a, b; \mu, \sigma^2)$ denotes the Gaussian distribution with mean μ and variance σ^2 truncated to the interval (a, b) .

Some properties of this bounded process follow immediately. First, using the symmetry of the Gaussian distribution, the unconditional mean is $\frac{b-a}{2}$ and the conditional expectation is

$$E_t[\tau_{t+1}] = \tau_t + \sigma_\tau \left[\frac{\phi\left(\frac{a-\tau_t}{\sigma_\tau}\right) - \phi\left(\frac{b-\tau_t}{\sigma_\tau}\right)}{\Phi\left(\frac{b-\tau_t}{\sigma_\tau}\right) - \Phi\left(\frac{a-\tau_t}{\sigma_\tau}\right)} \right] \text{ if } a \leq \tau_t \leq b$$

Note that for small values of σ_τ relative to $(b-a)$ the process has a conditional expectation that is almost identical to its current value if $|\tau_t - a| > 2\sigma_\tau$ and $|\tau_t - b| > 2\sigma_\tau$.

Next consider the question of bounding ρ_t . We assume $\varepsilon_t^\rho \sim TN(a_\rho - \rho_{t-1}, b_\rho - \rho_{t-1}; 0, \sigma_\rho^2)$. Similar features will apply to the bounded process for ρ_t as for trend inflation. For instance, we have

$$E_t[\rho_{t+1}] = \rho_t + \sigma_\rho \left[\frac{\phi\left(\frac{a_\rho - \rho_t}{\sigma_\rho}\right) - \phi\left(\frac{b_\rho - \rho_t}{\sigma_\rho}\right)}{\Phi\left(\frac{b_\rho - \rho_t}{\sigma_\rho}\right) - \Phi\left(\frac{a_\rho - \rho_t}{\sigma_\rho}\right)} \right] \text{ if } a_\rho \leq \rho_t \leq b_\rho. \quad (6)$$

The bounds in this case are determined by the requirement that the conditional expectation of the inflation gap process converges to zero as the forecast horizon increases. The simplest restriction is to limit ρ_t to be inside the unit circle or some other interval with constant limits (e.g. our empirical results use the interval $0 < \rho_t < 1$).

Our bounded model for trend inflation has the advantage that it satisfies the conditions (1) and (2), which other popular models of trend inflation do not. It also incorporates the implicit target bounds that central bankers may have. A drawback is that most central banks do not have explicit target ranges for inflation and, thus, the researcher does not know a and b . In this paper, we investigate two different treatments of this issue. First, we set a and b to constants that are subjectively selected so as to be reasonable. The researcher can try various choices for a and b and investigate how estimates of trend inflation change. Secondly, we treat a and b as unknown parameters which are estimated from the data.

A simple extension of the model above would be to allow for the bounds to vary over time. We do not pursue this extension, finding no compelling reason for thinking that, even in high inflation times it is likely that central bankers desired high trend inflation. Experiences such as the high inflation period of the 1970s are better thought of as times where deviations from the desired trend level of inflation were quite persistent. That is, central bankers were temporarily more tolerant of higher-than-desired trend inflation in these periods or less confident in their ability to quickly bring back inflation to the desired level).

3 Properties of Models of Trend Inflation

In order to further understand the properties of our model of bounded trend inflation relative to other options, we carry out a prior predictive analysis (see, e.g., Geweke, 2010). A prior predictive analysis involves simulating from the prior distribution and then, for each set of parameter values drawn from the prior, simulating an artificial data set. The properties of these artificial data sets can be compared to the properties of the actual data to see if the model is capable of generating the kinds of behavior observed in the data. The data consist of U.S. quarterly CPI from 1947Q1 to 2011Q3. Specifically, given the quarterly CPI figures z_t , we compute $y_t = 400(\log(z_t) - \log(z_{t-1}))$, and use it as the inflation rate.

To do a prior predictive analysis, we must first specify the set of models being compared and their priors. The next two sub-sections describe our choices.

3.1 Competing Models

Our bounded inflation model is given in (5). Since it specifies an AR process for deviations of inflation from trend, but places bounds on both trend inflation and the time-varying AR process, we refer to it as the **AR-trend-bound** model. We bound the time-varying AR coefficient to lie in the interval $(0, 1)$, but consider different treatments of the bounds for trend inflation (as outlined below).

In addition to this specification, we consider four other models. **AR-trend** is the same as the **AR-trend-bound** model, but without the bounds. That is, it specifies that all the innovations are Gaussian (instead of trun-

Table 1: A list of competing models.

Model	Description
Trend-SV	inflation trend model as in Stock and Watson (2007)
Trend	same as Trend-SV but the variance of state equation for τ_t is time-invariant
Trend-bound	same as Trend but $\tau_t \in (a, b)$:
AR-trend	autoregressive inflation trend model without any bounds
AR-trend-bound	autoregressive inflation trend model where $\tau_t \in (a, b)$ and $\rho_t \in (0, 1)$

cated Gaussian). **Trend** is the version of the UC-SV model given in (3) with a constant error variance in the state equation for trend inflation. That is, it assumes $\varepsilon_t^\tau \sim N(0, \sigma_\tau^2)$. **Trend-SV** is the version of the UC-SV model used in Stock and Watson (2007). Specifically, it replaces the assumption that $\varepsilon_t^\tau \sim N(0, \sigma_\tau^2)$ with (4). Note that adding stochastic volatility as in (4) can lead to a model which is difficult to estimate (even without adding an AR process with time-varying coefficient). Stock and Watson (2007) do not independently estimate σ_h^2 and σ_g^2 but place a restriction on them. In our model, adding this third latent process adds little and we omit this extension in our most general bounded inflation model for the sake of simplicity. Nevertheless, we include it in our set of competing models since it is a popular specification in the literature.

