Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

A Bounded Model of Time Variation in Trend Inflation, NAIRU and the Phillips Curve

---

CAMA Working Paper 10/2014
January 2014

Joshua C.C. Chan
Centre for Applied Macroeconomic Analysis, ANU

Gary Koop
University of Strathclyde

Simon M. Potter
Federal Reserve Bank of New York

Abstract

In this paper, we develop a bivariate unobserved components model for inflation and unemployment. The unobserved components are trend inflation and the non-accelerating inflation rate of unemployment (NAIRU). Our model also incorporates a time-varying Phillips curve and time-varying inflation persistence. What sets this paper apart from the existing literature is that we do not use unbounded random walks for the unobserved components, but rather use bounded random walks. For instance, trend inflation is assumed to evolve within bounds. Our empirical work shows the importance of bounding. We find that our bounded bivariate model forecasts better than many alternatives, including a version of our model with unbounded unobserved components. Our model also yields sensible estimates of trend inflation, NAIRU, inflation persistence and the slope of the Phillips curve.

Keywords

trend inflation, non-linear state space model, natural rate of unemployment, inflation targeting, Bayesian

JEL Classification

Address for correspondence:

(E) cama.admin@anu.edu.au

A Bounded Model of Time Variation in Trend Inflation, NAIRU and the Phillips Curve*

Joshua C.C. Chan
Centre for Applied Macroeconomic Analysis (CAMA)
Australian National University

Gary Koop
University of Strathclyde

Simon M. Potter
Federal Reserve Bank of New York

January 2014

Abstract: In this paper, we develop a bivariate unobserved components model for inflation and unemployment. The unobserved components are trend inflation and the non-accelerating inflation rate of unemployment (NAIRU). Our model also incorporates a time-varying Phillips curve and time-varying inflation persistence. What sets this paper apart from the existing literature is that we do not use unbounded random walks for the unobserved components, but rather use bounded random walks. For instance, trend inflation is assumed to evolve within bounds. Our empirical work shows the importance of bounding. We find that our bounded bivariate model forecasts better than many alternatives, including a version of our model with unbounded unobserved components. Our model also yields sensible estimates of trend inflation, NAIRU, inflation persistence and the slope of the Phillips curve.

Keywords: trend inflation, non-linear state space model, natural rate of unemployment, inflation targeting, 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 Centre for Economic Analysis.

1 Introduction

Parsimonious models of inflation and unemployment, inspired by the Phillips curve, have enjoyed great popularity for modeling latent states such as trend inflation or the non-accelerating inflation rate of unemployment (NAIRU) as well as for forecasting (see, among many others, Staiger, Stock and Watson, 1997 and Stella and Stock, 2012). In the inflation literature, univariate models such as the unobserved components stochastic volatility (UCSV) model of Stock and Watson (2007) are commonly-used. It is often found (e.g. Atkeson and Ohanian, 2001, Stella and Stock, 2012 and many others) that simple univariate methods forecast inflation as well as more complicated multivariate models, at least on average. However, it is also noted that multivariate models, often based on the Phillips curve relationship, can forecast better at some points in time, even if on average they do not beat univariate methods (e.g. Dotsey, Fujita, and Stark, 2010, Stock and Watson, 2010 and Stella and Stock, 2012 and many others). Such considerations motivate interest in bivariate models of inflation and unemployment, but not conventional linear constant coefficients models which assume the same Phillips curve relationship holds at each point in time. Instead the desire is for models where structures with economic interpretation such as the Phillips curve, trend inflation and underlying inflation can change over time. See, for instance, Clark and Doh (2011) which is a recent paper which surveys various approaches to modeling trend inflation. Related to this is the large literature on modeling or forecasting macroeconomic variables using models with time-varying coefficients (see, among many others, Cogley and Sargent, 2005, Primiceri, 2005, Sims and Zha, 2006, Cogley and Sbordone, 2008, Cogley, Primiceri and Sargent, 2010 and D'Agostino, Giannone and Gambetti, 2013).

These considerations motivate interest in models with three characteristics. First, they are bivariate models of inflation and unemployment. Second, some coefficients can change over time. Third, they are written in terms of latent state vectors that can be given an economic interpretation. Such a model is developed in Stella and Stock (2012), which is closely related to the model developed in this paper. This is a bivariate model with latent states which can be interpreted as a time-varying NAIRU and time-varying trend inflation. In addition, their model is based on a Phillips curve relationship but the slope of the Phillips curve can change over time.

The most important way that our approach differs from papers such as Stella and Stock (2012) is in its treatment of the latent states. Following most of the existing literature, Stella and Stock (2012) model trend inflation and the NAIRU as driftless random walks. Modeling trend inflation as a random walk is a component of many macroeconomic models (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), despite the fact that there are many reasons for thinking that trend inflation should not wander in an unbounded random-walk fashion. For instance, the existence of explicit or implicit inflation targets by central banks means that trend or underlying inflation will be kept within bounds and not allowed to grow in an unbounded fashion. Similar issues apply for NAIRU. In previous work, Chan, Koop and Potter (2013), we discussed these points in detail in the context of various theoretical models of trend inflation. We developed and presented evidence for a model for inflation which involved trend inflation following a bounded random walk. In this paper, we extend our earlier work to allow for trend inflation, NAIRU and the time-varying Phillips curve coefficients to all be bounded.

Models such as the UCSV or the one in Stella and Stock (2012) are normal linear state space models (apart from the stochastic volatility present in the errors). Standard econometric methods (e.g. involving the Kalman filter) exist for these models. Since standard econometric methods also exist for the treatment of stochastic volatility (e.g. the algorithm of Kim, Shephard and Chib, 1998), estimation of such models is theoretically straightforward. However, in practice, these models can be difficult to estimate without restrictions or strong prior information. For instance, Stella and Stock (2012) note that their likelihood function is flat in several dimensions and set seven parameters (error variances in measurement and state equations) to fixed constants. Alternatively, papers such as Cogley and Sargent (2005) and Primiceri (2005) use very informative priors (calibrated using a training sample of data). The priors on the error covariance matrices in their state equations are of particular importance.

The fact that our latent states are bounded mean our model is not a Normal linear state space model and, accordingly, conventional econometric methods cannot be used. Accordingly, we use an algorithm which is an extension of the ones developed in Chan and Jeliazkov (2009), Chan and Strachan (2012) and Chan, Koop and Potter (2013).

After developing and justifying our model and describing relevant econometric methods, we present empirical work using US data on CPI inflation and the unemployment rate. We find that our model forecasts better than many comparators, including an unbounded version of our model and an unrestricted VAR. Estimates of trend inflation, NAIRU, inflation persistence and the slope of the Phillips curve are found to be sensible.

2 A Bounded Trend Model for Inflation and Unemployment

We begin with a general bivariate model for inflation, π_t, and unemployment, u_t of the form:

$$\begin{array}{l} (\pi_t - \tau_t^\pi) = \rho_t^\pi \left(\pi_{t-1} - \tau_{t-1}^\pi\right) + \lambda_t \left(u_t - \tau_t^u\right) + \varepsilon_t^\pi \\ (u_t - \tau_t^u) = \rho_1^u \left(u_{t-1} - \tau_{t-1}^u\right) + \rho_2^u \left(u_{t-2} - \tau_{t-2}^u\right) + \varepsilon_t^u \\ \tau_t^\pi = \tau_{t-1}^\pi + \varepsilon_t^{\tau\pi} \\ \tau_t^u = \tau_{t-1}^u + \varepsilon_t^{\tau u} \\ \rho_t^\pi = \rho_{t-1}^\pi + \varepsilon_t^{\rho\pi} \\ \lambda_t = \lambda_{t-1} + \varepsilon_t^\lambda \end{array} . \tag{1}$$

Both dependent variables are written as deviations from trends, τ_t^π and τ_t^u. These trends are unobserved latent states which can be interpreted as trend (or underlying) inflation and the NAIRU. This model incorporates the properties that it is deviations of unemployment from NAIRU and deviations of inflation from its trend that drive the Phillips curve. These are features in common with the model of Stella and Stock (2012) and, for the inflation equation, with Stock and Watson (2007), Clark and Doh (2011) and Chan, Koop and Potter (2013). Thus, the first equation embodies a Phillips curve. The coefficients in the Phillips curve equation are time-varying and evolve according to random walks as in, e.g., Cogley and Sargent (2005), Primiceri (2005) and Stella and Stock (2012). Stella and Stock (2012) emphasize that time variation in λ_t is a useful extension of a conventional Phillips curve. If there are time periods when $\lambda_t = 0$ then the Phillips curve relationship does not exist. In general, a model where λ_t varies over time allows for the strength of the unemployment-inflation relation to vary over time, consistent with the episodic forecasting performance of the Phillips curve noted by Stella and Stock and others.

A feature of (1) which is not present in Stella and Stock (2012) is the time variation in ρ_t^π. This feature was incorporated in the univariate model of Chan, Koop and Potter (2013) and was found to be empirically important, allowing for differences in the way the Fed tolerates deviations of inflation from target. For instance, evidence in the historical study of Weise (2011), suggests that the high inflation period of the 1970s was not necessarily a time when the trend level of inflation increased dramatically (as would be implied by the version of our model with $\rho_t^\pi = 0$), but was a time when deviations from the desired level of inflation were quite persistent. That is, the Fed either was temporarily more tolerant of higher-than-desired inflation or less able to quickly return inflation to the desired level. After the 1970s, these characteristics were reversed. A model where ρ_t^π changes over time can model such features. Adding in the Fed's dual mandate also motivates the inclusion of unemployment and the Phillips curve relationship in our inflation equation.

The second equation implies AR(2) behavior for the unemployment rate. The AR(2) assumption is empirically sensible and commonly-used. Note that we are assuming constant coefficients in the unemployment equation. Stella and Stock (2012) also assume an AR(2) with constant coefficients for their unemployment equation. In our empirical work, we also consider a model where ρ_1^u and ρ_2^u vary over time, but find it to perform very poorly (see below). Accordingly, the main model we focus on does not have time-variation in the coefficients in the unemployment equation.

