


THE FEDERAL RESERVE BANK of KANSAS CITY
ECONOMIC RESEARCH DEPARTMENT

A Bayesian Evaluation of Alternative Models of Trend Inflation

Todd E. Clark and Taeyoung Doh

November 2011

RWP 11-16



RESEARCH WORKING PAPERS

A Bayesian Evaluation of Alternative Models of Trend Inflation *

Todd E. Clark
Federal Reserve Bank of Cleveland

Taeyoung Doh
Federal Reserve Bank of Kansas City

November 2011

Abstract

With the concept of trend inflation now widely understood as to be important as a measure of the public's perception of the inflation goal of the central bank and important to the accuracy of longer-term inflation forecasts, this paper uses Bayesian methods to assess alternative models of trend inflation. Reflecting models common in reduced-form inflation modeling and forecasting, we specify a range of models of inflation, including: AR with constant trend; AR with trend equal to last period's inflation rate; local level model; AR with random walk trend; AR with trend equal to the long-run expectation from the Survey of Professional Forecasters; and AR with time-varying parameters. We consider versions of the models with constant shock variances and with stochastic volatility. We first use Bayesian metrics to compare the fits of the alternative models. We then use Bayesian methods of model averaging to account for uncertainty surrounding the model of trend inflation, to obtain an alternative estimate of trend inflation in the U.S. and to generate medium-term, model-average forecasts of inflation. Our analysis yields two broad results. First, in model fit and density forecast accuracy, models with stochastic volatility consistently dominate those with constant volatility. Second, for the specification of trend inflation, it is difficult to say that one model of trend inflation is the best. Among alternative models of the trend in core PCE inflation, the local level specification of Stock and Watson (2007) and the survey-based trend specification are about equally good. Among competing models of trend GDP inflation, several trend specifications seem to be about equally good.

Keywords: Likelihood, model combination, forecasting

JEL Classifications: E31, E37, C11

* *Clark* (corresponding author): Economic Research Dept.; Federal Reserve Bank of Cleveland; P.O. Box 6387; Cleveland, OH 44101; todd.clark@clev.frb.org. *Doh*: Economic Research Dept.; Federal Reserve Bank of Kansas City; 1 Memorial Drive; Kansas City, MO 64198; taeyoung.doh@kc.frb.org. We gratefully acknowledge helpful comments from Gianni Amisano, Garland Durham, Francesco Ravazzolo, Norman Swanson, Shaun Vahey, and seminar or conference participants at the Federal Reserve Bank of Cleveland, 2011 Bayesian workshop at the Rimini Center for Economic Analysis, and 2011 Bank of Canada workshop on nowcasting and short-term forecasting. Comments from Ellis Tallman were especially helpful to the paper's exposition. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Banks of Cleveland or Kansas City or the Federal Reserve Board of Governors.

1 Introduction

The concept of trend inflation is now viewed as important for a variety of reasons. First, as discussed in such sources as Kozicki and Tinsley (2001a) and Faust and Wright (2011), trend inflation is generally thought of as a measure of the public’s perception of the inflation goal of the central bank. As such, trend inflation can provide an important measure of the credibility of monetary policy. Second, trend inflation plays an important role in inflation forecasts. As the horizon increases, a forecast from a model will converge to the trend captured by the model. Prior studies such as Kozicki and Tinsley (1998) and Clark and McCracken (2008) have shown that the concept of trend embedded in an inflation model plays a key role in longer-term inflation forecasts (in these studies, point forecasts). More recently, Faust and Wright (2011) have argued that the accuracy of inflation forecasts (point forecasts) is crucially dependent on the accuracy of the underlying trend estimate. In addition, Williams (2009) has shown that forecast probabilities of deflation can be very sensitive to the model of trend inflation. In his analysis, models in which the inflation trend is represented by the 10-year ahead inflation forecast from the Survey of Professional Forecasters (SPF) yield a lower risk of deflation than do models in which the inflation trend is simply a function of past inflation.

But as Williams (2009) suggests, exactly how the trend in inflation should be modeled is not clear. Studies such as Clark (2011), Faust and Wright (2011), and Wright (2011) use measures of long-run inflation expectations from surveys of forecasters (either the SPF or Blue Chip) to capture trend inflation and show that the survey-based trend improves the accuracy of model-based forecasts relative to models that treat trend inflation as constant. While some comparisons in Faust and Wright (2011) suggest a survey-based trend to be hard to beat, a survey-based trend may nonetheless not be ideal. One reason is that surveys may be upward biased. For most of the period since 1995, the SPF has projected long-term CPI inflation (10 years ahead) of 2.5 percent. Yet core inflation was generally well below that threshold for most of the period. Since late 2007, long-term forecasts of PCE inflation from the SPF have consistently been a bit higher than longer-term projections from the FOMC.¹ A second reason is that survey-based measures lack some of the positive aspects

¹For example, in late April 2010, the central tendency of FOMC participants’ longer-run projection of PCE inflation (the longer-run projections “represent each participant’s assessment of the rate to which each variable would be expected to converge over time under appropriate monetary policy and in the absence of further shocks”) was 1.5 to 2.0 percent. In the mid-May SPF, the median forecast of average PCE inflation for 2009-2018 was 2.3 percent.

of models. With a model, it is easier (compared to using a survey result) to know what determines the estimate of trend, replicate results, and obtain a forecast whenever it is needed.