Lastly, **Trend-bound** has the same setup as **Trend**, but bounds trend inflation. That is, τ_t lies within (a, b) and we set $a = 0$ and $b = 5$. We summarize all five specifications in Table 1.

3.2 The Prior

The model given in (5) involves three state equations which must be initialized. The state equations for τ_t , ρ_t and h_t are initialized with

$$\begin{aligned}\tau_1 &\sim TN(a, b; \tau_0, \omega_\tau^2), \\ \rho_1 &\sim TN(0, 1; \rho_0, \omega_\rho^2), \\ h_1 &\sim N(h_0, \omega_h^2),\end{aligned}$$

where τ_0 , ω_τ^2 , h_0 , ω_h^2 , ρ_0 and ω_ρ^2 are known constants. In particular we set $\tau_0 = h_0 = \rho_0 = 0$, $\omega_\tau^2 = \omega_h^2 = 5$ and $\omega_\rho^2 = 1$. The prior variances are set to

be relatively large, so that the initial distributions for the states are proper yet relatively non-informative.

We denote the remaining parameters in the model by $\theta = (a, b, \sigma_h^2, \sigma_\rho^2, \sigma_\tau^2)$ and specify their prior as $p(\theta) = p(a, b)p(\sigma_h^2)p(\sigma_\rho^2)p(\sigma_\tau^2)$ where:

1. For the prior predictive analysis we fix a and b to the constants specified below;
2. $\sigma_\tau^2 \sim IG(\underline{\nu}_\tau, \underline{S}_\tau)$;
3. $\sigma_\rho^2 \sim IG(\underline{\nu}_\rho, \underline{S}_\rho)$;
4. $\sigma_h^2 \sim IG(\underline{\nu}_h, \underline{S}_h)$;

and $IG(\cdot, \cdot)$ denotes the inverse-Gamma distribution.

For the prior hyper-parameters we choose relatively small (and, thus, relatively non-informative) values for the degrees of freedom parameters: $\underline{\nu}_\tau = \underline{\nu}_\rho = \underline{\nu}_h = 10$. We next set $\underline{S}_\tau = 0.18$, $\underline{S}_\rho = 0.009$ and $\underline{S}_h = 0.45$. These values imply $\sqrt{E(\sigma_\tau^2)} = 0.141$, $\sqrt{E(\sigma_\rho^2)} = 0.0316$, and $\sqrt{E(\sigma_h^2)} = 0.224$. The degree of freedom parameters are chosen to be small so that the prior variances of σ_τ^2 , σ_ρ^2 and σ_h^2 are large. As for the values of the prior means, they reflect the desired smoothness of the corresponding state transition. For example, we have $\sqrt{E(\sigma_\tau^2)} = 0.141$, which implies that with high probability the difference between consecutive trend inflation, $\tau_t - \tau_{t-1}$, lies within the values -0.3 and 0.3 . This is calibrated for quarterly data, and we deem this range appropriate. We also note that $\sqrt{E(\sigma_h^2)} = 0.224$, which is similar to the value 0.2 used in Stock and Watson (2007).

AR-trend, **Trend** and **Trend-bound** are all restricted versions of the **AR-trend-bound** model and we use the same priors for parameters which are present in both models. For **Trend-SV**, we assume $\sigma_g^2 \sim IG(\underline{\nu}_g, \underline{S}_g)$, and $\sigma_h^2 \sim IG(\underline{\nu}_h, \underline{S}_h)$, where $\underline{\nu}_g = \underline{\nu}_h = 10$, and $\underline{S}_g = \underline{S}_h = 0.45$ (so that $\sqrt{E(\sigma_h^2)} = \sqrt{E(\sigma_g^2)} = 0.224$). This prior is in line with the priors used for other models and similar to the one used in Stock and Watson (2007).

The results of the prior predictive analysis presented below suggests that these priors are sensible.

3.3 Results of Prior Predictive Analysis

3.3.1 Properties of Trend Inflation

Before doing a prior predictive analysis involving the data itself, we initially investigate the properties of different trend inflation specifications. We compare results using the **AR-trend-bound** model to **Trend** and **Trend-SV**. We compute the predictive densities for future trend inflation, τ_{T+k} , with $k = 20$ under each specification. To compute the predictive densities, we first fix σ_τ at its prior mean (i.e., $\sigma_\tau = 0.141$) and fix τ_T as described below. For **Trend-SV**, we additionally set $\sigma_g = 0.224$ and $g_T = -3$ (approximately the value of the posterior mean for g_T). Then given τ_T , we generate $\tau_{T+1}, \dots, \tau_{T+k-1}$ (and g_{T+1}, \dots, g_{T+k} where appropriate) according to the relevant trend specification. Conditional on these draws, the density for τ_{T+k} (which is truncated Gaussian or Gaussian) can be calculated over a fine grid. We repeat this $R = 10^4$ times and compute the average.

For **AR-trend-bound**, we consider two bounded inflation specifications. In each τ_t follows a random walk with constant variance, but $\tau_t \in (a, b)$ — where $a = 0, b = 5$ and $a = 1, b = 4.5$. We denote the two specifications as **Bound-1** and **Bound-2**, respectively.

The results are reported in Figure 1 for various initial values of trend inflation. By definition, the models with bounds on trend inflation keep it within the bounds. If τ_T lies near the middle of the bounds, then **Bound-1** and **Bound-2** yield very similar predictive densities for trend inflation. However, if τ_T lies near the edges of the interval, then the choice of a and b can have a substantial impact on future trend inflation.

However, the most important feature of Figure 1 is that the predictive densities for either of the unbounded UC-SV models (**Trend** and **Trend-SV**) are very dispersed, containing a wide range of implausible values for trend inflation. The inclusion of stochastic volatility in the state equation for trend inflation increases dispersion further still relative to the UC-SV model with homoskedastic errors in this state equation.