We assume that the errors in (1) (and the following equations) are independent with one another and at all leads and lags, but that the errors in the inflation equation exhibit stochastic volatility. Thus,

$$\begin{aligned} \varepsilon_t^\pi &\sim N(0, \mathrm{e}^{h_t}) \\ h_t &= h_{t-1} + \varepsilon_t^h, \\ \varepsilon_t^h &\sim N(0, \sigma_h^2). \end{aligned} \tag{2}$$

However, based on preliminary empirical work and following Stella and Stock (2012) we assume $\varepsilon_t^u \sim N(0, \sigma_u^2)$.

Thus far, we have specified a flexible bivariate model for unemployment and inflation which is similar to many in the literature. In particular, it incorporates the same features as Stella and Stock (2012) and is a bivariate generalization of unobserved components models used in Stock and Watson (2007), Clark and Doh (2011) and many others. However, the conventional literature would next assume that all the errors in the state equations for τ_t^π,τ_t^u, ρ_t^π and λ_t have normal distributions. It is with this assumption that we part with the existing literature.

As discussed in Chan, Koop and Potter (2013), unbounded random walk behavior for trend inflation is not plausible when there are inflation targets or anchored or constrained inflation expectations. Similarly, unbounded random walk behavior of the NAIRU is not consistent with much macroeconomic theory or the Fed's dual mandate. Thus, we model trend inflation and the NAIRU as bounded random walks. This can be done through the assumption that $\varepsilon_t^{\tau\pi}$ and $\varepsilon_t^{\tau u}$ have truncated normal distributions. Formally, we assume

$$\begin{array}{l}\varepsilon_t^{\tau\pi} \sim TN(a_\pi - \tau_{t-1}^\pi, b_\pi - \tau_{t-1}^\pi; 0, \sigma_{\tau\pi}^2) \\ \varepsilon_t^{\tau u} \sim TN(a_u - \tau_{t-1}^u, b_u - \tau_{t-1}^u; 0, \sigma_{\tau u}^2)\end{array} \tag{3}$$

where $TN(a, b; \mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 truncated to the interval (a, b). This specification ensures that τ_t^π lies in the interval (a_π, b_π) and τ_t^u lies in the interval (a_u, b_u). These bounds can either be set to particular values suggested by the underlying economics (e.g. if a central bank has an official inflation target or target interval) or estimated from the data. In our empirical work, we estimate the bounds a_π, b_π, a_u and b_u.

Similarly, we bound the time-varying coefficients of the Phillips curve ρ_t^π and λ_t. In time-varying parameter models it is common to use tight priors on the error variances (or covariance matrices) in the state equations which control the time variation in parameters (as done, e.g., in Cogley and Sargent, 2005 and Primiceri, 2005) or even to restrict them to particular values (as in Stella and Stock, 2012) to surmount problems caused by flat likelihoods in these parameter rich models. But the problems are often due to parameters wandering according to unbounded random walks that sometimes move into undesirable regions of the parameter space. We argue that this problem can be avoided by directly bounding the states in the state equations, thus avoiding the need for tight priors on the error variances in state equations. Thus, we bound ρ_t^π so that this coefficient is constrained to the interval $(0, 1)$ and never wanders into the explosive region of the parameter space. Similarly we bound λ_t to $(-1, 0)$ to avoid non-stable behavior and ensure that the Phillips curve has a negative slope. To be precise, we assume

$$\begin{array}{l}\varepsilon_t^{\rho\pi} \sim TN(-\rho_{t-1}^\pi, 1 - \rho_{t-1}^\pi; 0, \sigma_{\rho\pi}^2) \\ \varepsilon_t^{\lambda} \sim TN(-1 - \lambda_{t-1}, 0 - \lambda_{t-1}; 0, \sigma_\lambda^2).\end{array} \tag{4}$$

We also impose the stationary condition on the unemployment equation and assume $\rho_1^u + \rho_2^u < 1$, $\rho_2^u - \rho_1^u < 1$ and $|\rho_2^u| < 1$.

The Online Appendix for this paper describes our Bayesian estimation methods. In particular, given priors for the initial conditions and the other parameters, we derive a Markov chain Monte Carlo (MCMC) algorithm for simulating from the posterior. Most of the blocks in this algorithm are standard. The non-standard blocks are those for drawing the latent states. The Online Appendix also carries out a prior predictive analysis which shows that: i) our prior leads to sensible predictive densities for inflation and the unemployment rate, ii) our model is capable of explaining the main features (e.g. percentiles and dynamics) of the observed data, and iii) the bounds in our model receive strong support in Bayes factors constructed using the prior predictive density.

3 Empirical Work

3.1 Overview

We divide our empirical work into two sub-sections. The first is a forecasting exercise, comparing our bounded trend model of inflation and unemployment described in the preceding section to a range of alternative models for inflation and unemployment. We will use **Bi-UC** as an acronym for this bivariate unobserved components model. In the second, we present estimates of the trend inflation, the NAIRU and the other latent states produced by the **Bi-UC** model. Our data consist of quarterly CPI inflation rates and (civilian seasonally adjusted) unemployment rates from 1948Q1 to 2013Q1.

3.2 Forecasting Results

We begin by comparing our **Bi-UC** model to a range of alternative bivariate models for unemployment and inflation. All forecasts are out-of-sample and calculated recursively (i.e. forecasts for period $t+k$ are calculated using data from periods 1 through t). We measure forecast performance using the standard metrics of root mean squared forecast errors (RMSFEs) and the sums of log predictive likelihoods. Note that the latter is a common Bayesian model selection device. RMSFEs evaluate the performance only of the point forecasts whereas predictive likelihoods evaluate the quality of the entire predictive density. Most of the papers we cite, such as Stella and Stock (2012), focus on forecasting inflation. However, we present results for both inflation and unemployment. There is, of course, a large literature which attempts to forecast the unemployment rate using bivariate or multivariate specifications (e.g. Carruth et al., 1998) and/or specifications allowing for parameter change (e.g., Montgomery et al, 1998).

We compare the **Bi-UC** model to various alternatives, each designed to investigate some aspect of our specification. These models, along with their acronyms, are as follows:

1. **Bi-UC**: the model described in Section 2.

2. **Bi-UC-const-λ_t**: this is the restricted version of **Bi-UC** where λ_t is time-invariant.

3. **Bi-UC-const-ρ_t^π**: the restricted version of **Bi-UC** where ρ_t^π is time-invariant.

4. **Bi-UC-const-λ_t-ρ_t^π:** the restricted version of **Bi-UC** where both λ_t and ρ_t^π are time-invariant.

5. **Bi-UC-NoBound**: an unbounded version of **Bi-UC**, where all the states follow random walks without bounds.

6. **Bi-UC-NoBound**–τ_t^π-τ_t^u: a variation of an unbounded version of **Bi-UC**, where only τ_t^π and τ_t^u follow random walks without bounds.

7. **Bi-UC-NoSV**: the restricted version of **Bi-UC** where the measurement equation for π_t is homoskedastic (no stochastic volatility).

8. **Bi-UC-TVP-ρ^u**: an extension of **Bi-UC** where the AR coefficients in the unemployment equation (ρ_1^u and ρ_2^u) are time-varying.

9. **VAR(2)**: a standard, homoskedastic, VAR(2).

10. **VAR(2)-Minn**: VAR(2) with Minnesota prior.

11. **VAR(2)-SV**: A VAR(2) with heteroskedastic errors modeled using the stochastic volatility specification of Cogley and Sargent (2005).

12. **Bi-RW**: a bivariate random walk model.

13. **UCSV-AR(2)**: a model which is a univariate unobserved components for the inflation equation and an AR(2) for the unemployment equation.

14. **Stella-Stock**: the model in Stella and Stock (2012).

This list includes restricted versions of our **Bi-UC** model as well as a wide variety of bivariate specifications which have been used for inflation and unemployment. The latter include the model of Stella and Stock (2012) and less structural approaches based on VARs. TVP-VARs have also been used, with some success, for macroeconomic forecasting (see D'Agostino, Giannone and Gambetti, 2013) and this justifies our inclusion of the **Bi-UC-TVP-ρ^u** model which can be thought of as a fully TVP version of our model. The list also includes combinations of models which have been used individually for inflation and unemployment (i.e. the **UCSV-AR(2)** and bivariate random walk models).

The relatively non-informative prior we use for **Bi-UC** is given in the Technical Appendix. The priors for the models which restrict one or more latent states in **Bi-UC** to be constant are the same as those used in **Bi-UC**, except for restricted state(s). For the latter, the prior for the initial condition for the states used in **Bi-UC** becomes the prior for the constant coefficient. The model with unbounded states, **Bi-UC-NoBound**, simply removes the bounds. It can be interpreted as a restricted version of **Bi-UC** which sets lower/upper bounds such as a_{π}/b_{π} to $-/+$ infinity. For **Bi-UC-TVP-ρ^u**, we use the same priors for the initial conditions as we used for the constant coefficients, ρ_1^u and ρ_2^u, in **Bi-UC**. Complete details of all models, including the priors for the two versions of the **VAR(2)** and **VAR(2)-SV**, are given in the Online Appendix.

Tables 1, 2, 3 and 4 report RMSFEs and sums of log predictive likelihoods for unemployment and inflation individually. We also present, in Table 5, sums of log predictive likelihoods based on the joint predictive density for inflation and unemployment. Results are presented relative to the forecast performance of the **VAR(2)**. For RMSFEs, we take a ratio so that a number less than unity indicates a model is forecasting better than the **VAR(2)**. For sums of log predictive likelihoods (relative to **VAR(2)**) we take differences, so that a positive number indicates a model is forecasting better than the **VAR(2)**. Our forecast evaluation period begins in 1975Q1 and we consider forecast horizons of $k = 1, 4, 8, 12, 16$ quarters.

Overall, these tables indicate that our **Bi-UC** is forecasting well. Often (for any forecast metric or forecast horizon) it is the best model. For cases where it is not the best, it is forecasting almost as well as the best model. Furthermore, most of the other candidates forecast very poorly in at least one case. We elaborate on these points in detail below.

