

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

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

Simon M. Potter
Federal Reserve Bank of New York

November 19, 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, NAIRU 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. Joshua Chan would like to acknowledge financial support by the Australian Research Council via a Discovery Early Career Researcher Award (DE150100795).

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, 2013). 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, 2013 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, 2013 and many others). Furthermore, the researcher may wish to provide joint projections or forecasts, based on a multivariate predictive density, as opposed to producing marginal forecasts based on univariate densities.

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, Cogley, Primiceri and Sargent, 2010 and D’Agostino, Giannone and Gambetti, 2013).

The preceding discussion suggests researchers should be interested in models with three characteristics. First, they should be bivariate models of inflation and unemployment. Second, some coefficients should change over time. Third, they should be written in terms of latent state vectors that can be given an economic interpretation. Such a model is developed in Stella and Stock (2013), 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 (2013) is in its treatment of the latent states. Following most of the existing literature, Stella and Stock (2013) 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, Ireland, 2007, Stock and Watson, 2007 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. For NAIRU bounding the counterfactual implications of unrestricted movements in the random walk is perhaps more important. With the exception of abrupt changes in employment law or unemployment benefits, which are observable events, one would expect the forces determining the NAIRU to be slow moving and not lead to declines of unemployment to levels close to zero or levels above previous peaks in the unemployment rate driven by recessions. This is particularly the case because the unemployment rate by construction is a bounded variable and one would expect long-run equilibrium in the labor market to produce strong restrictions on its movement within this bounded interval. Further, by imposing the unrestricted random walk specification on NAIRU researchers by construction add excess uncertainty to the possible location of NAIRU in the past, present and future. Building on our previous work, Chan, Koop and Potter (2013), we embed this bounded model of NAIRU within a structure where the central bank keeps trend inflation well-contained, allowing us to better discriminate between cyclical and trend movements in inflation.

We also allow for time variation in important parameters such as those controlling the slope of the Phillips curve and the degree of persistence in deviations of inflation from trend. We place bounds on the degree of time variation to ensure that the statistical model does not imply explosive or counter-intuitive behavior. We find such bounding to be of importance in obtaining reasonable estimates and forecasts from time-varying parameter models such as the one used in this paper.

Models such as the UCSV or the one in Stella and Stock (2013) 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 (2013) 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. We argue that inequality restrictions such as our bounds can be another source of information which can lessen the need for such tight priors or equality restrictions.

The fact that our latent states are bounded implies our model is not a Gaussian 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:

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

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 (2013) 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., Primiceri (2005) and Stella and Stock (2013). Stella and Stock (2013) 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 (2013) 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, since it allows 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 a 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 (2013) 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,

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

Further, based on preliminary empirical work and following Stella and Stock (2013) 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 (2013) 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 $\tau_t^\pi, \tau_t^u, \rho_t^\pi$ 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 and well-contained long-run inflation expectations. Similarly, unbounded random walk behavior of the NAIRU is not consistent with much macroeconomic theory about the slow moving and bounded forces driving NAIRU and thus, the Fed's maximum employment mandate. To capture these restrictions, 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

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

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, 2013) 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

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

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 noting that, for policy purposes, out-of-sample forecasting or projections involving the joint distribution of unemployment and inflation are often used. For instance, recently several central banks have made statements of the general form: “we will not raise interest rates at least until the unemployment rate falls below a certain threshold, provided inflation has not risen above a certain threshold”.¹ And central banks increasingly present information in a probabilistic format about such events. See, for instance, the projections in Charts 2 and 4 (for inflation and unemployment, respectively) from the Bank of England’s August 2013 Inflation Report.

In this spirit, we first present out-of-sample joint forecasts and projections of the sort a policymaker might be interested in. Given the data up to 2013Q1 and using our **Bi-UC** model, we compute the joint density forecasts of inflation and unemployment for 2013Q2, 2013Q3 and 2013Q4. We also construct the highest density regions of probabilities 0.2, 0.4, 0.6 and 0.8. The results are reported in Figure 1.




Figure 1: Joint predictive density of inflation and unemployment for 2013Q2, 2013Q3 and 2013Q4. The concentric "ellipses" represent the highest density regions of probabilities 0.2, 0.4, 0.6 and 0.8, whereas the triangle denotes the actual observation.

It can be seen that, using bounding, these densities can be elongated or skewed. The realizations do occur in regions of high probability.