This figure clearly illustrates the impact of allowing for random walk behavior of trend inflation. With such an unbounded specification, even at medium-term horizons like $k = 20$, trend inflation can wander far from reasonable values. We stress that the results of this exercise are not simply an artifact of using relatively noninformative priors. In this sub-section, we are producing medium term forecasts of trend inflation such as would be

made by a forecaster who knew precisely what trend inflation was and the parameter values were at time τ . If we had allowed for parameter uncertainty or uncertainty over what current trend inflation was, the predictive densities would be even more disperse using the unbounded formulations.

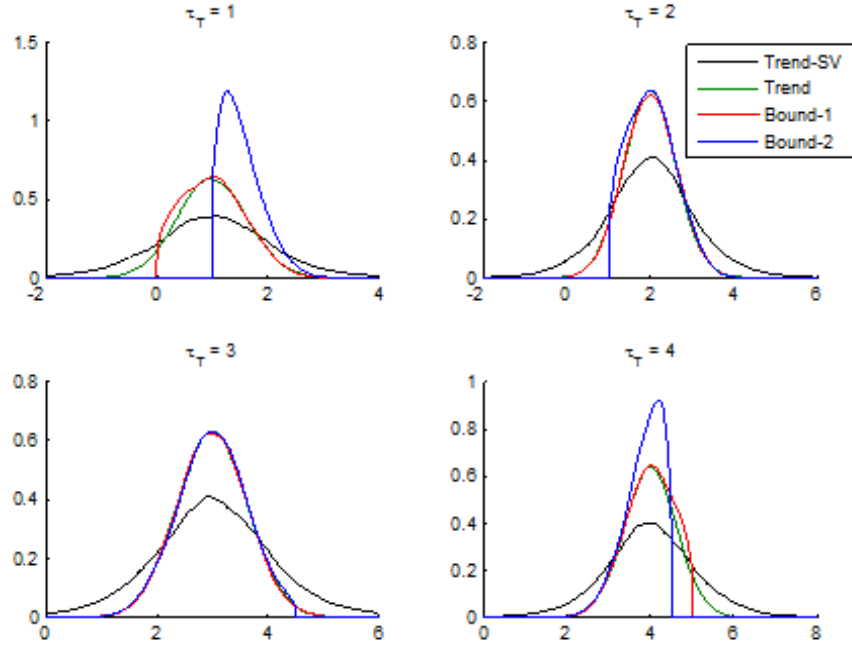


Figure 1: Predictive densities for τ_{T+k} under various trend inflation specifications where $k = 20$.

3.3.2 Prior Predictive Analysis for Data Features of Interest

We next perform a prior predictive analysis as suggested in Geweke (2010). Specifically, for each model, M_j , we first take a draw of its parameters from the prior and then simulate states from the state equations that are present in M_j . Given the drawn parameters, states and an initial value y_0 (which is fixed at 5.6739 — the CPI inflation rate in 1947Q2), we generate a dataset according to the measurement equation for M_j . Then, given the generated series, we compute various features of interest such as quantiles, variance,

autocorrelations, etc. We repeat this exercise for $R = 10^4$ times (each time with a new draw from the prior and the state equations), and construct the prior cumulative distribution function (cdf) for each of these features. That is, for each given feature z_i , we construct its empirical cdf under the model using the R draws $z_i^{(1)}, \dots, z_i^{(R)}$ generated as described above.

In this sub-section, we present results for all the models in Table 1. For models which bound trend inflation, we set $a = 0$ and $b = 5$.

The results are reported in Table 2 for 9 different features of interest. In each row we report the relevant figures of the stated feature of interest. The last row, labelled ‘‘MA coefficient’’, is the estimated moving average coefficient in an integrated MA(1) model. That is, it is the maximum likelihood estimate for ψ in the model: $\Delta y_t = u_t + \psi u_{t-1}$, where $u_t \sim N(0, \sigma^2)$. Columns 2–6 contain the prior cdf evaluated at the observed data under each of the models, i.e., $F_i(z_i^o | M_j) = \mathbb{P}(z_i \leq z_i^o | M_j)$ the probability that the feature under the prior and the model is less than the observed value.

Table 2: Prior cdfs of features.

Feature	Trend-SV	Trend	Trend-bound	AR-trend	AR-trend-bound
16%-tilde	0.833	0.856	0.734	0.767	0.757
median	0.678	0.889	0.816	0.754	0.801
84%-tilde	0.503	0.827	0.815	0.499	0.753
variance	0.205	0.690	0.707	0.348	0.635
fraction of $y_t < 0$	0.133	0.175	0.423	0.246	0.370
fraction of $y_t > 10$	0.464	0.812	0.794	0.465	0.731
lag 1 autocorrelation	0.315	0.771	0.814	0.615	0.540
lag 4 autocorrelation	0.227	0.638	0.687	0.300	0.550
MA coefficient	0.497	0.941	0.949	0.648	0.492

Values near zero or one indicate that the model is failing to mimic the relevant data feature well. It can be seen that all of the models are capable of fitting CPI inflation, although models with bounds on trend inflation do somewhat better (in the sense, that we are obtaining values of closer to 0.5). This is as might be expected for reasonably flexible models with relatively noninformative priors.