Relative to the model without bounds, a clear pattern emerges. At longer forecast horizons, the forecast performance of **Bi-UC-NoBound** deteriorates substantially, particularly for inflation. The unbounded random walk behavior of the latent states in this model is clearly leading to unreasonable long run forecasts. However, even with short run forecasts, **Bi-UC-NoBound** is almost always beaten by **Bi-UC**. In our introduction, we tried to argue that bounding latent states like trend inflation was economically sensible and more consistent with central bank behavior than assuming unbounded random walk behavior. Here we have also established empirically that bounding improves forecast performance. However, the improvements in long run forecast performance largely come through bounding ρ_t^{π} and λ_t, since **Bi-UC-NoBound**–τ_t^{π}-τ_t^u forecasts much better than **Bi-UC-NoBound**.

Similarly, **Bi-UC-TVP-ρ^u,** the extension of **Bi-UC** which allows for time-variation in the coefficients in the unemployment equation, forecasts poorly, particularly at longer forecast horizons. In fact, it often exhibits the worst forecast performance at medium and long horizons. However, even with short run forecasts, **Bi-UC-TVP-ρ^u** occasionally forecasts poorly and almost never beats **Bi-UC**. It is for this reason that in (1) we assumed constant coefficients in the unemployment equation. Empirically, allowing for time-varying coefficients in the unemployment equation is not warranted and contaminates forecasts.

Our **Bi-UC** also tends to forecast better than some standard implementation of VARs. In terms of MSFEs, the VAR with Minnesota prior forecasts unemployment well while the VAR(2) with stochastic volatility forecasts inflation well. However, if we look at log predictive likelihoods, **Bi-UC** (or restricted variants of it) are always forecasting substantial better than any of the VARs. This suggests that bounding, although useful for getting good point forecasts, is particularly useful for getting the dispersion and tails of the predictive distribution correct.

With regards to all the restricted versions of our **Bi-UC** model, none of them improves forecast performance greatly. Restricting λ_t to be constant is the best of these restricted models: imposing this restriction slightly improves forecasts of inflation (but causes forecasts of the unemployment rate to deteriorate slightly). Table 5 (which presents joint predictive likelihoods for inflation and unemployment) is perhaps the best single summary of forecasting performance we present. In this table, **Bi-UC-const-λ_t** forecasts best for short horizons, but **Bi-UC** forecasts best at medium and longer forecast horizons. Restricting ρ_t to be constant has a more substantive, negative, impact on forecast performance.

The model of Stella and Stock (2012) and the **UCSV-AR(2)** are also easily beaten by our **Bi-UC** in terms of forecast performance. Random walk forecasts do even worse, particularly at long forecast horizons. Finally, including stochastic volatility in the inflation equation is important since its inclusion leads to substantial improvements in inflation forecasts. This can be seen by comparing homoskedastic to heteroskedastic versions of the same model (e.g. comparing **Bi-UC** to **Bi-UC-NoSV**).

Figures 2 through 4 allow us to look at forecast performance of several of our main approaches in more detail by plotting cumulative sums of log predictive likelihoods for inflation, unemployment and the two variables jointly, respectively. Overall, it can be seen that **Bi-UC** model is consistently forecasting well. Furthermore, a few particular patterns are worth noting. For forecasting inflation for $k = 1$, it is clear that unobserved com-

Table 1: Relative RMSFEs (against **VAR(2)**) for forecasting inflation; 1975-2013.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	1.00	1.00	1.00	1.00	1.00
VAR(2)-Minn	1.00	1.00	0.99	0.98	0.98
VAR(2)-SV	1.01	0.93	0.93	0.94	0.94
Bi-UC	0.96	0.99	1.01	0.97	0.99
Bi-UC-const-λ_t	0.95	0.94	0.95	0.93	0.95
Bi-UC-const-ρ_t^π	1.00	0.99	0.94	0.90	0.91
Bi-UC-const-λ_t**-**ρ_t^π	0.99	0.96	0.93	0.90	0.90
Bi-UC-NoBound	0.98	1.07	1.54	2.58	5.64
Bi-UC-NoSV	0.97	0.98	1.00	0.97	0.99
Bi-UC-TVP-ρ^u	0.97	0.93	1.46	3.47	12.92
Bi-RW	1.09	1.12	1.11	1.13	1.15
UCSV-AR(2)	1.06	0.98	0.97	0.99	1.00
Bi-UC-NoBound–τ_t^π-τ_t^u	0.95	0.96	1.02	0.99	1.00
Stella-Stock	1.03	0.98	0.98	0.96	0.97

Table 2: Relative RMSFEs (against **VAR(2)**) for forecasting unemployment; 1975-2013.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	1.00	1.00	1.00	1.00	1.00
VAR(2)-Minn	0.96	0.90	0.91	0.96	1.00
VAR(2)-SV	1.00	1.01	1.03	1.01	0.99
Bi-UC	0.95	0.94	0.94	0.93	0.91
Bi-UC-const-λ_t	0.97	0.98	0.98	0.96	0.93
Bi-UC-const-ρ_t^π	0.97	0.99	1.00	0.97	0.94
Bi-UC-const-λ_t**-**ρ_t^π	0.97	0.99	1.00	0.98	0.94
Bi-UC-NoBound	0.95	0.93	0.95	0.97	0.97
Bi-UC-NoSV	0.96	0.98	1.00	0.98	0.95
Bi-UC-TVP-ρ^u	1.59	1.79	1.07	0.97	0.97
Bi-RW	1.27	1.07	1.09	1.15	1.14
UCSV-AR(2)	0.96	1.00	1.14	1.30	1.41
Bi-UC-NoBound–τ_t^π-τ_t^u	0.94	0.91	0.93	0.95	0.96
Stella-Stock	1.34	1.00	0.98	0.99	0.98

Table 3: Sum of log predictive likelihoods (against **VAR(2)**) for forecasting inflation; 1975-2013.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	0.0	0.0	0.0	0.0	0.0
VAR(2)-Minn	-1.7	1.2	1.1	1.8	4.0
VAR(2)-SV	2.6	25.2	18.6	16.6	15.2
Bi-UC	27.9	29.3	27.9	29.9	26.2
Bi-UC-const-λ_t	29.1	32.1	32.8	35.2	30.9
Bi-UC-const-ρ_t^π	21.4	18.5	20.3	23.3	23.8
Bi-UC-const-λ_t**-**ρ_t^π	22.6	20.3	20.6	22.9	24.7
Bi-UC-NoBound	25.4	25.0	15.3	13.1	6.6
Bi-UC-NoSV	6.0	10.6	15.8	18.4	18.4
Bi-UC-TVP-ρ^u	27.6	30.3	23.8	20.5	12.3
Bi-RW	-19.8	-50.9	-75.4	-97.3	-115.2
UCSV-AR(2)	19.1	19.5	14.5	12.5	9.0
Bi-UC-NoBound$-\tau_t^\pi$**-**τ_t^u	28.9	30.8	25.2	24.6	21.5
Stella-Stock	18.9	14.6	11.6	17.1	17.6

Table 4: Sum of log predictive likelihoods (against **VAR(2)**) for forecasting unemployment; 1975-2013.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	0.0	0.0	0.0	0.0	0.0
VAR(2)-Minn	3.2	12.9	17.8	14.6	10.3
VAR(2)-SV	-0.3	0.5	5.9	10.6	9.1
Bi-UC	-9.8	-1.3	13.4	26.3	32.1
Bi-UC-const-λ_t	0.5	0.2	4.3	12.2	17.1
Bi-UC-const-ρ_t^π	0.0	-2.7	0.6	9.9	15.7
Bi-UC-const-λ_t**-**ρ_t^π	-10.7	-5.6	5.3	18.0	25.5
Bi-UC-NoBound	-8.1	-2.9	8.0	15.8	19.1
Bi-UC-NoSV	-10.5	-5.6	5.0	17.3	24.2
Bi-UC-TVP-ρ^u	-61.4	-86.6	-13.2	9.3	11.4
Bi-RW	-51.8	-14.1	-13.3	-12.7	-7.1
UCSV-AR(2)	-10.1	-11.5	-19.9	-32.7	-46.3
Bi-UC-NoBound$-\tau_t^\pi$**-**τ_t^u	-10.1	-1.9	10.2	18.6	20.4
Stella-Stock	-62.0	-7.6	0.6	0.6	-1.9

Table 5: Sum of log predictive likelihoods (against **VAR(2)**) for jointly forecasting inflation and unemployment; 1975-2013.

	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	0.0	0.0	0.0	0.0	0.0
VAR(2)-Minn	1.0	12.6	22.0	25.7	24.8
VAR(2)-SV	2.4	26.0	23.0	21.1	7.9
Bi-UC	19.0	23.8	34.8	50.3	50.3
Bi-UC-const-λ_t	30.5	28.2	30.1	40.2	37.4
Bi-UC-const-ρ_t^π	22.3	7.6	7.9	18.8	26.0
Bi-UC-const-λ_t-ρ_t^π	12.8	9.9	20.5	34.9	42.6
Bi-UC-NoBound	18.4	23.0	26.0	31.5	23.0
Bi-UC-NoSV	-3.5	0.8	17.7	37.2	45.3
Bi-UC-TVP-ρ^u	-33.0	-58.8	13.1	34.7	25.4
Bi-RW	-70.7	-66.4	-79.8	-94.7	-104.1
UCSV-AR(2)	9.9	7.0	2.7	-5.0	-19.7
Bi-UC-NoBound–τ_t^π-τ_t^u	19.7	28.4	37.3	43.4	39.0
Stella-Stock	-41.1	12.7	18.5	13.4	6.6

ponents models (including unbounded and bounded versions of them), did better at capturing the abrupt shock of the financial crisis than the VARs. For unemployment, this statement also holds true, but to a lesser extent.

The poor performance of the unbounded trend model, **Bi-UC-NoBound**, in forecasting inflation at longer horizons mostly occurs in the relatively stable period from the early 1980s through early 1990s. However, for unemployment, this period of poor forecasting occurred later, from the mid 1990s through 2001. In general, this unbounded model is not forecasting poorly during unstable times, when its unbounded random walk states allow for rapid change, but in more stable times.