¹Both the Federal Reserve and Bank of England expressed their state contingent policies in terms of projected inflation. Since these were internal projections and difficult to replicate we use a realization of inflation over one/four quarters of 2 or 3 percent. This is similar to the so-called Evans rule, see Campbell, Evans, Fisher and Justiniano (2012).

In the top panels of Figure 2, we compute the probabilities that the inflation and unemployment rates in 2013Q2-2016Q1 (a 3 year period) are above chosen thresholds. For inflation, these thresholds are 2% and 3%. For the unemployment rate they are 6% and 6.5%. To see the effect of bounding, we present results for our **Bi-UC** model as well as **Bi-UC-NoBound** which does not impose bounds on any of the states.

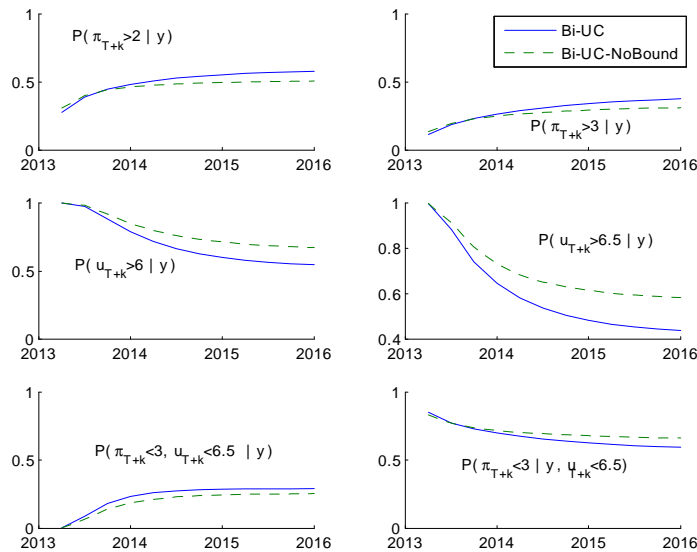


Figure 2: Marginal, joint and conditional probabilities involving inflation and unemployment rates

It can be seen that our projections are sensible, indicating an increasing probability that unemployment will fall below threshold and inflation rise above threshold. But note the effect of bounding. With our model, these results are coming through more strongly. Without bounding, results can differ substantially, particularly at longer horizons. For instance, without bounding the projections say that the probability that unemployment will remain above 6.5% in 2016Q1 is over 60%, but with bounding this number drops to almost 40%.

An advantage of a bivariate model is that joint and conditional projections involving both variables can be produced. To illustrate this, the bottom left panel of Figure 2 presents the joint probability of an event that might interest central bankers: the joint probability that unemployment and inflation

are both below threshold. The bottom right panel of Figure 2 presents a conditional probability: the probability that the inflation rate is below threshold given unemployment is below threshold. In this figure, by comparing **Bi-UC** to **Bi-UC-NoBound** estimates, we can also see the effects of bounding.

We next investigate how well our **Bi-UC** model would have forecast in the past relative to a range of alternative models. These are mostly bivariate models for inflation and unemployment. However, in order to compare our methods to those of our previous univariate work on inflation forecasting, Chan, Koop and Potter (2013), we also include two univariate forecasting models for inflation. All forecasts are pseudo-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. 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 (2013), 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- $-\tau_t^\pi$ - 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 according to random walks.
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 (2013).
15. **AR-Trend-bound**: the univariate model of Chan, Koop and Potter (2013) with bounds on trend inflation and the AR process.²
16. **Trend-bound**: restricts **AR-Trend-bound** to remove the AR process for inflation so that it is the UCSV model with bounded trend inflation.

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 (2013) 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) as well as two univariate models for inflation (**Trend-bound** and **AR-Trend-bound**) which were found to work well in Chan, Koop and Potter (2013).

The relatively non-informative prior we use for **Bi-UC** is given in the Online Appendix. The priors for the models which restrict one or more

²This model and the following one are univariate models for inflation so no results for the unemployment rate are presented.

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 VARs, **VAR(2)**, **VAR(2)-Minn** and **VAR(2)-SV**, are given in the Online Appendix.

Tables 1 and 2 report RMSFEs and sums of log predictive likelihoods for unemployment and inflation individually. We also present, in the bottom panel of Table 2, 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. For inflation, Table 3 present results for a shorter forecast evaluation period beginning in 1985Q1. We do not present post-1985 results for unemployment since results are similar to the post-1975 period.