Table 2 shows that the problem with any of these models is not necessarily that they are too restrictive, but that they may be too flexible (in the

sense of accommodating very unreasonable behavior as well as reasonable behavior). To gain more insight on this issue and quantitatively compare the different models, we carry out another prior predictive exercise. First, we approximate the joint prior density of the features $p(z | M_j) = p(z_1, \dots, z_q | M_j)$ using the artificially generated data and a Gaussian kernel (as detailed in Geweke, 2010, p. 84–85). We then evaluate the prior density at the observed value $p(z = z^o | M_j)$. To compare models M_j and M_k , one can simply compute the Bayes factor $p(z = z^o | M_j)/p(z = z^o | M_k)$ relating to a particular feature of interest (or combination of various features of interest). In Table 3 we report the results of this comparison exercise. Specifically, we divide the features into three groups: “Quantile” includes the first three features of interest (16%-tilde, median and 84%-tilde), “Spread and Drift” includes the next three (variance, fraction of $y_t < 0$, and fraction of $y_t > 10$), and “Dynamics” include the last three features of interest (lag 1 autocorrelation, lag 4 autocorrelation and MA coefficient). The second column reports the log prior densities at the observed values under the **Trend-SV** model. Columns 3-6 report the log Bayes factors in favor of each model over the **Trend-SV** model.

Table 3: Log Bayes factors in favor of each model over the trend model.

Feature	Trend-SV	Trend	Trend-bound	AR-trend	AR-trend-bound
Quantile	-12.640	6.008	6.820	-654.581	6.832
Spread and Drift	-11.474	3.027	2.876	$-\infty$	4.881
Dynamics	-0.319	-2.957	-2.414	-0.709	2.083
All	-23.584	4.308	2.713	$-\infty$	13.307

Table 3 provides strong evidence in favor of our proposed **AR-trend-bound** model. That is, if we use all 9 features of interest, we find the Bayes factor comparing the **AR-trend-bound** model to the **Trend-SV** model to be about 6×10^5 . Remember that the **Trend-SV** model is a UC-SV model of a standard sort. Clearly, the prior predictive analysis is providing strong evidence in favor of our bounded inflation model relative to this common benchmark.

The Bayes factor (using all 9 features of interest) comparing the **AR-trend-bound** model to the **Trend-bound** inflation is about 40000, indicating that the addition of the AR lag is providing substantial benefits (i.e.

bounding trend inflation yields great improvements, but adding AR lags provides additional improvements beyond this).

However, simply adding a time-varying AR lag to a standard UC-SV specification (without bounding trend inflation or the AR coefficient) does not seem to be a good way to go. That is, **AR-trend** (without bounds) performs very poorly in Table 3. The reason for this is that the **AR-trend** model, with our relatively noninformative prior, ends up generating many explosive series. One response to this might be to consider a tighter prior on parameters to lessen or eliminate these explosive draws. However, many researchers may prefer a strategy of bounding the AR coefficient to lie in the stationary region at each point in time, rather than spending more effort (and risking criticism of other researchers) in choosing a more informative subjective prior.

The preceding comments are based on a Bayes factor involving all 9 features of interest. The reader may be interested in the dimension in which the **AR-trend-bound** model performs better (i.e. is it better at picking up the dynamics of inflation? Or issues relating to dispersion and tails?). The remaining rows in Table 3 suggest that the **AR-trend-bound** model has advantages in all of these directions.

4 Posterior Simulation Methods

In the preceding section, we used prior predictive simulation methods to investigate the properties of our model of bounded trend inflation given in (5). In our posterior analysis, we use the same prior as in the prior predictive analysis plus one additional degree of flexibility. To describe this new aspect, note that in the prior predictive results we simply set a and b to fixed constants. In our empirical section, we present results treating a and b as unknown parameters. In this case, we need priors for a and b and we assume these 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 $\bar{b} = 5$.

We develop an MCMC algorithm which sequentially draws from:

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

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

where $y = (y_1, \dots, y_T)'$, $h = (h_1, \dots, h_T)'$, $\tau = (\tau_1, \dots, \tau_T)'$ and $\rho = (\rho_1, \dots, \rho_T)'$. Complete details are given in the Appendix. Here we describe the outlines of the algorithm.

Due to the presence of the inequality constraints, $p(\tau | y, h, \rho, \theta)$ and $p(\rho | y, \tau, h, \theta)$ are non-standard and conventional methods of inference in state space models cannot be used. Instead we use an approach developed in Chan and Strachan (2012) for posterior sampling in nonlinear state space models. An essential element of this algorithm is a Gaussian approximation to the conditional density $p(\tau | y, h, \rho, \theta)$. This uses a precision based algorithm which builds upon results derived for the linear Gaussian state-space model by Chan and Jeliazkov (2009). The Gaussian approximation is then used as a proposal density for an accept-reject Metropolis-Hasting (ARMH) step. We also use this algorithm to draw from $p(\rho | y, \tau, h, \theta)$ and $p(h | y, \tau, \rho, \theta)$.

Both $p(a | y, \tau, h, \rho, \sigma_h^2, \sigma_\rho^2, \sigma_\tau^2, b)$ and $p(b | y, \tau, h, \rho, \sigma_h^2, \sigma_\rho^2, \sigma_\tau^2, a)$ are one-dimensional non-standard densities, and draws from each of which can be obtained using the Griddy-Gibbs algorithm. To sample from $p(\sigma_\rho^2 | y, \tau, h, \rho, a, b)$ and $p(\sigma_\tau^2 | y, \tau, h, \rho, a, b)$, both of which are non-standard densities, we use an independence-chain MH algorithm. Finally, $p(\sigma_h^2 | y, \tau, h, a, b)$ is an inverse-Gamma density.

5 Empirical Results

In this section, we present empirical results for the models listed in Table 1 using the prior given in Section 3.2 and a and b treated as unknown parameters with prior given in Section 4. We use quarterly CPI inflation as described at the beginning of Section 3. All results below are based on 50,000 draws from our MCMC algorithm (after a burnin period of 5000).

We divide our results into three sub-sections. The first shows that our MCMC algorithm is efficient. The second presents empirical results for trend inflation and other features of interest. The third investigates the forecasting performance of our bounded inflation model.