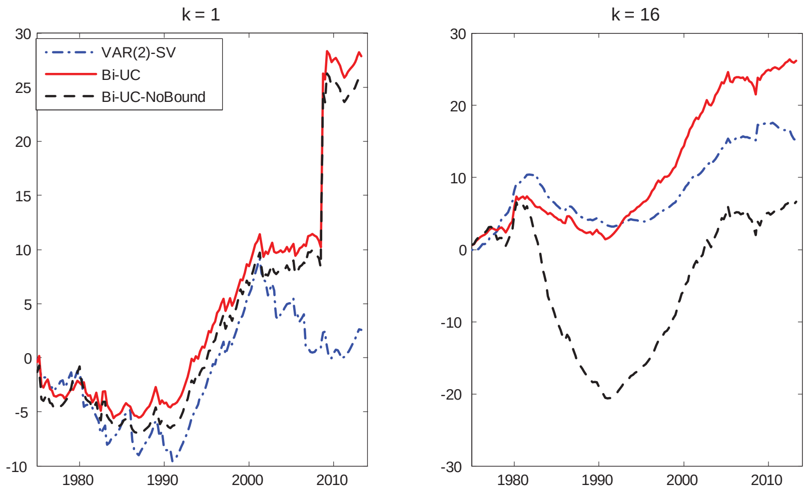

Figure 4: Cumulative sums of log predictive likelihoods for forecasting inflation (relative to **VAR(2)**).

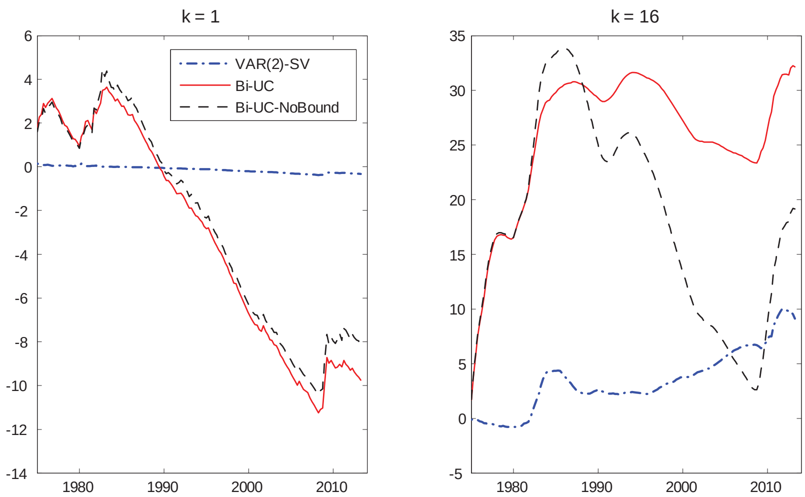

Figure 5: Cumulative sums of log predictive likelihoods for forecasting unemployment (relative to **VAR(2)**).

3.3 Estimates of Trend Inflation, NAIRU and Other Latent States

In previous work with a simple univariate model, Chan, Koop and Potter (2013), we showed the benefits of bounding in producing reasonable estimates of trend inflation. If trend inflation is left to evolve according to an unbounded random walk, it can track actual inflation too closely, leading to erratic (and, in the last 1970s very high) estimates of trend inflation. But by keeping trend inflation bounded in an interval consistent with beliefs about the behavior of central bankers, smoother and much more sensible behavior is obtained. In the present case, we obtain similar results for inflation. For unemployment, the use of bounding also helps avoid excessive swings in NAIRU.

Figure 5 plots the four main latent states, or τ_t^π,τ_t^u, ρ_t^π and λ_t, estimated using the full sample. That is, Figure 5 contains smoothed estimates, based on information available at time T (as opposed to filtered estimates, to be presented shortly). With regards to the observed increase in inflation in the late 1970s, it can be seen that our bounded model of trend inflation chooses to estimate it as a slight increase in trend inflation, but a much larger increase in persistence (i.e. ρ_t^π increases substantially at this time). This is consistent with a story where the Fed has kept a fairly low implicit inflation target, but was more willing to tolerate (or less able to correct) deviations from target in the 1970s than subsequently.

Our model also implies a smoothly evolving NAIRU. Despite large fluctuations in unemployment (see Figure 1), Figure 5 suggests NAIRU increases from roughly 5 to 7% the end of the 1980s before subsequently falling to the region of 6%. These numbers are consistent with the existing literature. For instance, Staiger, Stock and Watson (1997) present NAIRU estimates in the 5.5% to 5.9% range. The unemployment equation is very persistent, as the posterior means for ρ_1^u and ρ_2^u are estimated to be 1.617 and -0.674, respectively. It is also worth noting that our estimates of the bounds are reasonable: the posterior means of a_π, b_π, a_u and b_u are$0.447, 4.337, 3.956$ and 7.642, respectively.

The coefficient controlling the slope of the Phillips curve is, as expected, a negative number and tends to be around -0.4. Consistent with the evidence of the preceding sub-section, there is less evidence that it varies over time. However, there is some tendency for it to fall (become more negative) in the late 1970s and subsequently tend to move towards zero, indicating a weakening of the Phillips curve relationship. This is consistent with the finding of Stella and Stock (2012) that the Phillips curve was steeper in the

1970s than the 1990s. However, they also found a steepening after 2008, which we do not find to any great extent.

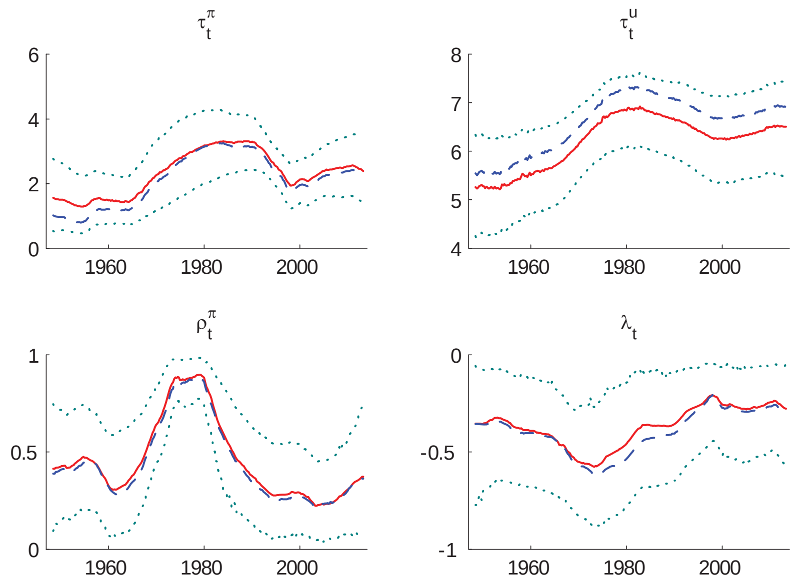

Figure 7: Posterior estimates for the latent states using the full sample. The solid lines represent the posterior means under the **Bi-UC** while the dotted lines are the 5 and 95 percentiles. The dashed lines represent the posterior means under the **Bi-UC-NoBound**–τ_t^π-τ_t^u.

Figure 5 contains smoothed estimates of the states. These are useful for a retrospective analysis using all available information. However, for some purposes, filtered estimates are useful. That is, it is also useful to consider a real time analysis, estimating the states at time t using data which were available at time t. These are presented in Figure 6. The broad patterns in Figure 6 are similar to Figure 5, suggesting that our bivariate unobserved components model can provide sensible real-time estimates of NAIRU and the trend inflation. As one would expect of filtered estimates, they tend to be slightly more erratic than smoothed estimates. In some ways, this is sensible. For instance, the smoothed estimates have the NAIRU rising as of shortly after 2000. This is due to a slight increase in unemployment after the

2001 slowdown and the much larger increase in unemployment in the recent recessions. Figure 5 smooths these two together as a gradual rise in NAIRU throughout the 2001-2011 period. But the filtered estimates which, in 2007 do not know the financial crisis which is about to occur, are still low in 2007 and only start rising after the financial crisis has hit. It is also interesting to note that the filtered estimates of trend inflation and ρ_t^{π} indicate that the 1980s are a time of decreasing inflation persistence and it is only later that declines in trend inflation occurred.

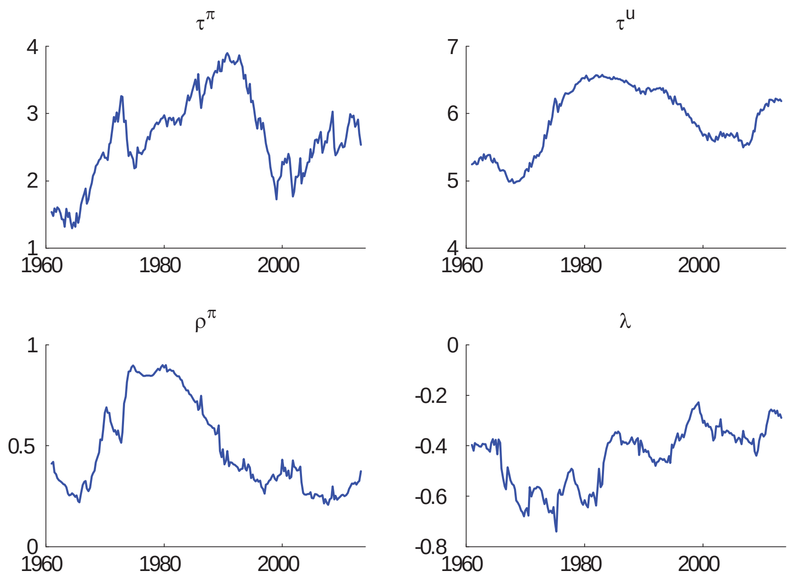

Figure 8: Estimates of the filtered states $E(\eta_t \,|\, Data_t)$, where $\eta_t = \tau_t^{\pi}, \tau_t^{u}, \rho_t^{\pi}, \lambda_t$ and $Data_t$ is the data from the beginning of the sample to time t.

4 Conclusions

In this paper, we have developed a bivariate unobserved components model for inflation and unemployment based on the Phillips curve. The model is written in terms of deviations of inflation from its trend and deviations of unemployment from its natural rate. One result of our paper is that such unobserved components models are attractive since they directly provide estimates of trend inflation and the NAIRU and can also forecast as well or better as reduced form models such as VARs. However, the literature contains papers with similar unobserved components models making similar points about their advantages (e.g. Stella and Stock, 2012). Relative to this literature, the main innovation of the present paper is the use of bounding of latent states such as trend inflation and NAIRU. The existing literature assumes these states evolve according to unbounded random walks, despite the fact that this assumption is inconsistent with much underlying economics (e.g. the fact that central bankers have implicit or explicit inflation targets). This paper develops a model which incorporates bounded random walks (where the bounds are estimated from the data). We find that this addition not only leads to more sensible estimates of trend inflation and NAIRU, but also forecasts better than a range of other approaches.