Overall, these tables indicate that our **Bi-UC** is forecasting well. The bottom panel of Table 2, which contains our preferred measure of joint forecast performance of the entire predictive distribution for inflation and unemployment, shows **Bi-UC** to be the best forecasting model by a substantial margin at medium and long horizons. At shorter horizons, **Bi-UC** with simple restrictions on them forecast best. For instance, one quarter ahead the **Bi-UC** with the slope of the Phillips curve restricted to be constant over time forecasts best. A similar pattern holds to a lesser extent when looking at marginal predictive likelihoods for inflation and unemployment individually. **Bi-UC** often performs best. For cases where it is not the best, it is forecasting almost as well as the best model and this best model is usually a restricted version of **Bi-UC**. Furthermore, most of the other candidates forecast very poorly in at least one case. For MSFEs the good forecast performance found with predictive likelihoods is repeated for the unemployment rate, but not the inflation rate over the post-1975 period. However, an interesting pattern emerges where, over the post-1985 period (see Table 3), the **Bi-UC** once again emerges as the best (or nearly best) model as measured

by MSFE. 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 (which does not bound trend inflation nor the NAIRU) 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 substantially 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). The joint predictive likelihood for inflation and unemployment is perhaps the best single summary of forecasting performance we present. For this, **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 per-

formance.

The model of Stella and Stock (2013) 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. 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**).

It is difficult to compare forecasts from a bivariate model, estimated to fit a two equation system, with those from a univariate model. But, with this qualification in mind, it can be seen that there is no clear winner between our new **Bi-UC** model and the univariate models for bounded trend inflation of Chan, Koop and Potter (2013). At some forecasting horizons, **Bi-UC** forecasts best, but at others (particularly longer horizons) the **AR-Trend-bound** exhibits superior forecast performance, thus reinforcing a view expressed in Stella and Stock (2013) that it is often hard to beat a good univariate time series model. But this should not blunt interest in our bivariate model since univariate forecasting is not the only thing the researcher is interested. Our bivariate model does forecast well and, as we have illustrated, many things a policymaker might be interested in involve the joint predictive density of inflation and unemployment. In such cases, individual univariate models are not sufficient.

3.3 Estimates of Trend Inflation, NAIRU and Other Latent States

In previous work with a 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 late 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 3 plots the four main latent states, $\tau_t^\pi, \tau_t^u, \rho_t^\pi$ and λ_t , estimated using the full sample. That is, Figure 3 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

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

Inflation					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
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-$\lambda_t-\rho_t^\pi$	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-$-\tau_t^\pi-\tau_t^u$	0.95	0.96	1.02	0.99	1.00
Stella-Stock	1.03	0.98	0.98	0.96	0.97
Trend-bound	1.33	1.09	0.93	0.88	0.89
AR-Trend-bound	0.94	1.06	0.91	0.93	0.93
Unemployment					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
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-$\lambda_t-\rho_t^\pi$	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-$-\tau_t^\pi-\tau_t^u$	0.94	0.91	0.93	0.95	0.96
Stella-Stock	1.34	1.00	0.98	0.99	0.98

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

Inflation					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
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-$\lambda_t-\rho_t^\pi$	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-$-\tau_t^\pi-\tau_t^u$	28.9	30.8	25.2	24.6	21.5
Stella-Stock	18.9	14.6	11.6	17.1	17.6
Trend-bound	3.6	19.5	30.9	41.9	35.8
AR-Trend-bound	28.2	16.3	36.2	38.8	36.7
Unemployment					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
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-$\lambda_t-\rho_t^\pi$	-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-$-\tau_t^\pi-\tau_t^u$	-10.1	-1.9	10.2	18.6	20.4
Stella-Stock	-62.0	-7.6	0.6	0.6	-1.9
Joint					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
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-$\lambda_t-\rho_t^\pi$	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	-70.716	-66.4	-79.8	-94.7	-104.1
UCSV-AR(2)	9.9	7.0	2.7	-5.0	-19.7
Bi-UC-NoBound-$-\tau_t^\pi-\tau_t^u$	19.7	28.4	37.3	43.4	39.0
Stella-Stock	-41.1	12.7	18.5	13.4	6.6

Table 3: Relative RMSFEs and sum of log predictive likelihoods (against **VAR(2)**) for forecasting inflation; 1985-2013.