Other research suggests a range of possible and reasonable models of trend inflation.² A number of studies and still-used forecasting models measure inflation expectations, or in effect, trend inflation, with past inflation (e.g., Brayton, Roberts, and Williams 1999, Gordon 1998, and Macroeconomic Advisers 1997).³ Another array of studies has modeled trend inflation as following a random walk (e.g., Cogley, Primiceri, and Sargent 2010, Cogley and Sbordone 2008, Ireland 2007, Kiley 2008, Kozicki and Tinsley 2006, Mertens 2011, Piger and Rasche 2008, and Stock and Watson 2007, 2010). Inflation less trend follows an autoregressive process in some of these studies (e.g., Cogley, Primiceri, and Sargent 2010) but not in others (Stock and Watson 2007). In addition, some research using random walk trends (Stock and Watson 2007, Mertens 2011, and Cogley, Primiceri, and Sargent 2010) has found that the variability of the trend component in inflation has varied over time.

Accordingly, this paper uses Bayesian methods to assess alternative models of trend inflation. We first use Bayesian metrics to compare the fits of alternative models. We then use Bayesian methods of model averaging to account for uncertainty surrounding the model of trend inflation, to obtain an alternative estimate of trend inflation in the U.S. and to generate medium-term, model-average forecasts of inflation. Reflecting models common in reduced-form inflation modeling and forecasting, we specify a range of models of inflation.⁴ We use predictive likelihoods to weight each model and forecast and construct probability-weighted average estimates of trend inflation and inflation forecasts. For forecasting, we consider not only point predictions but also density forecasts (specifically, deflation probabilities and average log predictive scores). Morley and Piger (2010) follow a broadly similar Bayesian model averaging method to estimate the trend and business cycle component of GDP.

In the interest of simplicity, we focus primarily, although not exclusively, on univariate models of inflation.⁵ Studies such as Atkeson and Ohanian (2001), Clark and McCracken

²The survey of Faust and Wright (2011) provides a longer list of studies that have considered time-varying inflation trends.

³In a related formulation, Cogley (2002) develops an exponential smoothing model of trend.

⁴Our modeling approach differs from that of Mertens (2011) in that we consider multiple models of trend, while Mertens considers multiple indicators with a single, random walk trend model.

⁵Multivariate model results reported in the paper include the results for the model using a survey-based inflation trend (in the presentation, we treat this as a univariate model) and results for inflation models including an unemployment gap. While not reported in the paper, we also considered a bivariate model

(2008, 2010), Dotsey, Fujita, and Stark (2011), and Stock and Watson (2003, 2007) have found that, in U.S. data since at least the mid-1980s, univariate models of inflation typically forecast better (in terms of point forecast accuracy) than do multivariate models. Our set of models incorporates significant differences in the trend specification: AR with constant trend; AR with trend equal to last period's inflation rate; local level model; AR with random walk trend; AR with trend equal to the long-run expectation from the SPF; and AR with time-varying parameters. Finally, we consider versions of the models with constant shock variances and with time-varying shock variances (stochastic volatility). To further highlight the importance of trend specification to inflation inferences, in assessing deflation probabilities we also consider some specifications that include an unemployment gap. Our analysis extends Faust and Wright (2011) by considering a wider array of trend models and inflation models with both constant volatilities and stochastic volatilities, using Bayesian methods to compare overall model fit, and examining density forecasts in addition to point forecasts.

Consistent with Clark (2011), one clear finding of our analysis is that, in terms of model fit and density forecast accuracy, the models with stochastic volatility dominate those with constant volatility. The incorporation of stochastic volatility also materially affects estimates of the probability of deflation, sharply lowering them in the past decade, a period in which deflation became a concern of policymakers and others.

As to the specification of trend inflation, our results show that it is difficult to say that one model of trend inflation is the best. Among alternative models of the trend in core inflation, the local level specification of Stock and Watson (2007) and the SPF specification are about equally good. Our Bayesian measures of model fit for the full sample indicate the local level specification of trend is best, with the SPF specification next best and reasonably close. However, we also find that model fit has evolved considerably over the sample. For example, relative to other models, the fit of the local level model with stochastic volatility improved significantly in the last several years of the sample. In out-of-sample forecasting (both point and density), the local level model is slightly more accurate than the SPF-based model, but not by a significant margin. Both in-sample and out-of-sample, the other trend models are inferior, although sometimes only modestly so.

in inflation and the 10-year Treasury bond yield, with a random walk trend in inflation accounting for the trends in both variables. As this model did not fit the data any better than the other models considered, we have omitted these results in the interest of brevity.