5 Technical Appendix

We remind the reader that our model is defined by:

$$\begin{array}{l} (\pi_t - \tau_t^\pi) = \rho_t^\pi \left(\pi_{t-1} - \tau_{t-1}^\pi\right) + \lambda_t \left(u_t - \tau_t^u\right) + \varepsilon_t^\pi \\ (u_t - \tau_t^u) = \rho_1^u \left(u_{t-1} - \tau_{t-1}^u\right) + \rho_2^u \left(u_{t-2} - \tau_{t-2}^u\right) + \varepsilon_t^u \\ \tau_t^\pi = \tau_{t-1}^\pi + \varepsilon_t^{\tau\pi} \\ \tau_t^u = \tau_{t-1}^u + \varepsilon_t^{\tau u} \\ \rho_t^\pi = \rho_{t-1}^\pi + \varepsilon_t^{\rho\pi} \\ \lambda_t = \lambda_{t-1} + \varepsilon_t^\lambda \end{array}, \quad (5)$$

and

$$\begin{array}{rcl} \varepsilon_t^\pi & \sim & N(0, \mathrm{e}^{h_t}) \\ h_t & = & h_{t-1} + \varepsilon_t^h, \\ \varepsilon_t^h & \sim & N(0, \sigma_h^2), \\ \varepsilon_t^u & \sim & N(0, \sigma_u^2). \end{array} \quad (6)$$

Trends are bounded through:

$$\begin{array}{l} \varepsilon_t^{\tau\pi} \sim TN(a_\pi - \tau_{t-1}^\pi, b_\pi - \tau_{t-1}^\pi; 0, \sigma_{\tau\pi}^2) \\ \varepsilon_t^{\tau u} \sim TN(a_u - \tau_{t-1}^u, b_u - \tau_{t-1}^u; 0, \sigma_{\tau u}^2) \end{array} \quad (7)$$

and time varying parameters through

$$\begin{array}{l} \varepsilon_t^{\rho\pi} \sim TN(-\rho_{t-1}^\pi, 1 - \rho_{t-1}^\pi; 0, \sigma_{\rho\pi}^2) \\ \varepsilon_t^\lambda \sim TN(-1 - \lambda_{t-1}, 0 - \lambda_{t-1}; 0, \sigma_\lambda^2). \end{array} \quad (8)$$

We also impose the stationary condition on the unemployment equation and assume $\rho_1^u + \rho_2^u < 1$, $\rho_2^u - \rho_1^u < 1$ and $|\rho_2^u| < 1$.

We use notation where $\pi = (\pi_1, \ldots, \pi_T)'$, $u = (u_1, \ldots, u_T)'$, $y = (\pi', u')'$, $\tau^\pi = (\tau_1^\pi, \ldots, \tau_T^\pi)'$, $\tau^u = (\tau_1^u, \ldots, \tau_T^u)'$, $\rho^\pi = (\rho_1^\pi, \ldots, \rho_T^\pi)'$, $\lambda = (\lambda_1, \ldots, \lambda_T)'$, $h = (h_1, \ldots, h_T)'$, and define $\varepsilon^\pi, \varepsilon^u, \varepsilon^{\tau\pi}, \varepsilon^{\tau u}, \varepsilon^h, \varepsilon^{\rho\pi}, \varepsilon^\lambda$ similarly. In addition, let θ denote the model parameters, i.e., $\theta = (\sigma_u^2, \sigma_{\tau\pi}^2, \sigma_{\tau u}^2, \sigma_h^2, \sigma_{\rho\pi}^2, \sigma_\lambda^2, a_\pi, b_\pi, a_u, b_u, \rho_1^u, \rho_2^u)'$.

5.1 The Prior

We require a prior for the initial condition in every state equation and these are:

$$\begin{aligned}
\tau_1^\pi &\sim TN(a_\pi, b_\pi; \tau_0^\pi, \omega_{\tau\pi}^2), \\
\tau_1^u &\sim TN(a_u, b_u; \tau_0^u, \omega_{\tau u}^2), \\
\rho_1^\pi &\sim TN(0, 1; \rho_0^\pi, \omega_{\rho\pi}^2), \\
\lambda_1 &\sim TN(-1, 0; \lambda_0, \omega_\lambda^2), \\
h_1 &\sim TN(h_0, \omega_h^2),
\end{aligned}$$

where τ_0^π, $\omega_{\tau\pi}^2$, τ_0^u, τ_{-1}^u, $\omega_{\tau u}^2$, ρ_0^π, $\omega_{\rho\pi}^2$, λ_0, $\omega_{\rho u}^2$, h_0, and ω_h^2 are known constants. In particular we choose the relatively non-informative values of $\tau_0^\pi = 3$, $\tau_0^u = \tau_{-1}^u = 5$, $h_0 = \rho_0^\pi = \lambda_0 = 0$, $\omega_{\tau\pi}^2 = \omega_{\tau u}^2 = \omega_h^2 = 5$ and $\omega_{\rho\pi}^2 = \omega_{\rho u}^2 = \omega_\lambda^2 = 1$. Note that the need for τ_{-1}^u arises from the use of an AR(2) specification in the unemployment equation.

The prior for the model parameters is specified as

$$p(\theta) = p(\sigma_u^2)p(\sigma_h^2)p(\sigma_{\tau\pi}^2)p(\sigma_{\tau u}^2)p(\sigma_{\rho\pi}^2)p(\sigma_\lambda^2)p\left(a_\pi\right)p\left(b_\pi\right)p\left(a_u\right)p\left(b_u\right)p\left(\rho_1^u\right)p\left(\rho_2^u\right),$$

where $\sigma_u^2 \sim IG(\underline{\nu}_u, \underline{S}_u)$, $\sigma_h^2 \sim IG(\underline{\nu}_h, \underline{S}_h)$, $\sigma_{\tau\pi}^2 \sim IG(\underline{\nu}_{\tau\pi}, \underline{S}_{\tau\pi})$, $\sigma_{\tau u}^2 \sim IG(\underline{\nu}_{\tau u}, \underline{S}_{\tau u})$, $\sigma_{\rho\pi}^2 \sim IG(\underline{\nu}_{\rho\pi}, \underline{S}_{\rho\pi})$, $\sigma_\lambda^2 \sim IG(\underline{\nu}_\lambda, \underline{S}_\lambda)$, and $IG(\cdot, \cdot)$ denotes the inverse-Gamma distribution. We choose relatively small values for the degrees of freedom parameters, which imply large prior variances, i.e., the priors are relatively non-informative. Specifically, we set $\underline{\nu}_u = \underline{\nu}_h = \underline{\nu}_{\tau\pi} = \underline{\nu}_{\tau u} = \underline{\nu}_{\rho\pi} = \underline{\nu}_{\rho u} = \underline{\nu}_\lambda = 10$. We then choose values for the scale parameters so that the parameters have the desired prior means. We set $\underline{S}_u = \underline{S}_h = 0.9$, which imply prior means $E(\sigma_u^2) = E(\sigma_h^2) = 0.1$. Next, we set $\underline{S}_{\tau\pi} = 0.18$ and $\underline{S}_{\tau u} = 0.09$, which imply $E(\sigma_{\tau\pi}^2) = 0.02$ and $E(\sigma_{\tau u}^2) = 0.01$. These values are chosen to reflect the desired smoothness of the corresponding state transition. For example, the prior mean for $\sigma_{\tau\pi}^2$ implies that with high probability the difference between consecutive trend inflation, $\tau_t^\pi - \tau_{t-1}^\pi$, lies within the values -0.3 and 0.3. We set $\underline{S}_{\rho\pi} = \underline{S}_\lambda = 0.018$, which imply prior means $E(\sigma^2_{\rho\pi}) = E(\sigma^2_{\rho u}) = E(\sigma^2_{\lambda}) = 0.002$. These values imply a relatively smooth transition for the relevant states.

For the bounds we use uniform priors: $a_\pi \sim U(0,2), b_\pi \sim U(3,5), a_u \sim U(3,5), b_u \sim U(6,8)$. The priors for ρ_1^u and ρ_2^u are jointly normal with mean $(1.8, -0.8)'$ and covariance matrix $5I_2$.

5.2 MCMC Algorithm

We extend the MCMC sampler developed in Chan, Koop and Potter (2013) which in turn is an adaptation of the algorithm introduced in Chan and Strachan (2012).