Relative RMSFEs					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	1.00	1.00	1.00	1.00	1.00
VAR(2)-Minn	1.00	0.99	0.99	1.01	1.02
VAR(2)-SV	1.02	0.99	0.94	0.91	0.95
Bi-UC	0.90	0.85	0.87	0.86	0.85
Bi-UC-const-λ_t	0.90	0.85	0.87	0.84	0.85
Bi-UC-const-ρ_t^π	0.96	0.89	0.87	0.84	0.84
Bi-UC-const-$\lambda_t-\rho_t^\pi$	0.97	0.89	0.87	0.84	0.84
Bi-UC-NoBound	0.92	0.87	0.95	1.16	1.89
Bi-UC-NoSV	0.93	0.87	0.89	0.86	0.84
Bi-UC-TVP-ρ^u	0.91	0.87	0.90	1.17	4.47
Bi-RW	1.11	1.21	1.16	1.13	1.30
UCSV-AR(2)	0.95	0.91	0.96	1.04	1.19
Bi-UC-NoBound-$-\tau_t^\pi-\tau_t^u$	0.90	0.86	0.89	0.88	0.85
Stella-Stock	0.95	0.97	1.09	1.09	1.16
Trend-bound	0.91	0.89	0.88	0.81	0.86
AR-Trend-bound	0.91	0.85	0.85	0.84	0.84
Sum of log predictive likelihoods					
	$k = 1$	$k = 4$	$k = 8$	$k = 12$	$k = 16$
VAR(2)	0.0	0.0	0.0	0.0	0.0
VAR(2)-Minn	-1.9	1.2	1.4	0.3	1.2
VAR(2)-SV	9.1	20.8	12.6	11.9	8.8
Bi-UC	33.0	35.2	28.7	27.8	21.6
Bi-UC-const-λ_t	33.6	36.3	31.2	31.1	25.1
Bi-UC-const-ρ_t^π	26.6	26.6	21.3	20.4	17.4
Bi-UC-const-$\lambda_t-\rho_t^\pi$	26.9	26.9	21.2	20.0	17.6
Bi-UC-NoBound	31.7	33.5	25.6	24.1	15.4
Bi-UC-NoSV	9.0	12.3	13.9	14.9	14.2
Bi-UC-TVP-ρ^u	32.8	34.1	29.5	29.2	20.6
Bi-RW	-15.5	-43.6	-68.8	-87.2	-102.3
UCSV-AR(2)	30.6	30.0	21.7	18.3	3.3
Bi-UC-NoBound-$-\tau_t^\pi-\tau_t^u$	32.8	34.3	26.2	23.6	18.0
Stella-Stock	29.0	22.5	13.2	14.7	7.7
Trend-bound	26.7	31.6	35.0	43.6	33.7
AR-Trend-bound	30.9	29.7	33.5	37.9	33.1

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, Figure 3 suggests NAIRU increases from roughly 5 to 7% from the late 1960s to 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 (2013) 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.

Our forecasting results suggest that, at least for long run forecasts, the benefits of bounding largely arise through the bounding of ρ_t^π and λ_t , rather than bounding trend inflation or NAIRU. To investigate this point more deeply, Figure 3 also presents estimates of the latent states produced by **Bi-UC-NoBound**— τ_t^π — τ_t^u . There are some slight differences between estimates produced by our model (which bounds all latent states) and those produced by a model which only bounds ρ_t^π and λ_t , but these are not large. This reinforces the story of the forecasting section that it is the bounding of ρ_t^π and λ_t which is particularly important. However, a slightly different story emerges when we look at filtered estimates of the latent states.