5.1 Evidence on Efficiency of the MCMC Algorithm

A common diagnostic of MCMC efficiency is the inefficiency factor, defined as:

$$1 + 2 \sum_{l=1}^L \phi_l,$$

where ϕ_l is the sample autocorrelation at lag length l , and L is chosen large enough so that the autocorrelation tapers off. To interpret it, note that independent draws from the posterior would give an inefficiency factor of 1. Inefficiency factors indicate how many extra draws need to be taken to give results equivalent to independent draws. For instance, if we take 50,000 draws of a parameter and find an inefficiency factor of 100, then these draws are equivalent to 500 independent draws from the posterior.

We report the inefficiency factors for the parameters and states for the five models in Table 4. Note that, for example, there are a total of T inefficiency factors associated with τ (one for each τ_t) and, hence, there are too many states for us to present inefficiency factors for all of them. Instead, we sort the inefficiency factors for τ , and report the 25th, 50th and 75th percentiles, denoted as $\tau_{(25\%)}$, $\tau_{(50\%)}$ and $\tau_{(75\%)}$ respectively. We adopt a similar strategy for the other states.

It is also worth noting that when we bound trend inflation and the time-varying AR coefficients, we no longer have a Gaussian linear state space models and, hence, computation is potentially much slower (e.g. the algorithm of Koop and Potter, 2011 is very inefficient relative to that of an unrestricted state space model). A sensible measure of performance is how much more inefficient an algorithm is than its unbounded variant. For instance, if we compare results for **AR-trend-bound** to **AR-trend**, we find the former to have higher inefficiency factors for the states which are bounded. However, the inefficiency factors tend to be only roughly 10 times higher than for the unbounded version of the model. Hence, to achieve a desired degree of accuracy, the bounded algorithm takes roughly 10 times as long as the unbounded variant. This is an appreciable increase in computation time. But, in the context of low-dimensional models such as the ones we are working with in this paper, this increase is not a substantial burden. Our algorithm is fast and efficient enough for easy use.

Table 4: Inefficiency factors of selected parameters.

Parameter	Trend-SV	Trend	Trend-bound	AR-trend	AR-trend-bound
$\tau_{(25\%)}$	1.1	1.0	87.9	1.2	24.4
$\tau_{(50\%)}$	2.0	1.8	125.3	2.3	34.6
$\tau_{(75\%)}$	10.8	12.2	166.5	2.6	98.5
$h_{(25\%)}$	16.9	5.9	4.3	1.9	2.3
$h_{(50\%)}$	38.6	9.0	6.1	2.1	2.7
$h_{(75\%)}$	85.8	24.4	9.9	2.3	3.0
$\rho_{(25\%)}$	–	–	–	1.5	16.2
$\rho_{(50\%)}$	–	–	–	1.6	19.4
$\rho_{(75\%)}$	–	–	–	1.8	24.2
$g_{(25\%)}$	45.6	–	–	–	–
$g_{(50\%)}$	69.7	–	–	–	–
$g_{(75\%)}$	201.2	–	–	–	–
σ_{τ}^2	–	120.9	586.0	28.8	188.5
σ_h^2	74.9	30.7	33.6	26.0	24.6
σ_{ρ}^2	–	–	–	25.2	62.9
σ_g^2	67.4	–	–	–	–
a	–	–	–	–	65.9
b	–	–	–	–	49.7

5.2 Estimates of Trend Inflation, Persistence and Volatility

Figure 2 presents estimates (posterior means) of trend inflation for the five models. It can be seen that large differences exist between the unbounded UC-SV models (**Trend-SV** and **Trend**) and the other models. The former are much more erratic and yield more extreme results than the latter. Trend inflation estimates from the unbounded UC-SV models tend to track actual inflation fairly closely. Especially for the **Trend-SV** model of Stock and Watson (2007) we are finding very high values of trend inflation (over 10% in some periods). Furthermore, trend inflation is far from being smooth in the unbounded UC-SV models, exhibiting rapid changes over short periods. We find these properties of unbounded UC-SV models to be counter-intuitive. As we have argued previously, the high inflation in the 1970s is better inter-

preted as reflecting large and persistent deviations from a fairly low trend rather than large increases in trend. Furthermore, abrupt changes in trend inflation seem inconsistent with theoretical or common-sense ideas of what trend inflation is.

Figure 2 also indicates the role played by the time-varying AR coefficients. Results for **AR-trend** indicate that, even without bounds, allowing for a time-varying AR coefficient has a large impact on estimates of trend inflation. **AR-trend** allows for estimation of a random walk in trend inflation and time variation in persistence of deviations from trend. In high inflation periods, this model attributes much of the inflation increase as reflecting the latter rather than the former. That is, given the choice, the econometric model estimates trend inflation as being fairly constant and allocates most of the change in inflation to the transitory component. We should stress, however, that in Section 3.3 we found the **AR-trend** model to have undesirable properties that were not present with the **AR-trend-bound** model.

It is also worth noting that the **Trend-bound** model (which assumes $\rho_t = 0$) is yielding estimates of trend inflation which are very different than the unbounded UC-SV models. This shows that it is not simply the inclusion of a time-varying AR component that is important in achieving sensible measures of trend inflation, the bounding is also playing an important role.

Figure 3 plots the posterior mean of trend inflation along with a credible interval for our **AR-trend-bound** model. The credible interval is not that wide, indicating our preferred model is estimating trend inflation fairly precisely.