Specifically, we sequentially draw from (we suppress the dependence on π_0, u_0 and u_{-1}):

1. $p(\tau^\pi \,|\, y, \tau^u, \rho^\pi, \lambda, h, \theta)$
2. $p(\tau^u \,|\, y, \tau^\pi, \rho^\pi, \lambda, h, \theta)$
3. $p(\rho^\pi \,|\, y, \tau^\pi, \tau^u, \lambda, h, \theta)$
4. $p(\lambda \,|\, y, \tau^\pi, \tau^u, \rho^\pi, h, \theta)$
5. $p(h \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, \theta)$
6. $p\left(\theta | y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h\right)$

Step 1: To derive the conditional distribution $p(\tau^\pi \,|\, y, \tau^u, \rho^\pi, \lambda, h, \theta)$, we first rewrite the inflation equation as

$$K_\pi \pi = \mu_\pi + K_\pi \tau^\pi + \varepsilon^\pi, \quad \varepsilon^\pi \sim N(0, \Omega_\pi),$$

where $\Omega_\pi = \mathrm{diag}(\mathrm{e}^{h_1}, \ldots, \mathrm{e}^{h_T})$ and

$$\mu_\pi = \begin{pmatrix} \rho_1^\pi(\pi_0 - \tau_0^\pi) + \lambda_1(u_1 - \tau_1^u) \\ \lambda_2(u_2 - \tau_2^u) \\ \lambda_3(u_3 - \tau_3^u) \\ \vdots \\ \lambda_T(u_T - \tau_T^u) \end{pmatrix}, \quad K_\pi = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -\rho_2^\pi & 1 & 0 & \cdots & 0 \\ 0 & -\rho_3^\pi & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \cdots & -\rho_T^\pi & 1 \end{pmatrix}.$$

Since $|K_\pi| = 1$ for any ρ^π, K_π is invertible. Therefore, we have

$$(\pi \,|\, u, \tau^u, \rho^\pi, \lambda, h, \theta) \sim N(K_\pi^{-1}\mu_\pi + \tau^\pi, (K_\pi' \Omega_\pi^{-1} K_\pi)^{-1}),$$

i.e.,

$$\log p(\pi \mid u, \tau^u, \rho^\pi, \lambda, h, \theta) \propto -\frac{1}{2}\iota_T' h - \frac{1}{2}(\pi - K_\pi^{-1}\mu_\pi - \tau^\pi)' K_\pi' \Omega_\pi^{-1} K_\pi (\pi - K_\pi^{-1}\mu_\pi - \tau^\pi), \quad (9)$$

where ι_T is a $T \times 1$ column of ones. Similarly, rewrite the state equation for τ^π as

$$H\tau^\pi = \alpha_\pi + \varepsilon^{\tau\pi},$$

where

$$\alpha_\pi = \begin{pmatrix} \tau_0^\pi \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad H = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \cdots & -1 & 1 \end{pmatrix}.$$

That is, the prior density for τ^π is given by

$$\log p(\tau^\pi \mid \sigma^2_{\tau\pi}) \propto \frac{1}{2}(\tau^\pi - H^{-1}\alpha_\pi)' H' \Omega_{\tau\pi}^{-1} H (\tau^\pi - H^{-1}\alpha_\pi) + g_{\tau\pi}(\tau^\pi, \sigma^2_{\tau\pi}), \quad (10)$$

where $a_\pi < \tau_t^\pi < b_\pi$ for $t = 1, \dots, T$, $\Omega_{\tau\pi} = \text{diag}(\omega^2_{\tau\pi}, \sigma^2_{\tau\pi}, \dots, \sigma^2_{\tau\pi})$ and

$$\begin{aligned} g_{\tau\pi}(\tau^\pi, \sigma^2_{\tau\pi}) = & -\log\left(\Phi\left(\frac{b_\pi}{\omega_{\tau\pi}}\right) - \Phi\left(\frac{a_\pi}{\omega_{\tau\pi}}\right)\right) \\ & - \sum_{t=2}^{T} \log\left(\Phi\left(\frac{b_\pi - \tau^\pi_{t-1}}{\sigma_{\tau\pi}}\right) - \Phi\left(\frac{a_\pi - \tau^\pi_{t-1}}{\sigma_{\tau\pi}}\right)\right). \end{aligned}$$

Combining (9) and (10), we obtain

$$\begin{aligned} \log p(\tau^\pi \mid y, \tau^u, \rho^\pi, \rho^u, \lambda, h, \theta) \\ \propto & -\frac{1}{2}(\pi - K_\pi^{-1}\mu_\pi - \tau^\pi)' K_\pi' \Omega_\pi^{-1} K_\pi (\pi - K_\pi^{-1}\mu_\pi - \tau^\pi) \\ & -\frac{1}{2}(\tau^\pi - H^{-1}\alpha_\pi)' H' \Omega_{\tau\pi}^{-1} H (\tau^\pi - H^{-1}\alpha_\pi) + g_{\tau\pi}(\tau^\pi, \sigma^2_{\tau\pi}), \\ \propto & -\frac{1}{2}(\tau^\pi - \hat{\tau}^\pi)' D_{\tau\pi}^{-1} (\tau^\pi - \hat{\tau}^\pi) + g_{\tau\pi}(\tau^\pi, \sigma^2_{\tau\pi}), \end{aligned}$$

where $a_\pi < \tau_t^\pi < b_\pi$ for $t = 1, \dots, T$, and

$$\begin{aligned} D_{\tau\pi} &= \left(H'\Omega_{\tau\pi}^{-1} H + K_\pi' \Omega_\pi^{-1} K_\pi\right)^{-1}, \\ \hat{\tau}^\pi &= D_{\tau\pi}(H'\Omega_{\tau\pi}^{-1}\alpha_\pi + K_\pi'\Omega_\pi^{-1} K_\pi(\pi - K_\pi^{-1}\mu_\pi)). \end{aligned}$$

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

Step 2: To implement Step 2, note that information about τ^u comes from three sources: the two measurement equations for inflation and unemployment and the state equation for τ^u. We derive an expression for each component in turn. First, write the inflation equation as:

$$z = \Lambda\tau^u + \varepsilon^\pi, \quad \varepsilon^\pi \sim N(0, \Omega_\pi),$$

where $z_t = (\pi_t - \tau_t^\pi) - \rho_t^\pi(\pi_{t-1} - \tau_{t-1}^\pi) - \lambda_t u_t$, $z = (z_1, \dots, z_T)'$, and $\Lambda = \text{diag}(-\lambda_1, \dots, -\lambda_T)$. Hence, ignoring any terms not involving τ^u, we have

$$\log p(\pi \,|\, u, \tau^u, \rho^\pi, \lambda, h, \theta) \propto -\frac{1}{2}(z - \Lambda\tau^u)'\Omega_\pi^{-1}(z - \Lambda\tau^u). \qquad (11)$$

The second component comes from the unemployment equation, which can be written as:

$$K_u u = \mu_u + K_u\tau^u + \varepsilon^u, \quad \varepsilon^u \sim N(0, \Omega_u),$$

where $\Omega_u = I_T \otimes \sigma_u^2$ and

$$\mu_u = \begin{pmatrix} \rho_1^u(u_0 - \tau_0^u) + \rho_2^u(u_{-1} - \tau_{-1}^u) \\ \rho_2^u(u_0 - \tau_0^u) \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \quad K_u = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ -\rho_1^u & 1 & 0 & 0 & \cdots & 0 \\ -\rho_2^u & -\rho_1^u & 1 & 0 & \cdots & 0 \\ 0 & -\rho_2^u & -\rho_1^u & 1 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & -\rho_2^u & -\rho_1^u & 1 \end{pmatrix}.$$

Thus, ignoring any terms not involving τ^u, we have

$$\log p(u \,|\, \tau^u, \theta) \propto -\frac{1}{2}(u - K_u^{-1}\mu_u - \tau^u)'K_u'\Omega_u^{-1}K_u(u - K_u^{-1}\mu_u - \tau^u). \quad (12)$$

The third component is contributed by the state equation for τ^u :

$$\log p(\tau^u \,|\, \sigma_{\tau u}^2) \propto -\frac{1}{2}(\tau^u - H^{-1}\alpha_u)'H'\Omega_{\tau u}^{-1}H(\tau^u - H^{-1}\alpha_u) + g_{\tau u}(\tau^u, \sigma_{\tau u}^2), \qquad (13)$$

where $\alpha_u = (\tau_0^u, 0, \ldots, 0)'$, $a_u < \tau_t^u < b_u$ for $t = 1, \ldots, T$, $\Omega_{\tau u} = \text{diag}(\omega_{\tau u}^2, \sigma_{\tau u}^2, \ldots, \sigma_{\tau u}^2)$ and

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

Combining (11), (12) and (13), we obtain

$$\begin{aligned}
&\log p(\tau^u \,|\, y, \tau^\pi, \rho_1^u, \rho_2^u, \lambda, h, \theta) \\
&\quad \propto -\frac{1}{2}(z - \Lambda\tau^u)'\Omega_\pi^{-1}(z - \Lambda\tau^u) \\
&\qquad -\frac{1}{2}(u - K_u^{-1}\mu_u - \tau^u)'K_u'\Omega_u^{-1}K_u(u - K_u^{-1}\mu_u - \tau^u) \\
&\qquad -\frac{1}{2}(\tau^u - H^{-1}\alpha_u)'H'\Omega_{\tau u}^{-1}H(\tau^u - H^{-1}\alpha_u) + g_{\tau u}(\tau^u, \sigma_{\tau u}^2), \\
&\quad \propto -\frac{1}{2}(\tau^u - \hat{\tau}^u)'D_{\tau u}^{-1}(\tau^u - \hat{\tau}^u) + g_{\tau u}(\tau^u, \sigma_{\tau u}^2),
\end{aligned}$$

where $a_u < \tau_t^u < b_u$ for $t = 1, \ldots, T$, and

$$\begin{aligned}
D_{\tau u} &= \left(H'\Omega_{\tau u}^{-1}H + K_u'\Omega_u^{-1}K_u + \Lambda'\Omega_\pi^{-1}\Lambda\right)^{-1}, \\
\hat{\tau}^u &= D_{\tau u}(H'\Omega_{\tau u}^{-1}\alpha_u + K_u'\Omega_u^{-1}K_u(u - K_u^{-1}\mu_u) + \Lambda'\Omega_\pi^{-1}z).
\end{aligned}$$

Again, we sample τ^u via an ARMH step with candidate draws obtained from $N(\hat{\tau}^u, D_{\tau u})$.

Step 3: Next, we derive an expression for $p(\rho^\pi \,|\, y, \tau^\pi, \tau^u, \lambda, h, \theta)$. First, let $\pi_t^* = \pi_t - \tau_t^\pi$, $u_t^* = u_t - \tau_t^u$, $\pi^* = (\pi_1^*, \ldots, \pi_T^*)'$, and $u^* = (u_1^*, \ldots, u_T^*)'$. Then the measurement equation for inflation can be rewritten as

$$\pi^* + \Lambda u^* = X_\pi \rho^\pi + \varepsilon^\pi, \quad \varepsilon^\pi \sim N(0, \Omega_\pi),$$

where $X_\pi = \text{diag}(\pi_0^*, \ldots, \pi_{T-1}^*)$ and $\Lambda = \text{diag}(-\lambda_1, \ldots, -\lambda_T)$. From the state equation for ρ^π we also have

$$H\rho^\pi = \varepsilon^{\rho\pi}.$$

Therefore, using a similar argument as before, we have

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

where $0 < \rho^\pi_t < 1$ for $t = 1, \ldots, T$,

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

$$D_{\rho\pi} = \left(H'\Omega^{-1}_{\rho\pi}H + X'_\pi\Omega^{-1}_\pi X_\pi\right)^{-1}, \quad \hat{\rho}^\pi = D_{\rho\pi}X'_\pi\Omega^{-1}_\pi(\pi^* + \Lambda u^*),$$

and $\Omega_{\rho\pi} = \mathrm{diag}(\omega^2_{\rho\pi}, \sigma^2_{\rho\pi}, \ldots, \sigma^2_{\rho\pi})$. As before, we implement an ARMH step with approximating density $N(\hat{\rho}^\pi, D_{\rho\pi})$.

