A New Model of Trend Inflation*

Joshua C.C. Chan

Australian National University

Gary Koop

University of Strathclyde

Simon M. Potter

Federal Reserve Bank of New York

September 27, 2012

Abstract: This paper introduces a new model of trend 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. 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. This research was supported by the ESRC under grant RES-062-23-2646. The online appendix associated with this paper is available at <http://personal.strath.ac.uk/gary.koop/research.htm>.

1 Introduction

Numerous studies have analyzed the behavior of inflation. 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, Cogley, Primiceri and Sargent, 2010 and Faust and Wright, 2012). For instance, many researchers have found the persistence of inflation declined 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

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, but offer some discussion on this topic below.

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 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 prior for the 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.

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 for inflation. However, the basic issues discussed, involving imposing constraints on trend inflation, will hold in multivariate models such as TVP-VARs.

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. The fact that our states are bounded leads to an algorithm which is a slight extension of Chan and Strachan (2012). We find this algorithm to work very well, providing accurate estimates of trend inflation in a 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 the 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 equation:

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

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

³We follow the working paper version which is Cogley, Primiceri and Sargent (2008).

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:

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

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.

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

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

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

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

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 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. The historical study of Weise (2011) provides evidence of such toleration of higher inflation, but not necessarily of tolerance of higher levels of trend inflation. Indeed, Volcker's October 1979 revolution could be thought of as part of an effort to speed up the adjustment towards low levels of trend inflation. Note that it might take a new policymaker such as Volcker some time to bring down the transitory component of inflation if it is highly persistent and big shocks had been hitting the economy. Adding in the Fed's dual mandate, in such a situation there could be a sustained period

of inflation above the upper bound. It is also interesting to note that there is evidence that Volcker did not seem particularly perturbed by moderately high inflation, but wanted to make sure long forward inflation expectations were low.

In short, our bounded model of trend inflation is consistent with the behavior of both inflation and Fed officials during the 1970s and 1980s. Of course, just because a model has these properties does not necessarily mean it is the literal truth. When estimating unobserved latent processes such as trend inflation (or the output gap, or trend-cycle decompositions, etc.), characteristics of the model which cannot be tested are always important in obtaining results. But these considerations hold as well in other models, such as those with random walk behavior for trend inflation (i.e. the fact the trend inflation is assumed to follow a random walk heavily influences the estimates of trend inflation which are obtained). Nevertheless, we feel that the idea that policymakers desire to keep trend inflation reasonably low is a sensible one and that bounds can be used to model this desire. If nothing else, investigating the implications of such a model is of interest. Furthermore, evidence from our forecasting exercise indicates that including bounds does lead to some improvements in forecast performance relative to models without bounds.

It is also worth noting that the idea of bounded trend inflation is reasonable in the current monetary policy regime that focuses on price stability. But there is always some chance that the policy regimes will change in the future. For example, Leeper and Walker (2012) develop a model where there is an upward drift in trend inflation if the debt level is not stabilized through

fiscal policy, but instead monetary means are used. At present, empirical evidence (including measures of inflation expectations) indicates such changes in monetary policy regime are not occurring, but if they occur models based on the idea that central bankers aim at long run price stability may not be applicable.

Of course, it may be sensible to say underlying inflation lies within a bounded interval $[a, b]$, but the choice of a and b may still be controversial. 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. We use priors on these parameters. It is likely that the prior for a will be relatively uncontroversial. Given the Fed's desire to avoid deflation, having a prior for a which ranges from zero to a low value is reasonable. Having a prior for b which is above the upper bound of the prior for a is an obvious thing to do. These prior choices are not controversial and we have found empirical results to be insensitive to precise choice of, e.g., what is meant by a "low value" for the lower bound. But what about the upper bound of the prior for b ? We focus on the case where this upper bound is 5%. We stress that we are not imposing an upper bound of 5% on trend inflation, but only saying that this is the highest value it can possibly be (our point estimate is actually below 5%). This is a fairly high value that is consistent with a sensible range of beliefs by policymakers. However, the online appendix associated with this paper contains a prior sensitivity analysis which shows estimates of trend inflation are not very

sensitive to choice of the upper bound of the prior for b (e.g. allowing for upper bound of the prior for b to be up to 10% has little impact on estimates of trend inflation).

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 truncated 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.

Table 1: A list of competing models for prior predictive analysis (top panel) and forecasting (bottom panel).