For inflation in the GDP price index, several of the trend specifications we consider seem to be about equally good. Our Bayesian measures of model fit for the full sample indicate the model with time-varying parameters fits best, with the lagged inflation and SPF-based models next (with likelihoods that imply model weights of about 20 percent, compared to 45 percent for the best model). In out-of-sample forecasting, many of the models perform similarly in both point and density prediction.

Broadly, these results highlight the importance of considering predictions from a range of models — as opposed to a single model — that fit the historical data about as well as one another. This is especially the case for the evaluation of risks such as the chance of deflation.

The paper proceeds as follows. Section 2 describes the data. Sections 3 and 4 present the models and estimation methodology, respectively. Section 5 presents the results. An appendix provides details of the estimation algorithms, priors, and computation of predictive likelihoods. Section 6 concludes.

2 Data

We focus on modeling and forecasting core PCE inflation, which refers to the price index for personal consumption expenditures excluding food and energy, commonly viewed as the Federal Reserve’s preferred measure of inflation. But we also include some results for a broader measure of inflation, the GDP price index, often used in assessments of inflation dynamics and forecast comparisons. Inflation rates are computed as annualized log percent changes ($400 \ln(P_t/P_{t-1})$).

Our time period of focus (for our estimates) is 1960 through mid-2011. To obtain training sample estimates (section 4 provides more detail on the samples), we use other data to extend the core PCE and GDP price indexes back in time. Specifically, we merge the published core PCE inflation rate (beginning in 1959:Q2) with: (1) a constructed measure of core inflation for 1947:Q2 through 1959:Q1 that excludes energy goods (for which data go back to 1947) but not energy services (for which data only go back to 1959); and (2) overall CPI inflation for 1913:Q2 through 1947:Q1.⁶ In analysis of inflation in the GDP price index, for which published data extend back through 1947, we merge the published

⁶We used the methodology of Whelan (2002) to construct the measure of core inflation from raw data series on total PCE and energy goods spending and prices. As to the historical CPI, we obtained a 1967 base year index from the website of the Bureau of Labor Statistics and seasonally adjusted it with the X11 algorithm.

GDP inflation rate (beginning in 1947:Q2) with overall CPI inflation for 1913:Q2 through 1947:Q1.

For our models that rely on survey forecasts of long-run inflation as the measure of trend inflation, we use the survey-based long-run (5- to 10-year-ahead) PCE inflation expectations series of the Federal Reserve Board of Governor’s FRB/US econometric model. The FRB/US measure splices econometric estimates of inflation expectations from Kozicki and Tinsley (2001a) early in the sample to 5- to 10-year-ahead survey measures compiled by Richard Hoey and, later in the sample, to 10-year-ahead values from the Federal Reserve Bank of Philadelphia’s Survey of Professional Forecasters. In presenting the results, we refer to this series as PTR, using the notation of the FRB/US model.

Finally, for our analysis of deflation probabilities from models that include an economic activity indicator, we use an unemployment gap, defined as the unemployment rate for men between the ages of 25 and 54 less a one-sided estimate of the trend in that unemployment rate. Some may view the trend as a measure of the natural rate of unemployment. The use of an unemployment rate for prime-age men helps to reduce (but not eliminate) the influence of long-term demographic trends. Following Clark (2011), we use exponential smoothing to estimate the trend, with a smoothing coefficient of 0.02.

With the exception of the CPI index obtained from the Bureau of Labor Statistics, we obtained all of the data from the FAME database of the Federal Reserve Board of Governors.

3 Models

3.1 Constant volatility

Our baseline model specification relates current inflation to past inflation and current and past rates of trend inflation:

$$\pi_t - \pi_t^* = b(L)(\pi_{t-1} - \pi_{t-1}^*) + v_t, \quad (1)$$

where π_t denotes actual inflation and π_t^* denotes trend inflation. The trend corresponds to the moving endpoint concept of Kozicki and Tinsley (1998, 2001a,b): $\pi_t^* = \lim_{h \rightarrow \infty} E_t \pi_{t+h}$. In the absence of deterministic terms, this concept of trend inflation is the same as in the Beveridge and Nelson (1981) decomposition.