Step 4: Using the same argument as before, we have

$$p(\lambda \,|\, y, \tau^\pi, \tau^u, \rho^\pi, h, \theta) \propto -\frac{1}{2}(\lambda - \hat{\lambda})'D^{-1}_\lambda(\lambda - \hat{\lambda}) + g_\lambda(\lambda, \sigma^2_\lambda),$$

where $-1 < \lambda_t < 0$ for $t = 1, \ldots, T$,

$$g_\lambda(\lambda, \sigma^2_\lambda) = -\sum_{t=2}^{T} \log\left(\Phi\left(\frac{-\lambda_{t-1}}{\sigma_\lambda}\right) - \Phi\left(\frac{-1-\lambda_{t-1}}{\sigma_\lambda}\right)\right),$$

$$D_\lambda = \left(H'\Omega^{-1}_\lambda H + X'_u\Omega^{-1}_\pi X_u\right)^{-1}, \quad \hat{\lambda} = D_\lambda X'_u\Omega^{-1}_\pi w,$$

$X_u = \mathrm{diag}(u^*_0, \ldots, u^*_{T-1})$, $w = (\pi^*_1 - \rho^\pi_1\pi^*_0, \ldots, \pi^*_T - \rho^\pi_T\pi^*_{T-1})'$, and $\Omega_\lambda = \mathrm{diag}(\omega^2_\lambda, \sigma^2_\lambda, \ldots, \sigma^2_\lambda)$. As before, we implement an ARMH step with approximating density $N(\hat{\rho}^u, D_{\rho u})$.

Step 5: For Step 5, we use the algorithm in Chan and Strachan (2012) to sample from $p(h \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, \theta)$.

Step 6: We draw from θ in separate blocks, mainly using standard results for the regression model. We use notation where θ_{-x} for all parameters in θ except for x.

Using standard linear regression results, it can be shown that $\rho^u = (\rho^u_1, \rho^u_2)'$ is a bivariate truncated normal:

$$p(\rho^u \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\rho^u}) \propto -\frac{1}{2}(\rho^u - \hat{\rho}^u)'D^{-1}_{\rho u}(\rho^u - \hat{\rho}^u) + g_{\rho u}(\rho^u),$$

with the stationarity constraints $\rho^u_1 + \rho^u_2 < 1$, $\rho^u_2 - \rho^u_1 < 1$ and $|\rho^u_2| < 1$, where

$$D_{\rho u} = \left(V^{-1}_{\rho u} + X'_u X_u/\sigma^2_u\right)^{-1}, \quad \hat{\rho}^u = D_{\rho u}X'^*_u/\sigma^2_u, \quad X_u = \begin{pmatrix} u^*_0 & u^*_{-1} \\ u^*_1 & u^*_0 \\ \vdots & \vdots \\ u^*_{T-1} & u^*_{T-2} \end{pmatrix}.$$

A draw from this truncated normal distribution can be obtained via acceptance-rejection sampling with proposal from $N(\hat{\rho}^u, D_{\rho u})$.

To sample from the error variances, first note that they are conditionally independent given the data and the states. Hence, we can sample each element one by one.

Now, both $p(\sigma^2_u \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\sigma^2_u})$ and $p(\sigma^2_h \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\sigma^2_h})$ are standard inverse-Gamma densities, respectively:

$$(\sigma^2_u \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\sigma^2_u}) \sim IG\left(\underline{\nu}_u + \frac{T}{2}, \underline{S}_u + \frac{1}{2}\sum_{t=1}^{T}(u^*_t - \rho^u_1 u^*_{t-1} - \rho^u_2 u^*_{t-2})^2\right)$$

$$(\sigma^2_h \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \rho^u, \lambda, h, \theta_{-\sigma^2_h}) \sim IG\left(\underline{\nu}_h + \frac{T-1}{2}, \underline{S}_h + \frac{1}{2}\sum_{t=2}^{T}(h_t - h_{t-1})^2\right).$$

Next, the log conditional density for $\sigma^2_{\tau\pi}$ is given by

$$\log p(\sigma^2_{\tau\pi} \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\sigma^2_{\tau\pi}}) \propto \\ -(\underline{\nu}_{\tau\pi} + 1)\log\sigma^2_{\tau\pi} - \frac{\underline{S}_{\tau\pi}}{\sigma^2_{\tau\pi}} - \frac{T-1}{2}\log\sigma^2_{\tau\pi} - \frac{1}{2\sigma^2_{\tau\pi}}\sum_{t=2}^{T}(\tau^\pi_t - \tau^\pi_{t-1})^2 + g_{\tau\pi}(\tau^\pi, \sigma^2_{\tau\pi}),$$

which is a non-standard density. To proceed, we implement an MH step with the proposal density

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

Similarly, the log conditional density for $\sigma^2_{\tau u}$ is given by

$$\log p(\sigma^2_{\tau u} \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\sigma^2_{\tau u}}) \propto \\ -(\underline{\nu}_{\tau u} + 1)\log\sigma^2_{\tau u} - \frac{\underline{S}_{\tau u}}{\sigma^2_{\tau u}} - \frac{T-1}{2}\log\sigma^2_{\tau u} - \frac{1}{2\sigma^2_{\tau u}}\sum_{t=2}^{T}(\tau^u_t - \tau^u_{t-1})^2 + g_{\tau u}(\tau^u, \sigma^2_{\tau u}).$$

Again, a draw from $p(\sigma^2_{\tau u} \,|\, y, \tau^\pi, \tau^u, \rho^\pi, \lambda, h, \theta_{-\sigma^2_{\tau u}})$ is obtained via an MH step with the proposal density

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

The remaining error variances, $\sigma^2_{\rho\pi}$ and σ^2_λ, are sampled analogously.

To draw from the bounds a_π, b_π, a_u and b_u, we use a Griddy-Gibbs sampler which is the same as the one used in Chan, Koop and Potter (2013). Specifically, since each of the bounds has finite support, we can evaluate its conditional density on a fine grid, which can then be used to construct the associated cumulative distribution function. Finally, a draw from the conditional density can be obtained via the inverse-transform method. The reader is referred to our earlier paper for details.

5.3 Specification and Estimation Details for Other Models

The other models used in our forecast comparisons are mostly restricted special cases of **Bi-UC** and all specification and prior details are identical to **Bi-UC** except that the relevant restriction is imposed. The exceptions to this are discussed in this sub-section.

The VAR(2) is specified as:

$$y_t = \mu + B_1 y_{t-1} + B_2 y_{t-2} + \varepsilon_t,$$

where $y_t = (\pi_t, u_t)'$ and $\varepsilon_t = (\varepsilon_t^\pi, \varepsilon_t^u)' \sim N(0, \Sigma)$. We use a relatively noninformative prior. For μ the prior is $N(0, 100)$. For the VAR coefficients, we assume each is N(0,1) and all are, a priori, uncorrelated with one another. The prior for Σ is $IW\left(10, \begin{bmatrix} 1.4 & 0 \\ 0 & 0.7 \end{bmatrix}\right)$ so that the prior mean of the error variances in the two equations are 0.2 and 0.1, respectively.

We also use a VAR(2) with stochastic volatility:

$$y_t = \mu + B_1 y_{t-1} + B_2 y_{t-2} + A^{-1}\varepsilon_t,$$

where $\varepsilon_t^\pi \sim N(0, \mathrm{e}^{h_t})$, $\varepsilon_t^u \sim N(0, \sigma_u^2)$ and

$$A = \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix}.$$

The VAR coefficients have the same prior as the VAR(2) without stochastic volatility. All details (including prior) relating to h_t and σ_u^2 are exactly as in **Bi-UC** and a has a $N(0, 10)$ prior.

The VARs are estimated using MCMC methods as outlined, e.g., in Koop and Korobilis (2009).

For **Bi-UC-TVP-**ρ^u, all details are identical to **Bi-UC** except that we need to specify the initial conditions for ρ_{1t}^u and ρ_{2t}^u which are now time-varying. These are both assumed to be $N(0, 5)$. The priors for the two error variances in the two state equations are both $IG(10, 0.009)$, a relatively noninformative choice implying prior means of 0.001.

The VAR(2) with Minnesota Prior model is specified as follows. Let $\beta = \text{vec}((\mu, B_1, B_2)')$ and $X_t = I_2 \otimes (1, y'_{t-1}, y'_{t-2})$, and let $\widehat{b}$ denote the least squares estimator of β. We fix $\Sigma = \widehat{\Sigma}$, where $\widehat{\Sigma} = T^{-1} \sum_{t=1}^{T} (y_t - X_t\widehat{b})'(y_t - X_t\widehat{b})$. For β, we consider the prior $\beta \sim N(0, \underline{V}_{\text{mn}})$, where

$$\underline{V}_{\text{mn}} = \text{diag}\left(\underline{a}_3 s_1^2, \underline{a}_1, \frac{\underline{a}_2 s_1^2}{s_2^2}, \frac{\underline{a}_1}{4}, \frac{\underline{a}_2 s_1^2}{4s_2^2}, \underline{a}_3 s_2^2, \frac{\underline{a}_2 s_2^2}{s_1^2}, \underline{a}_1, \frac{\underline{a}_2 s_2^2}{4s_1^2}, \frac{\underline{a}_1}{4}\right)$$

and s_i^2 is the i-th diagonal element of $\widehat{\Sigma}$, $i = 1, 2$. We set $\underline{a}_1 = 0.1$, $\underline{a}_2 = 0.05$, and $\underline{a}_3 = 0.1$.