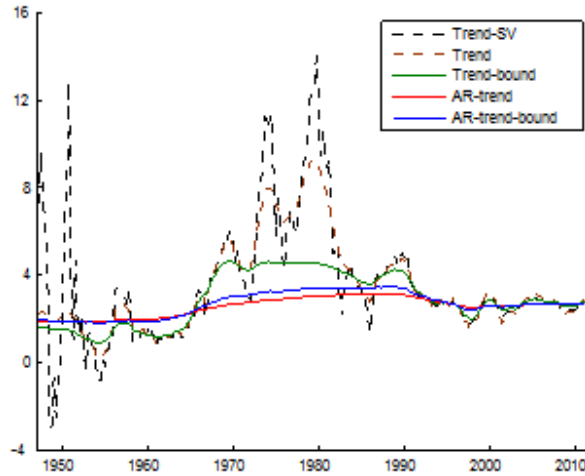


Figure 2: Posterior means of τ for various models.

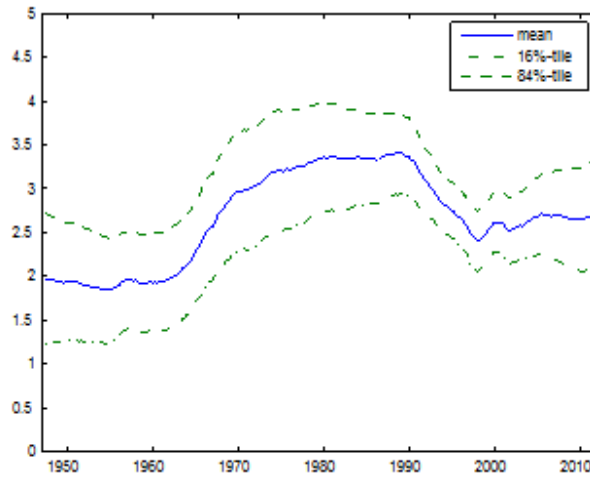


Figure 3: Posterior mean and quantiles of τ under the **AR-trend-bound** model.

Figures 4 and 5 are comparable in format as Figures 2 and 3, except for

volatilities (h_t) instead of trend inflation. Figure 4 shows that **AR-trend** and **AR-trend-bound** are producing volatility estimates which are similar to one another. However, the three UC-SV models without an autoregressive structure (**Trend**, **Trend-SV** and **Trend-bound**) are producing volatility estimates which exhibit some differences from the AR-trend models and from each other. **Trend-SV** (and to some extent **Trend**) are leading to lower volatilities than the other models since it is allocating much of the variation in the data to trend inflation (and, in the case of **Trend-SV** to g_t). **Trend-bound** is finding a much higher increase in volatility in the late 1970s and early 1980s than the other models. This is due to the fact that this model is producing a very stable estimate of trend inflation and does not allow for time-varying persistence in the transitory component. Thus, much of variability in the data is ascribed to h_t .

Figure 4 illustrates the sensitivity of volatility estimates to modeling assumptions. We argue that working with an **AR-trend** model is more sensible since it allows the data to decide whether variation is due to changes in trend inflation, persistence in the transitory component or volatility. Our previous results show the importance of bounding and, hence, argues for **AR-trend-bound**. Figure 5 shows that this model is estimating the volatilities in a precise fashion.

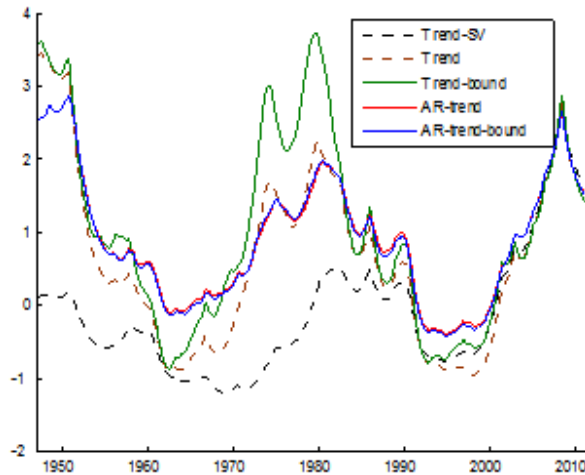


Figure 4: Posterior means of h under various models.

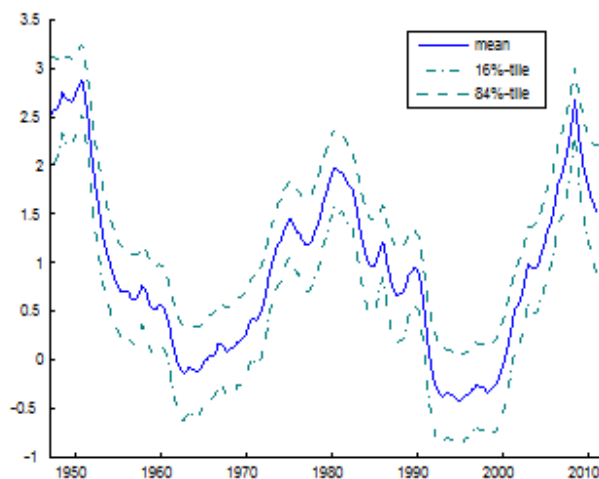


Figure 5: Posterior mean and quantiles of h under the **AR-trend-bound** model.

Figure 6 plots the estimates of ρ_t for the models which allow for time-varying persistence in deviations of inflation from trend. We are finding substantial changes in ρ_t over time. A fair degree of persistence (around 0.6 to 0.7) is found in the 1960s. This increases gradually through the 1960s and 1970s before peaking at around 0.9 in the late 1970s and early 1980s. Subsequently, during the time of the Great Moderation, persistence fell steadily. After the financial crisis there is some evidence that it began to increase slightly. The **AR-trend** and **AR-trend-bound** models are producing similar results. However, it is worth noting that, especially during the late 1970s, the estimate of ρ_t is closer to one with **AR-trend** and the posterior includes values greater than one. **AR-trend-bound** rules out this explosive region of the parameter space.

Figure 7 indicates that ρ_t is estimated fairly precisely by our model.