Models for the prior predictive analysis	
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 $\rho_t \in (0, 1)$
Additional models for forecasting	
Model	Description
AR-S-trend	same as AR-trend but the trend is stationary
AR-constant-trend	same as AR-trend but $\tau_t = \tau$ with prior $\tau \sim N(2.5, 1)$
AR-trend-bound-ρ_t	same as AR-trend but $\rho_t \in (0, 1)$
TVP-AR	time-varying parameter autoregressive model with $p = 2$
TVP-AR-R	same as TVP-AR , but any forecasts that are above 20% or below -15% are discarded
Random-walk	naive random walk forecast rule
Constant-2.5%	constant forecast rule where the forecast is fixed at 2.5%

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

$$\tau_1 \sim TN(a, b; \tau_0, \omega_\tau^2),$$

$$\rho_1 \sim TN(0, 1; \rho_0, \omega_\rho^2),$$

$$h_1 \sim N(h_0, \omega_h^2),$$

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: 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). A prior predictive analysis can be of interest for two reasons. First, it can be used to see if a model (i.e. likelihood function plus prior) is capable of generating data which replicates observed features of interest (e.g. to see if a model is capable of generating the observed moments or persistence of the data). In this sense, it can be used as a preliminary model selection device. Models which are incapable of generating observed behavior in empirically-important dimensions can be weeded out or expanded on at an early stage. Second, it can be used to see if a model is not parsimonious enough. If a model is capable of generating the observed features of interest, but is not parsimonious (in the sense that it is also capable of generating data sets which are very different from what is observed) the prior predictive analysis can reveal this through a prior Bayes factor. This can be used to choose between models. For instance, if one has two models, both of which can generate the observed features of interest, but one model is much more parsimonious, then the researcher would wish to choose the more parsimonious model. The Bayes factor produced in a prior predictive analysis can be used to make this choice.

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 the models listed in upper panel of 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 $p(z \mid M_j) = p(z_1, \dots, z_q \mid M_j)$ using the artificially generated data and a Gaussian kernel (as detailed in Geweke, 2010, pp. 84–85). We then evaluate the prior density at the observed value $p(z = z^o \mid M_j)$. To compare models M_j and M_k , one can simply compute the Bayes factor $p(z = z^o \mid M_j)/p(z = z^o \mid M_k)$ 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 $\bar{b} = 5$. In the online appendix we investigate sensitivity to the choice of a uniform distribution and as well as the prior hyperparameter choices.

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 online 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 burn-in 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 the models listed in Table 1 and some additional forecasting models.

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 listed in Table 1. It can be seen that large differences exist between the unbounded UC-SV models (**Trend-SV** and **Trend**) and the other mod-

els. 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 interpreted 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.

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.

With the **AR-trend-bound** model, we are finding trend inflation to be quite smooth with a point estimate that rarely goes much above 3%. In the **AR-trend-bound** model, long run inflation expectations should be equal to trend inflation. However, long run inflation expectations as measured by surveys, have occasionally been much higher than 3%. For instance, in the spring of 1981, the Blue Chip forecast of inflation 6-10 years ahead stood at 7 percent. How do we reconcile this apparent contradiction? In the early 1980s, our **AR-trend-bound** model yields estimates of ρ_t which are close to 1. This implies that inflation is very persistent and even forecasts several

years in the future may not have converged to trend inflation. This point is illustrated in Figure 9 which uses data up to 1980Q1 to construct the posterior densities for inflation expectations for 1986, 1988 and 1990, i.e., the densities for $\mathbb{E}(\pi_{t+h} | \pi_t, \dots, \pi_1)$, where $t = 1980Q1$ and $h = 24, 32, 40$. By 1990Q1, it can be seen that the predictive density is becoming less dispersed, being pulled back towards trend inflation. However, the predictive density in 1986Q1 is still quite dispersed and the value of 7 percent (reported in the Blue Chip forecast) is in a region of high probability. Thus, in the early 1980s, a forecaster who believed trend inflation to be low could easily have reported long-term inflation expectations which were much higher due to uncertainty about how persistent deviations from target would be tolerated by the Fed. These features are accommodated in our **AR-trend-bound** model.

Figure 9: Inflation expectations for 1986, 1988 and 1990 using data up to 1980Q1.

Given that the Fed has no explicit inflation target during the sample period under study, it is likely that the choice of prior for a and b will be thought most controversial. In the online appendix, we present results using different priors for these parameters, including allowing $\bar{b} = 10$. We find results to be fairly insensitive to prior choice. That is, estimating a and b from the data appears to be the most crucial choice. Even if we allow for unreasonably large choices such as $\bar{b} = 10$, the posterior mode of b remains between 4 and 5.