Building on this baseline, we consider a range of models, each incorporating a different specification of the trend in inflation. These specifications, drawn from studies such as those described in the introduction, include:

- $\pi_t^* = \text{constant}$ (*constant trend*)
- $\pi_t^* = \pi_{t-1}$ (π_{t-1} *trend*)
- $\pi_t^* = \text{long-run survey forecast}$ (*PTR trend*)
- $\pi_t^* = \pi_{t-1}^* + n_t$ (*random walk trend*)

The first model is a stationary AR model, written in demeaned form; Villani (2009) develops an approach for estimating such models with a steady state prior. The second model is a stationary AR model in the first difference of inflation. The third and fourth models are AR models in detrended inflation, where trend is defined as the PTR trend in one case and a random walk process in the other.

We consider two other related model formulations. The first is the local level model, an unobserved components specification considered in such studies as Stock and Watson (2007):

$$\pi_t = \pi_t^* + v_t, \quad \pi_t^* = \pi_{t-1}^* + n_t. \quad (2)$$

This model simplifies the version of equation (1) with a random walk trend by imposing AR coefficients of zero. The local-level model implies a filtered trend estimate that could be computed by exponential smoothing. The AR model with a random walk trend in inflation generalizes this local-level model by allowing autoregressive dynamics in the deviation from trend; this form of the AR model is equivalent to the trend-cycle model of Watson (1986).

The second alternative model we consider is an AR specification with time variation in all coefficients, considered in such studies as Cogley and Sargent (2005):

$$\pi_t = b_{0,t} + \sum_{i=1}^p b_{i,t} \pi_{t-i} + v_t, \quad b_t = b_{t-1} + n_t, \quad (3)$$

where the vector b_t contains the intercept and slope coefficients of the AR model and $\text{var}(n_t) = Q$. For this time-varying parameters (TVP) model, we follow Cogley and Sargent (2005) and others in the growing TVP literature in estimating trend inflation as the instantaneous mean, as implied in each period by the intercept and slope coefficients.

In our examination of the role of trend inflation in assessing the probability of deflation, we supplement our analysis with two models — the π_{t-1} trend and PTR trend specifications — augmented to include the unemployment gap. Letting y_t denote the vector containing $\pi_t - \pi_t^*$ and u_t , where u denotes the unemployment gap, we use a conventional VAR (without

intercept):

$$y_t = b(L)y_{t-1} + v_t. \quad (4)$$

Finally, as to the lag order of the models with autoregressive dynamics (every model except the local level), for computational tractability we fix the lag order at two for all models. As a robustness check, we estimated most of the models with stochastic volatility using four lags and obtained very similar results (for four lags versus two) for model fit and forecast performance.

3.2 Stochastic volatility

In light of existing evidence of sharp changes over time in the volatility of inflation (e.g., Clark 2011, Cogley and Sargent 2005, and Stock and Watson 2007), we consider versions of the models above supplemented to allow for time variation in the residual variances. As emphasized in Clark (2011) and Jore, Mitchell, and Vahey (2010), modeling changes in volatility is likely to be essential for accurate forecasts of measures that require the entire forecast density — e.g., the probability of deflation.

With stochastic volatility, the basic model specification is:

$$\begin{aligned} \pi_t - \pi_t^* &= b(L)(\pi_{t-1} - \pi_{t-1}^*) + v_t, \\ v_t &= \lambda_t^{0.5} \epsilon_t, \quad \epsilon_t \sim N(0, 1), \\ \log(\lambda_t) &= \log(\lambda_{t-1}) + \nu_t, \quad \nu_t \sim N(0, \phi). \end{aligned} \quad (5)$$

This specification applies to the constant trend, π_{t-1} trend, PTR trend, and random walk trend models. In the case of the random walk trend specification, the model also includes an equation for the trend:

$$\pi_t^* = \pi_{t-1}^* + n_t, \quad n_t \sim N(0, \sigma_n^2). \quad (6)$$

The local level model with stochastic volatility takes the form given in Stock and Watson (2007):

$$\begin{aligned} \pi_t &= \pi_t^* + v_t, \quad \pi_t^* = \pi_{t-1}^* + n_t, \\ v_t &= \lambda_{v,t}^{0.5} \epsilon_{v,t}, \quad \epsilon_{v,t} \sim N(0, 1), \\ n_t &= \lambda_{n,t}^{0.5} \epsilon_{n,t}, \quad \epsilon_{n,t} \sim N(0, 1), \\ \log(\lambda_{v,t}) &= \log(\lambda_{v,t-1}) + \nu_{v,t}, \quad \nu_{v,t} \sim N(0, \phi_v), \\ \log(\lambda_{n,t}) &= \log(\lambda_{n,t-1}) + \nu_{n,t}, \quad \nu_{n,t} \sim N(0, \phi_n). \end{aligned} \quad (7)$$

In this paper, we generalize the approach of Stock and Watson (2007) by actually estimating the key parameters of the model, the variances of shocks to log volatilities, rather than treating them as fixed. However, we obtained very similar results from the model when we effectively fixed these variances at the Stock-Watson values (of 0.04) by setting the prior means of ϕ_v and ϕ_n at 0.04 and the prior degrees of freedom at 10,000.