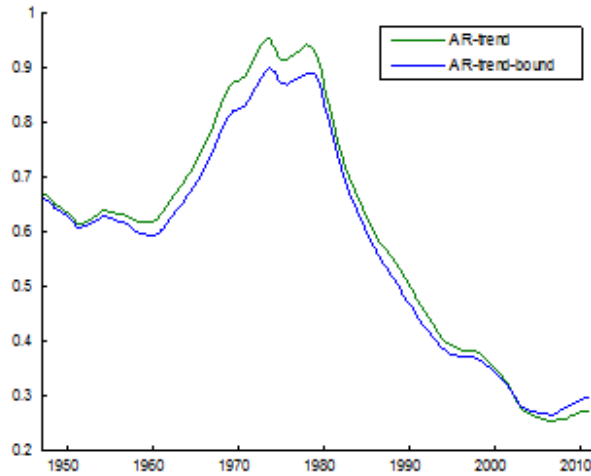


Figure 6: Posterior means of ρ_t under the **AR-trend** and **AR-trend-bound** models.

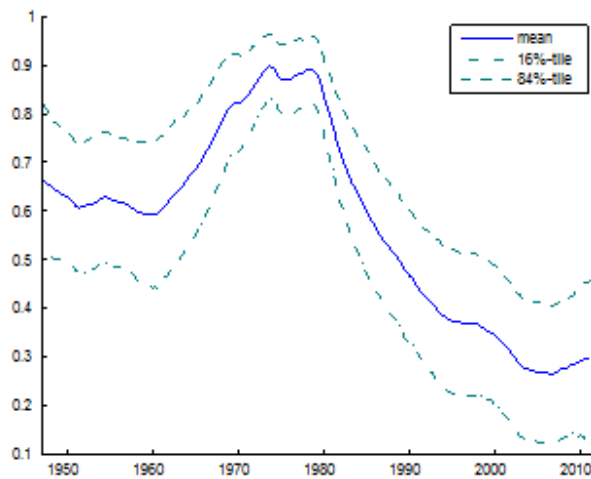


Figure 7: Posterior mean and quantiles of ρ_t under the **AR-trend-bound** model.

Finally, Figure 8 presents the posteriors for the bounds of the interval for trend inflation in the **AR-trend-bound** model. It can be seen that they are relatively flat over the ranges we allow for in the prior. We have found our results to be quite robust to treatment of a and b and Figure 8 shows why this is so. It is not the precise choice of values for a or b which is leading to the smooth and reasonable estimates of trend inflation we are finding.

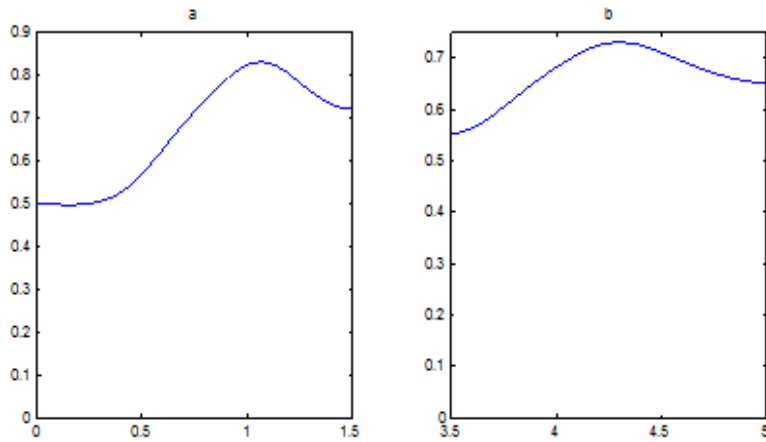


Figure 8: Kernel density estimates for $p(a|y)$ (left panel) and $p(b|y)$ (right panel) under the **AR-trend-bound** model.

5.3 Forecasting

We now evaluate the forecast performance of six models for forecasting the quarterly CPI inflation rate at different horizons. The first five models are the ones we considered previously: **Trend-SV**, **Trend**, **Trend-bound**, **AR-trend** and **AR-trend-bound**. We also add one another model which was used by Clark and Doh (2011). This is the time-varying-parameter autoregressive model, referred to as **TVP-AR**, specified as:

$$\begin{aligned}
 y_t &= \phi_{0t} + \phi_{1t}y_{t-1} + \dots + \phi_{pt}y_{t-p} + \varepsilon_t, \\
 \phi_t &= \phi_{t-1} + \varepsilon_t^\phi, \\
 h_t &= h_{t-1} + \varepsilon_t^h,
 \end{aligned}$$

where $\varepsilon_t \sim N(0, e^{h_t})$, $\varepsilon_t^\phi \sim N(0, \Omega)$, $\varepsilon_t^h \sim N(0, \sigma_h^2)$, $\phi_t = (\phi_{0t}, \phi_{1t}, \dots, \phi_{pt})'$ and $\Omega = \text{diag}(\omega_0, \omega_1, \dots, \omega_p)$. We fix $p = 2$ and evaluate our forecasts over the period 1975Q1 to 2011Q3 using root mean squared forecast error (RMSFE) and the average of log predictive likelihoods evaluated at the observed value.⁴ We consider five different forecast horizons: $k = 1, 4, 8, 12, 16$. We use iterated forecasts calculated with predictive simulation. To be precise, when forecasting using information through time t , predictive simulation is done of future values of the states and the dependent variable.