5.3 Forecasting

We now evaluate the forecast performance of several 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 several others. The first is a model which was used by Clark and Doh (2011). This is the time-varying-parameter autoregressive model, referred to as **TVP-AR**, specified as:

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

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 $p = 2$.

To investigate whether explosive draws of the AR coefficients can contaminate forecasts, we also produce forecasts from **TVP-AR**, but discard

any forecasts that are above 20% or below -15% . This version is referred to as **TVP-AR-R** and, for this ad hoc procedure, we do not produce density forecasts.

We also include a naive no-change forecast rule (**Random-walk**) where we simply use the current observation as the h -step-ahead forecast, i.e., $\hat{\pi}_{t+h} = \pi_t$. In addition, we also evaluate the simple constant forecast rule (**Constant-2.5%**) of fixing the h -step-ahead forecast at 2.5%: $\hat{\pi}_{t+h} = 2.5$. These two approaches do not produce density forecasts.

We also add several special cases of the models in Table 1 so as to investigate particular issues. The first model is a special case of **AR-trend-bound**, where the trend inflation $\tau_t = \bar{\tau}$ is constant over time with an informative prior $\bar{\tau} \sim \mathcal{N}(2.5, 1)$. We add this case in order to investigate the costs of simply assuming trend inflation, which we have found to be fairly smooth in models which include AR dynamics, to be constant. This version is referred to as **AR-constant-trend**.

The second model is a variant of **AR-trend**, where the transition equation for τ_t is stationary:

$$\tau_t = \mu + \varphi(\tau_{t-1} - \mu) + \epsilon_t^\tau,$$

and we set $\mu = 2.5$ and φ has the truncated normal prior $TN(0.95, 0.999; 0.98, 1)$ so that the trend process is stationary, but strongly persistent. This model is referred to as **AR-S-trend**.

The third model is another variation of **AR-trend**, where we impose the restriction $\rho_t \in (0, 1)$, but let τ_t remain an unrestricted random walk. This

version is referred to as **AR-trend-bound- ρ_t** , and is introduced so as to investigate the relative roles of bounding ρ_t and τ_t . We 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. We also present results of a nonparametric sign test of forecast performance (see Section 1.2.1 of Diebold and Mariano, 1995, for details). The null hypothesis of the sign test is that a benchmark model predicts equally well as a comparator. The test statistic asymptotically follows the standard normal distribution.⁵ We choose the benchmark model to be **AR-trend-bound**. This test can be calculated based on the RMSFE or the log predictive likelihoods. For the RMSFE comparison, a positive test statistic indicates that the benchmark model preforms better than the comparator. But for comparing the average log predictive likelihoods, the reverse is true: a negative test statistic indicates that the benchmark model preforms better.

Tables 5 and 6 present these forecast metrics and test statistics. Using the average of the log predictive likelihoods as the forecast metric, our

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.

⁵This limiting distribution should be viewed as an approximate guide since its derivation does not formally account for forecasting model parameter estimation error. The asymptotics of Giacomini and White (2006) justify use of this test when using rolling samples, but not recursive samples.

Table 5: RMSFEs for forecasting quarterly CPI. The test statistics of the sign test against **AR-trend-bound** are in parenthesis.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
Trend-SV	2.168 (-0.6)	2.643 (1.2)	3.290 (2.2)	3.592 (2.2)	3.635 (1.4)
Trend	2.332 (1.4)	2.703 (0.1)	3.112 (2.1)	3.354 (1.7)	3.412 (1.6)
Trend-bound	3.032 (2.6)	3.067 (3.1)	3.079 (1.9)	3.148 (3.5)	3.140 (1.1)
AR-trend	2.139 (-0.6)	2.866 (2.4)	5.248 (2.1)	10.536 (2.2)	26.945 (3.2)
AR-S-trend	2.135 (-0.1)	2.861 (0.6)	5.267 (-0.4)	10.592 (2.4)	26.996 (1.7)
AR-constant-trend	2.127 (1.1)	2.825 (1.6)	5.330 (-0.4)	11.166 (2.1)	29.704 (2.2)
AR-trend-bound-ρ_t	2.094 (0.2)	2.419 (1.2)	2.878 (0.2)	3.068 (0.2)	3.121 (1.1)
AR-trend-bound	2.089	2.430	2.916	3.116	3.168
TVP-AR	— 2.156 (0.7)	— 2.826 (1.4)	— 4.464 (2.4)	— 6.761 (2.1)	— 11.636 (2.9)
TVP-AR-R	2.158 (0.6)	2.642 (1.2)	3.133 (1.6)	3.242 (1.1)	3.153 (1.2)
Random-walk	2.440 (1.7)	3.072 (3.2)	3.563 (2.7)	3.709 (2.2)	3.780 (1.1)
Constant-2.5%	3.479 (4.0)	3.479 (1.7)	3.479 (1.7)	3.479 (3.9)	3.479 (3.5)