The TVP model with stochastic volatility takes the form given in Cogley and Sargent (2005), simplified to a univariate process:

$$\begin{aligned}\pi_t &= b_{0,t} + \sum_{i=1}^p b_{i,t} \pi_{t-i} + v_t, \\ b_t &= b_{t-1} + n_t, \quad \text{var}(n_t) = Q, \\ v_t &= \lambda_t^{0.5} \epsilon_t, \quad \epsilon_t \sim N(0, 1), \\ \log(\lambda_t) &= \log(\lambda_{t-1}) + \nu_t, \quad \nu_t \sim N(0, \phi).\end{aligned}\tag{8}$$

Finally, for the bivariate VAR models that include both inflation less trend and the unemployment gap, letting y_t denote the vector of variables included, the bivariate VAR with stochastic volatility takes the form:

$$\begin{aligned}y_t &= B(L)y_{t-1} + v_t, \\ v_t &= A^{-1} \Lambda_t^{0.5} \epsilon_t, \quad \epsilon_t \sim N(0, I_2), \quad \Lambda_t = \text{diag}(\lambda_{1,t}, \lambda_{2,t}), \\ \log(\lambda_{i,t}) &= \log(\lambda_{i,t-1}) + \nu_{i,t}, \quad \nu_{i,t} \sim N(0, \phi_i) \quad \forall i = 1, 2,\end{aligned}\tag{9}$$

where A = a lower triangular matrix with ones on the diagonal and a coefficient a_{21} in row 2 and column 1 (such that A reflects a Choleski structure). Again, for simplicity, for all models with autoregressive dynamics (every model except the local level), we fix the lag order at two.

4 Estimation

4.1 Algorithms

Focusing on the 1960:Q1 to 2011:Q2 period, we estimate the models described above using Bayesian Markov Chain Monte Carlo (MCMC) methods. More specifically, to limit the influence of priors, we use empirical Bayes methods, specifying the priors for our model estimates on the basis of posterior estimates obtained from an earlier sample. For this purpose, we divide our data sample into three pieces: an estimation sample of 1960:Q1-2011:Q2, an intermediate estimation sample of 1947-59 for which we generate posterior

estimates that form the basis of our priors for the estimation sample, and a training sample of 1913-1946 that we use to set priors for estimation in the intermediate sample.

We use Gibbs samplers to estimate the models with constant volatilities. For the constant trend model, we use the algorithm detailed in Villani (2009). For the π_{t-1} trend and PTR trend specifications, the algorithm takes the Normal-diffuse form of Kadiyala and Karlsson (1997). Estimation of the local level and random walk trend models is a straightforward application of Gibbs sampling, using state-space representations of the models. Estimation of the TVP model is described in sources such as Cogley and Sargent (2001); the algorithm for the local level model is effectively the same, with a lag length of 0 in the AR model.⁷ In all cases (except the local level model), we follow the approach of Cogley and Sargent (2005), among others, in discarding draws with explosive autoregressive roots (and re-drawing).

We use Metropolis-within-Gibbs MCMC algorithms to estimate the models with stochastic volatilities, combining some of the key Gibbs sampling steps for the constant volatility models with Cogley and Sargent’s (2005) Metropolis algorithm (taken from Jacquier, Polson, and Rossi 1994) for stochastic volatility. For the constant trend, π_{t-1} trend, and PTR trend specifications with stochastic volatility, our algorithms are the same as those used in Clark (2011) and Clark and Davig (2011). For the TVP model with stochastic volatility, our algorithm takes the form described in Cogley and Sargent (2005). Again, for all models, we discard draws with explosive autoregressive roots. The appendix provides more detail on all of the algorithms for estimation with stochastic volatility.

All of our reported results are based on samples of 5000 posterior draws. To ensure the reliability of our results, we estimated each model with a large number of MCMC draws, obtained by first performing burn-in draws and then taking additional draws, from which we retained every k ’th draw to obtain a sample of 5000 draws. Skipping draws is intended to reduce correlation across retained posterior draws. Koop and Potter (2008) show that MCMC chains for VARs with TVP can be quite slow to mix, a finding confirmed in Clark and Davig (2011). Drawing on this previous evidence of convergence properties and our own selected checks of convergence properties, we use larger burn-in samples and higher skip intervals for models with latent states (unobserved trends or time-varying volatility)

⁷For all of our models with time-varying trends or coefficients, we use the algorithm of Durbin and Koopman (2002) for the backward smoothing and simulation, rather than the Carter and Kohn (1994) smoother used by Cogley and Sargent (2005), because the Durbin-Koopman algorithm is faster in the software we used.

than models without latent states. The appendix includes a table with the burn-in counts and skip intervals used for each model.