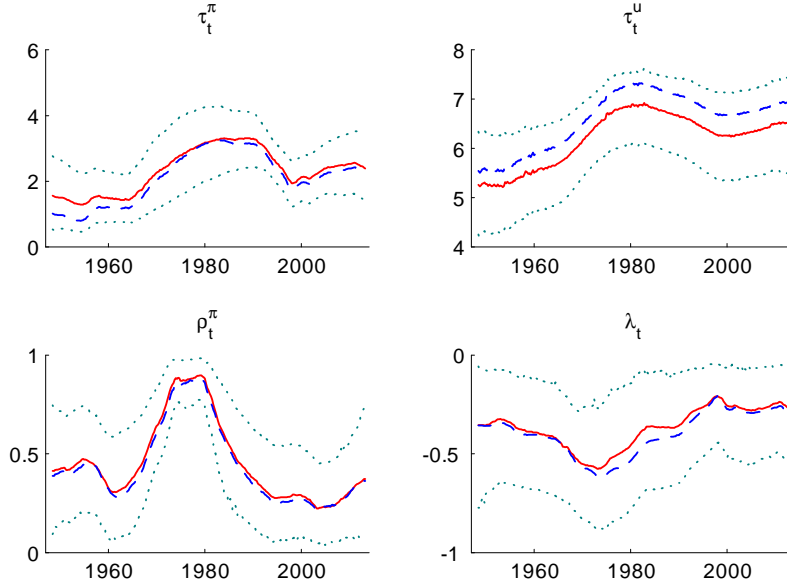


Figure 3: Posterior estimates for the latent states using the full sample. The solid lines represent the posterior means under the **Bi-UC** while the dotted lines are the 5 and 95 percentiles. The dashed lines represent the posterior means under the **Bi-UC-NoBound** $-\tau_t^\pi - \tau_t^u$.

Figure 3 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 . Results are presented in Figure 4 which plot the filtered states for our **Bi-UC** compared to **Bi-UC-NoBound** (which has no bounding at all) and **Bi-UC-NoBound** $-\tau_t^\pi - \tau_t^u$ which bounds only parameters ρ_t^π and λ_t .

Before considering the role of bounding, we note that the broad patterns in Figure 4 for **Bi-UC** are similar to Figure 3, suggesting that our bivariate unobserved components model can provide sensible real-time estimates of NAIRU and 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 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 3 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 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.

From a comparison of the lines in Figure 4, it can be seen that the most important effect of bounding is in ensuring smooth behavior of the NAIRU. That is, without bounding the NAIRU counter-intuitively spikes at nearly 8% in the early 1980s. This occurs even if we only bound ρ_t^π and λ_t . Previously, we said that it appeared that the bounding of ρ_t^π and λ_t was of particular importance. It is. However, from Figure 4 it can be seen that it is not enough to ensure sensible filtered estimates of the NAIRU. To achieve this goal, we must additionally bound NAIRU itself. A similar finding occurs, to a lesser extent, for trend inflation.

Bounding ρ_t^π and λ_t can be seen to have an important impact on the filtered estimates of ρ_t^π . Without this bounding, its estimate becomes greater than one in the late 1970s and early 1980s. This is, no doubt, driving some results discussed previously in our forecasting sections. Forecasts use filtered estimates and if the filtered estimate of ρ_t^π is evolving into the explosive region, it will negatively impact on forecast performance.

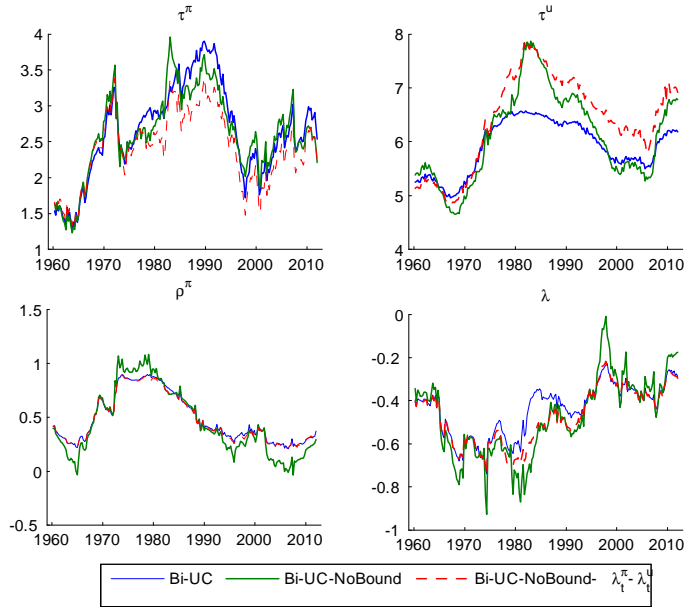


Figure 4: Estimates of the filtered states $E(\eta_t | Data_t)$ where $\eta_t = \tau_t^\pi, \tau_t^u, \rho_t^\pi, \lambda_t$ and $Data_t$ is the data from the beginning of the sample to time t .

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 the NAIRU. 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, 2013). 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

and that the fundamentals driving the NAIRU should not have unrestricted variation when expressed in unemployment rate space). 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.

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.

Campbell, J., Evans, C., Fisher, J. and Justiniano, A. (2012). "Macroeconomic effects of Federal Reserve forward guidance," *Brookings Papers on Economic Activity*, Spring 2012, 1-80.

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.

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.

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. (2013). “A state-dependent model for inflation forecasting,” Board of Governors of the Federal Reserve System, International Finance Discussion Papers, Number 1062.

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.