The bivariate random walk model is:

$$\begin{aligned}\pi_t &= \pi_{t-1} + \epsilon_t^{\pi},\\ u_t &= u_{t-1} + \epsilon_t^{u},\end{aligned}$$

where $\epsilon_t^{\pi} \sim N(0, \sigma_{\pi}^2)$ and $\epsilon_t^{u} \sim N(0, \sigma_u^2)$ are independent.

This model is a special case of **Bi-UC**, where $\rho_1^u = 1$, $\rho_2^u = 0$, $\tau_t^{\pi} = \tau_t^u = \lambda_t = 0$, $\rho_t^{\pi} = 1$ for $t = 1, \ldots, T$, and the errors are independent and homoskedastic.

The UCSV-AR(2) model is:

$$\begin{aligned}\pi_t &= \tau_t^{\pi} + \epsilon_t^{\pi},\\ u_t &= \rho_1^u u_{t-1} + \rho_2^u u_{t-2} + \epsilon_t^{u},\end{aligned}$$

where $\epsilon_t^{\pi} \sim N(0, e^{h_t})$ and $\epsilon_t^{u} \sim N(0, \sigma_u^2)$ are independent. The inflation trend τ_t^{π} and log-volatility h_t follow independent random walks. This model is a special case of **Bi-UC**, where $\tau_t^u = \lambda_t = \rho_t^{\pi} = 0, = 1$ for $t = 1, \ldots, T$.

5.4 Prior Predictive Analysis

To convince the reader of the sensibility of our model and prior, this subsection presents results from a prior predictive analysis.

We begin by computing the predictive densities for future trend inflation and future NAIRU, τ_{T+k}^{π} and τ_{T+k}^{u} respectively, with $k = 20$ under the **Bi-UC** model. The results are reported in Figure 5.4 and 5.4. These figures should that the predictive densities produced by our model are sensible and show the role of the bounds tightening these densities in a sensible manner.

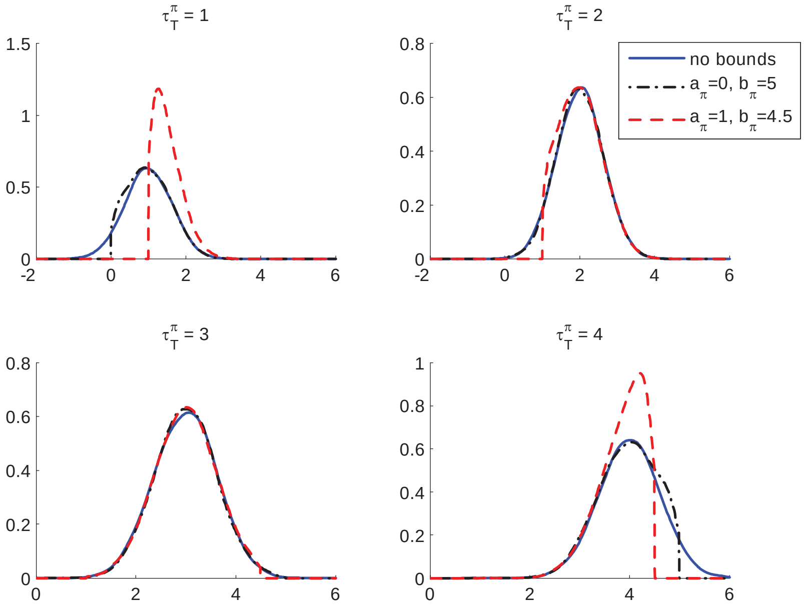

Figure 1: Predictive densities for τ^{π}_{T+k} under the **Bi-UC** model where $k = 20$.

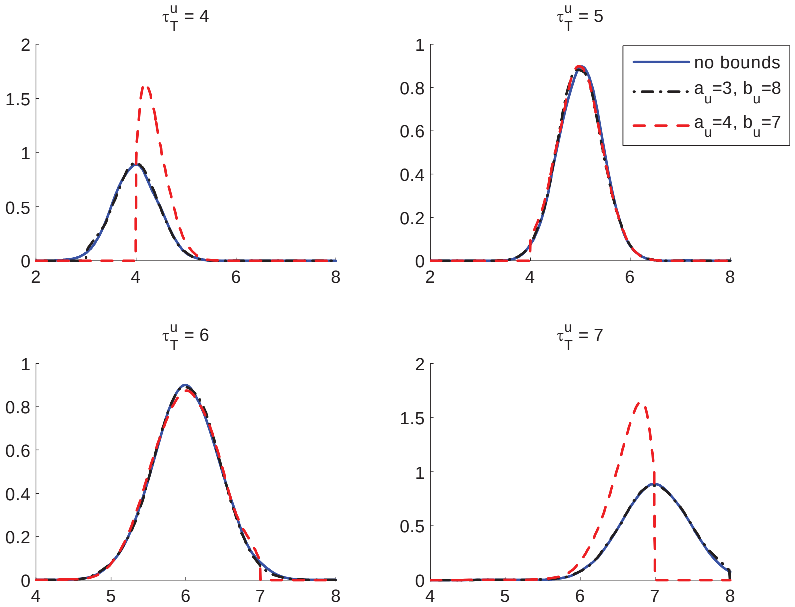

Figure 2: Predictive densities for τ^u_{T+k} under the **Bi-UC** model where $k = 20$.

Next, we preform a prior predictive analysis. This involves taking a draw from the prior (using the prior described in the Technical Appendix) and then simulating from the state equations. Given the drawn parameters, states and an initial value for inflation and the unemployment rate (we set $\pi_0 = 3$ and $u_0 = u_{-1} = 5.5$), an artificial dataset of inflation and unemployment can be generated. This is repeated 10^4 times and, for each generated data set, we compute various features of interest such as quantiles, variance, autocorrelations, etc. in order to build up the prior cumulative distribution function (cdf) for each of these features. Tables 1 and 2 present these cdf's evaluated at the feature of interest calculated for the observed data set. It can be seen that all of the features of the data can be well explained by our model. Our model does worst at explaining the autocorrelation patterns in the unemployment rate series. However, even for this case, it does as well as an unbounded version of our model.

To formally compare **Bi-UC** and **Bi-UC-NoBound**, we compute the log Bayes factors using the prior predictive densities for various combinations

Table 6: Prior cdfs of features for π_t.

Feature	**Bi-UC**	**Bi-UC-NoBound**
16%-tilde	0.576	0.574
median	0.557	0.547
84%-tilde	0.663	0.585
variance	0.598	0.520
fraction of $\pi_t < 0$	0.431	0.417
fraction of $\pi_t > 10$	0.687	0.620
lag 1 autocorrelation	0.634	0.706
lag 4 autocorrelation	0.726	0.581
MA coefficient	0.539	0.713

Table 7: Prior cdfs of features for u_t.

Feature	**Bi-UC**	**Bi-UC-NoBound**
16%-tilde	0.694	0.708
median	0.715	0.721
84%-tilde	0.725	0.703
variance	0.626	0.587
fraction of $u_t < 4$	0.273	0.261
fraction of $u_t > 8$	0.746	0.719
lag 1 autocorrelation	0.965	0.965
lag 4 autocorrelation	0.922	0.903
MA coefficient	0.765	0.764

of the features of interest considered in Tables 1 and 2. In particular, 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 $y_t > 10$ for $y_t = \pi_t$ or u_t), and "Dynamics" include the last three features of interest (lag 1 autocorrelation, lag 4 autocorrelation and MA coefficient). The results are reported in Table 8. The fact that the inclusion of bounds leads to a more parsimonious model (without causing the fit of the model to deteriorate), leads to log Bayes factors strongly in favor of our bounded model for inflation. For the unemployment rate, the bounds play less of a role and our model is performing roughly as well as its unbounded version.

Table 8: Log Bayes factors in favor of **Bi-UC** against **Bi-UC-NoBound**.

	Quantile	Spread and drift	Dynamics	All
π_t	286.67	328.76	1.45	613.85
u_t	0.47	0.34	-0.53	0.15

References

Atkeson, A. and Ohanian, L. (2001). "Are Phillips curves useful for forecasting inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review*, 25(1), 2-11.

Carruth, A., Hooker, M. and Oswald, A. (1998). "Unemployment equilibria and input prices: Theory and evidence from the United States," *The Review of Economics and Statistics*, 80, 621-628.

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 2012-13.

Chan, J., Koop, G. and Potter, S. M. (2013). "A new model of trend inflation," *Journal of Business and Economic Statistics*, 31, 94-106.

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.

D’Agostino, A., Giannone, D. and Gambetti, L. (2013). “Macroeconomic forecasting and structural change,” *Journal of Applied Econometrics*, 28, 81-101.

Dotsey, M., Fujita, S. and Stark, T. (2010). ”Do Phillips curves conditionally help to forecast inflation?”, manuscript, Federal Reserve Bank of Philadelphia.

Ireland, P. (2007). “Changes in the Federal Reserve’s inflation target: Causes and consequences,” *Journal of Money, Credit and Banking*, 39, 1851-1882.

Kim, S., Shephard, N. and Chib, S. (1998). “Stochastic volatility: likelihood inference and comparison with ARCH models,” *Review of Economic Studies*, 65, 361-93.

Koop, G. and Korobilis, D. (2009). “Bayesian multivariate time series methods for empirical macroeconomics,” *Foundations and Trends in Econometrics*, 3, 267-358.

Montgomery, A., Zarnowitz, V. Tsay, R. and Tiao, G. (1998). “Forecasting the U.S. unemployment rate,” *Journal of the American Statistical Association*, 92, 478-
493.

Primiceri. G. (2005). “Time varying structural vector autoregressions and monetary policy,” *Review of Economic Studies*, 72, 821-852.

Sims, C. and Zha, T. (2006). “Were there regime switches in macroeconomic policy?” *American Economic Review*, 96, 54-81.

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.

Staiger, D., Stock, J. and Watson, M. (1997). “The NAIRU, unemployment and monetary policy,” *Journal of Economic Perspectives*, 11, 33-49.

Stella, A. and Stock, J. (2012). “A state-dependent model for inflation forecasting,” manuscript available at
http://scholar.harvard.edu/files/stock/files/statedependentmodelinflation.pdf.

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.