4.2 Model assessment

Similar to Morley and Piger’s (2010) approach to assessing models of the business cycle component of GDP, we construct a measure of trend inflation which addresses model uncertainty by averaging over different models of trend inflation. We assume equal prior weight for each model and use the following posterior model probability to average across different models:

$$w_i(\pi^{(T)}) = \frac{p(M_i|\pi^{(T)})}{\sum_{i=1}^n p(M_i|\pi^{(T)})}, \quad (10)$$

where M_i denotes model i and $\pi^{(T)}$ denotes the time series of inflation up to period T .

To assess the congruency of each model with the data and compute the posterior model probabilities that determine the model weights, we follow Geweke and Amisano (2010) in using 1-step ahead predictive likelihoods. Sources such as Geweke (1999) and Geweke and Whiteman (2006) emphasize the close relationship between the predictive likelihood and marginal likelihood: as stated in Geweke (1999, p.15), “... the marginal likelihood summarizes the out-of-sample prediction record... as expressed in ... predictive likelihoods.” Following Geweke and Amisano (2010), we use the log predictive likelihood defined as

$$\log \text{PL}(M_i) = \sum_{t=t_0}^T \log p(\pi_t^o | \pi^{(t-1)}, M_i), \quad (11)$$

where π_t^o denotes the observed outcome for inflation in period t and $\pi^{(t-1)}$ denotes the history of inflation up to period $t - 1$. Following studies such as Bauwens, et al. (2011), we compute $p(\pi_t^o | \pi^{(t-1)}, M_i)$ from the simulated predictive density, using a kernel smoother to estimate the empirical density (from draws of forecasts). Finally, in computing the log predictive likelihood, we sum the log values over different samples, detailed below.

In averaging trends and forecasts from their posterior densities, we follow the mixture of distributions approach described in Bjornland, et al. (2010). Specifically, from the posterior sample of 5000 draws, we sample 5000 draws with replacement, taking the draw from model i ’s density with probability M_i . We then form the statistics of interest (median trend, etc.) from this mixture distribution.

4.3 Forecasting

To assess the role of the trend model in medium-term forecasting, we consider the accuracy of forecasts of horizons from 1 to 16 quarters. The role of trend inflation in the forecast will increase with the forecast horizon, at a rate that depends on the persistence of inflation relative to trend. We assess the accuracy of point forecasts using mean errors and root mean square errors (RMSEs) and the accuracy of density forecasts using average log predictive scores (computed using the density estimated from forecast draws). The predictive scores provide the broadest possible measure of the calibration of the density forecasts. We evaluate forecasts over the period 1985:Q1 to 2011:Q2. For the purpose of model evaluation and model combination based on real-time model fit, we begin forecasting in 1975:Q1. In forming the first forecast, for 1975:Q1, we estimate each model using data from 1960:Q1 through 1974:Q4, and form forecasts from horizons 1 through 16. We then proceed to move forward a quarter, estimate each model using 1960:Q1-1975:Q1 data, and form forecasts for horizons 1 through 16. We continue with this recursive approach to forecasting through the rest of the sample.

We also consider another specific aspect of the density forecast that may be particularly dependent on the specification of the trend in inflation: the probability of deflation, defined as inflation (over 4 quarters) next year of less than 0. More specifically, our deflation probability is defined as the probability of an average inflation rate less than 0 for periods $t + 4$ through $t + 7$, when forecasting starting in period t with models estimated using data through period $t - 1$.

For each model, for each (retained) draw in the MCMC chain, we draw forecasts from the posterior distribution using an approach like that of Cogley, Morozov, and Sargent (2005). For models with time-varying coefficients (TVP), trends (local level, random walk), or stochastic volatility, we simulate the latent variables over the forecast horizon, using their random walk structure. For example, to incorporate uncertainty associated with time variation in λ_t over the forecast horizon, we sample innovations to $\log \lambda_{t+h}$ from a normal distribution with variance ϕ , and use the random walk specification to compute $\log \lambda_{t+h}$ from $\log \lambda_{t+h-1}$. For each period of the forecast horizon, we then sample shocks v_{t+h} to the equation for inflation with a variance of λ_{t+h} and compute the forecast draw of π_{t+h} from the AR structure and drawn shocks. The appendix provides more detail on our simulation of predictive densities, for the models with stochastic volatility.