AR-trend-bound is exhibiting the best forecast performance at all forecast horizons. Using RMSFEs, either **AR-trend-bound** or **AR-trend-bound- ρ_t** forecasts best. Models which include AR dynamics do tend to forecast better than those which do not. The gains from bounding are particularly notable at longer forecast horizons. The fact that **AR-trend-bound- ρ_t** tends to forecast well indicates the particular importance of bounding ρ_t . Our results indicate that explosive draws in the AR process can contaminate

long-run forecast performance. Naive strategies, such as that **TVP-AR-R** (which simply discard extreme forecasts) do improve RMSFEs relative to **TVP-AR**, but formal bounding is found to work better. However, forecast improvements are also achieved by bounding trend inflation. And allowing for trend inflation to be constant or stationary, but very persistent, does not lead to forecasts which are better than those with random walk trend inflation. Our **AR-trend-bound** model also forecasts appreciably better than the no-change or constant forecasting strategy at all horizons by either forecast metric.

Table 6: Average log predictive likelihood for forecasting quarterly CPI. The test statistics of the sign test against **AR-trend-bound** are in parenthesis.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
Trend-SV	-2.052 (-2.7)	-2.323 (-3.7)	-2.494 (-5.5)	-2.562 (-5.4)	-2.624 (-5.0)
Trend	-2.088 (-3.4)	-2.332 (-2.6)	-2.490 (-4.2)	-2.548 (-4.4)	-2.592 (-2.4)
Trend-bound	-2.221 (-3.4)	-2.341 (-1.1)	-2.395 (-1.1)	-2.434 (-2.7)	-2.425 (-1.4)
AR-trend	-2.041 (-2.2)	-2.264 (-4.2)	-2.426 (-5.4)	-2.471 (-5.0)	-2.531 (-5.7)
AR-S-trend	-2.038 (-0.4)	-2.249 (-1.7)	-2.407 (-2.1)	-2.467 (-2.9)	-2.527 (-3.4)
AR-constant-trend	-2.035 (-1.4)	-2.305 (-3.4)	-2.496 (-2.9)	-2.534 (-2.7)	-2.603 (-3.5)
AR-trend-bound-ρ_t	-2.026 (-2.2)	-2.220 (-2.4)	-2.346 (-1.7)	-2.367 (-2.2)	-2.411 (-2.2)
AR-trend-bound	-2.025 —	-2.214 —	-2.339 —	-2.358 —	-2.404 —
TVP-AR	-2.040 (-1.1)	-2.250 (-0.4)	-2.394 (-1.6)	-2.413 (-2.2)	-2.472 (-2.9)

Are the forecasting improvements provided by the **AR-trend-bound** model statistically significant? Our nonparametric sign test of forecast per-

formance indicates that they often are, particularly when using log predictive likelihoods as the forecast metric.

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 per-

sistence 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,” CAMA Working Paper Series 2012-08, 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.
- Diebold, F. and Mariano, R. (1995). “Comparing Predictive Accuracy,” *Journal of Business Economics and Statistics*, 13, 134-144.
- Faust, J. and Wright, J. (2012). “Forecasting inflation,” manuscript available at http://www.econ2.jhu.edu/People/Wright/inf_forecast.pdf.
- Geweke, J. (2010). *Complete and Incomplete Econometric Models*. Princeton University Press, Princeton.
- Giacomini, R. and White, H. (2006). “Tests of conditional predictive ability,” *Econometrica*, 74, 1545-1578.
- 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.

Leeper, E. and Walker, T. (2012). “Perceptions and misperceptions of fiscal inflation,” NBER Working Paper 17903.

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.

Weise, C. (2011). “Political pressures on monetary policy during the US Great Inflation,” *American Economic Journal: Macroeconomics*, 4, 33-64.

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