Table 5: RMSFEs for forecasting quarterly CPI inflation

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
Trend-SV	2.168	2.644	3.290	3.592	3.636
Trend	2.332	2.703	3.112	3.354	3.412
Trend-bound	3.032	3.067	3.079	3.148	3.140
AR-Trend	2.139	2.866	4.686	10.536	26.945
AR-trend-bound	2.089	2.430	2.916	3.116	3.168
TVP-AR	2.156	2.826	4.464	6.761	11.637

Table 6: Average log predictive likelihood for forecasting quarterly CPI inflation

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
Trend-SV	-2.052	-2.323	-2.494	-2.562	-2.624
Trend	-2.088	-2.332	-2.490	-2.548	-2.592
Trend-bound	-2.221	-2.341	-2.395	-2.434	-2.425
AR-Trend	-2.041	-2.264	-2.426	-2.471	-2.531
AR-trend-bound	-2.025	-2.214	-2.339	-2.358	-2.404
TVP-AR	-2.040	-2.250	-2.394	-2.413	-2.472

⁴The only data revisions in the CPI are produced by changing seasonal factors. Thus, our exercise is effectively a pseudo real time one. It should also be noted that one source of time variation in the parameters will be the various methodological changes in the CPI that have taken place over the last 60 years. Studies that use the latest vintage of PCE inflation or the GDP deflator will not have variation due to methodology changes.

Tables 5 and 6 present these forecast metrics. Regardless of forecast horizon or metric, with one exception our **AR-trend-bound** is exhibiting the best forecast performance. The one exception occurs at $k = 16$ where the RMSFE of the trend-bound model is slightly lower. We take these results as suggesting that bounding is useful for improving forecast performance. This is particularly true at longer horizons where RMSFEs for some of the models which do not bound trend inflation are very high due to the random walk behavior of trend inflation (or, in the case of the **TVP-AR** model, some explosive draws in the AR process). However, even with relatively short forecast horizons such as $k = 4$, substantial benefits of bounding trend inflation are appearing.

6 Conclusions

In this paper, we have introduced a new model that restricts trend inflation to lie in bounds. We have argued that this is a reasonable property in that central banks have implicit ranges for trend inflation which are deemed to be desirable and that these ranges are fairly constant over time. The most popular models of trend inflation used in the literature allow for trend inflation to follow a random walk. In theory, such models allow for trend inflation to grow in a counter-intuitively unbounded fashion. In practice, such models tend to yield trend estimates which follow actual inflation fairly closely leading to erratic estimates of trend inflation (e.g. estimating trend inflation as being very large in the 1970s).

A concern with our bounded trend inflation model is that it can no longer be estimated using methods for linear Gaussian state space models. Accordingly, we have investigated the use of an alternative algorithm for nonlinear state space models proposed by Chan and Strachan (2012) and found it to work well.

Our empirical results, based on quarterly CPI inflation, show the advantages of the bounded inflation model. Most importantly, it is yielding estimates of trend inflation which are very different from the popular UC-SV model of Stock and Watson (2007). We argue that our estimates are more sensible. A second finding is the importance of allowing for time-varying persistence in the transitory component of inflation where we find results similar to CPS but using a more reliable model of trend inflation. A final finding is that bounding trend inflation leads to improvements in inflation forecasting,

particularly at longer horizons.

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," Technical report, Research School of Economics, Australian National University, Canberra.
- Clark, T. and Davig, T. (2008). "An empirical assessment of the relationships among inflation and short- and long-Term expectations," *Federal Reserve Bank of Kansas City Research Working Paper* 08-05.
- Clark, T. and Doh, T. (2011). "A Bayesian evaluation of alternative models of trend inflation," manuscript.
- Cogley, T., Primiceri, G. and Sargent, T. (2010). "Inflation-gap persistence in the U.S.," *American Economic Journal: Macroeconomics*, 2, 43-69.
- Cogley, T. and Sargent, T. (2005). "Drifts and volatilities: Monetary policies and outcomes in the post WWII U.S.," *Review of Economic Dynamics* 8, 262-302.
- Cogley, T. and Sbordone, A. (2008). "Trend inflation, indexation, and inflation persistence in the new Keynesian Phillips curve," *American Economic Review*, 98, 2101-2126.
- Geweke, J. (2010). *Complete and Incomplete Econometric Models*. Princeton University Press, Princeton
- Ireland, P. (2007). "Changes in the Federal Reserve's inflation target: Causes and consequences," *Journal of Money, Credit and Banking*, 39, 1851-1882.
- Koop, G. and Potter, S. (2007). "Estimation and forecasting in models with multiple breaks," *Review of Economic Studies*, 74, 763-789.
- Koop, G. and Potter, S. (2011). "Time varying VARs with inequality restrictions," *Journal of Economic Dynamics and Control*, 35, 1126-1138
- Smets, F. and Wouters, R. (2003). "An estimated dynamic stochastic general equilibrium model of the Euro area," *Journal of the European Economic Association*, 1, 1123-1175.
- Stock, J. and Watson, M. (2007). "Why has U.S. inflation become harder to forecast?," *Journal of Money, Credit and Banking* 39, 3-33.
- Stock, J. and Watson, M. (2010). "Modeling inflation after the crisis," manuscript prepared for the Federal Reserve Bank of Kansas City Symposium, "Macroeconomic Policy: Post-Crisis and Risks Ahead," Jackson Hole,

Wyoming, August 26-28.

Williams, John C. (2009), "The risk of deflation," FRBSF Economic Letter, March 27.

Appendix: 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. $p(\rho | y, \tau, h, \theta)$;
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.,

$$\log p(y | \rho, h, y_0) \propto -\frac{1}{2}\iota_T' h - \frac{1}{2}(y - K^{-1}\mu_0 - \tau)' K' \Omega_y^{-1} K (y - K^{-1}\mu_0 - \tau), \quad (7)$$

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, (7) 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

$$\log p(\tau | \sigma_\tau^2, a, b) \propto -\frac{1}{2}\tau' H' \Omega_\tau^{-1} H \tau + g_\tau(\tau, a, b, \sigma_\tau^2), \quad (8)$$

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

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

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

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

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:

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

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

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

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

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

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

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