For all of the models with latent variables, the forecast distributions computed with this approach account for uncertainty about the latent variables — i.e., in many of the model specifications, the evolution of trend inflation. In the case of the model using the PTR measure of trend, for simplicity we fix the trend over the forecast horizon of $t + 1$ through $t + H$ at the value of PTR in period t .⁸ As a result, the predictive density from the PTR trend model abstracts from uncertainty about the evolution of trend over the forecast horizon. However, based on comparisons conducted in the research for Clark (2011), this simplification has little effect on the results.

5 Results

This section proceeds by first presenting model parameter estimates for the full sample of data and then reporting model fit for the full sample. The subsequent sections present the trend estimates and forecast results (with forecasting results for core PCE inflation and GDP inflation presented in separate subsections).

5.1 Full sample parameter estimates

Tables 1 and 2 provide parameter estimates (posterior means and standard deviations) for full sample estimates of the constant volatility and stochastic volatility models, respectively. In the interest of brevity, we report results only for core PCE inflation; results for inflation in the GDP price index are qualitatively very similar.

Consistent with some prior studies (e.g., Kozicki and Tinsley 2002), the AR models that allow time variation in mean or trend inflation yield modestly lower sums of AR coefficients, reflecting reduced persistence of inflation relative to mean or trend. For example, with constant volatility, the sum of coefficients for the constant trend AR model is about 0.95, while the sum of coefficients for the PTR trend and random trend models is a little above 0.8. The same pattern is evident in the models with stochastic volatility. Within the set of constant volatility models, the local level specification yields modestly larger shocks to trend ($\sigma_n^2 = 0.481$) than the noise component ($\sigma_v^2 = 0.214$). Consistent with some prior studies (e.g., Stock and Watson 2007), the local level model attributes much of the movement in inflation to the trend component. By comparison, the variance of shocks to trend is considerably smaller (although still measurable) in the random walk trend model, which

⁸This approach allows us to use a univariate rather than bivariate model, simplifying the model specification and speeding up the computations.

incorporates autoregressive dynamics.

Consistent with recent evidence from studies such as Clark (2011) and Cogley and Sargent (2005), all estimates of the models with stochastic volatility imply sizable variation over time in the variance of shocks to inflation. For example, in the PTR trend model, the posterior mean estimate of the variance of shocks to log volatility is 0.073. Figure 1’s plot of the posterior median of the time series of volatility (specifically, Figure 1 plots the median and 70% credible set for $\lambda_t^{0.5}$) confirms the considerable movement of volatility over the sample, dominated by the Great Moderation and a rise in volatility in the last decade of the sample. The estimates of the variance of shocks to log volatility (all shown in Table 2) and the time series of volatilities (in the interest of brevity, Figure 1 provides estimates for just the PTR model) are similar across all models with autoregressive dynamics. Estimates differ somewhat for the local level model, which yields less variability in the size of shocks to the noise component of the model. Instead, the local level model yields more sizable variation in the size of shocks to trend.

5.2 Model fit

To assess overall model likelihood and fit, Table 3 reports log predictive likelihoods (specifically, sums of 1-step ahead log likelihoods for 1975:Q1-2011:Q2). With the core PCE measure of inflation, among models with constant volatilities, the PTR trend specification fits the data far better than most of the other models, with a log predictive likelihood of -172.960 for the 1975-2011 sample. Based on this measure of Bayesian fit, the other models fit the data much worse (recall that, for model comparison, the differences in log predictive likelihoods would be exponentiated, so a difference in logs of a few points is a very large difference in probability). The fits of the constant trend, random walk trend, and TVP models are not as good, with log predictive likelihoods ranging from about -177 to -179.5. The local level model, of particular interest in light of the Stock and Watson (2007) evidence in support of a version with time-varying volatility, ranks next to worst among constant-volatility models of core inflation, with a log predictive likelihood of -181.104. By a small margin, the π_{t-1} trend specification yields the worst fit among models with constant volatility.

However, allowing for stochastic volatility yields a somewhat different view of congruency with the data. For each specification of inflation trend and AR dynamics, the log predictive likelihood is considerably better with stochastic volatility than constant volatility.

For example, with the PTR trend model of core PCE inflation, the version with stochastic volatility yields a log predictive likelihood of -161.884, compared to -172.960 for the version with constant volatility. In some cases, differences in likelihoods across models with stochastic volatility are smaller than differences in likelihoods across models with constant volatility. In addition, the model rankings change somewhat. With stochastic volatility, the best-fitting model is the local level specification. But the PTR trend model fits the data almost as well as the local level model. The other models don't fit nearly as well, with log predictive likelihoods 1.8 to 4.3 points lower.

Translated into model probabilities as described in section 4, the full-sample evidence gives a 56.9% probability to the local level-SV model and a 28.5% probability to the PTR trend-SV specification. The TVP-SV and random walk trend-SV specifications receive smaller, non-trivial weights, of 9.2% and 3.7%, respectively. The other two models with stochastic volatility receive weights of less than 1%, while the constant volatility models have weights of essentially 0.⁹

Results for the GDP price index are similar in some important respects. Among models with constant volatilities, the PTR trend and local level models yield, respectively, the best and worst predictive likelihoods. Models with stochastic volatility fit the data much better than the models with constant volatilities. For example, the PTR trend specification yields a log predictive likelihood of -206.904, while the same model with stochastic volatility has a likelihood of -191.246.

The key difference in the results for the GDP price index compared to the core inflation results is that, among models with stochastic volatility, the differences in model fit are more modest with the GDP measure of inflation. In the estimates based on the GDP price index, the local level-SV model does not beat other models as it does in estimates based on the core PCE price index. With the GDP price index, the TVP-SV model fits the data best, and the π_{t-1} trend-SV model and PTR trend-SV models are next best and comparable in fit. The log predictive likelihoods give the following probabilities to the models: 45.1% for TVP-SV, 21.3% for PTR-SV, 19.5% for π_{t-1} trend-SV, 9.4% for constant trend-SV, and 4.1% for local level-SV.

⁹For two inflation models that include an unemployment gap, we report log predictive likelihoods but not model weights. We use these models later in this section to further investigate the role of trend inflation in determining deflation probabilities. Because the overall fit of the models is of interest, we report their likelihoods. However, in light of our focus on models of trend inflation, we do not include these specifications in the model combinations considered below. Accordingly, we do not report model weights for these specifications.

To further assess congruence of the models with data over time, we follow examples such as Geweke and Amisano (2010) in plotting the time series of cumulative log predictive likelihoods. To facilitate this graphical assessment, we take the local level model with stochastic volatility (the Stock and Watson 2007 model) as the benchmark, and form, for every other model, the difference between its cumulative log predictive likelihood and the benchmark cumulative log predictive likelihood.

The results for core PCE inflation provided in Figure 2 indicate that there have been important changes over time in the relative fit of the models. Through the late 1980s, these relative likelihoods were positive for most of the constant volatility models, indicating that the local level-SV model did not fit the data as well as most of the constant volatility models. From the early-1990s onward, the ranking of the constant volatility models changed little, remaining near the full sample ranking (PTR trend best, etc.). Relative to constant volatility models, the stochastic volatility versions rapidly gain advantage starting in the early-1990s. This pattern likely reflects the influences of the Great Moderation on volatility and the success of the stochastic volatility models in capturing those influences. Moving forward in time, the rankings of the stochastic volatility models shift around some. Most notably, based on cumulative likelihoods from the mid-1990s through 2005, the PTR trend-SV model usually fits the data slightly better than the local level-SV model, but for the remainder of the sample, the local level model fits the data slightly better.

Qualitatively, the results for GDP inflation provided in Figure 3 are similar in important respects. Most notably, relative to constant volatility models, the stochastic volatility versions rapidly gain advantage starting in the early-1990s. In addition, over time, the rankings and fits of the stochastic volatility models shift around. For example, since the mid-1990s, the fit of the random walk trend-SV model has gradually deteriorated relative to the fit of the other models with stochastic volatility. However, the results for GDP inflation differ from the results for core PCE inflation in that there are smaller differences in the fits of some of the competing trend models (among models with stochastic volatility).

In light of the changes over time in the rankings of models, from a forecasting perspective it may be desirable to use model weights computed at each point in time from a rolling window of predictive likelihoods. Studies such as Jore, Mitchell, and Vahey (2010) and Kascha and Ravazzolo (2010) use rolling windows of likelihoods to compute model weights for forecasting. Accordingly, in this paper, in combining forecasts, we will consider model

averages based on 10-year rolling windows of predictive likelihoods (these are the weights available for pseudo-real time forecasting).¹⁰ Figure 4 presents these weights for the models of core PCE inflation (we omit the corresponding figure for GDP inflation in the interest of brevity). In the mid-1980s, the random walk trend model (with constant volatility) receives the most weight, peaking at 77% in mid-1986. From the mid-1990s until about 2009, the constant volatility models receive essentially no weight. But in the last few years of the sample, the PTR trend and TVP models with constant volatility receive some weight. From the late 1980s onward, most of the stochastic volatility models receive at least some weight, with rankings that move around over time. Starting in about 2006, the weight given to the local level-SV model rises sharply, to as much as 90%, before trailing off in the last few years of the forecast sample. Section 6 will examine how using these rolling sample-determined weights affect forecast accuracy.

Collectively, these results on model fit offer a somewhat different and more complicated picture than do the results of Stock and Watson (2007) and Faust and Wright (2011). Stock and Watson (2007) find the local level-SV model to yield the most accurate point forecasts, while Faust and Wright (2011) conclude that, for point forecasts, a survey-based trend works best.¹¹ In our estimates for core PCE inflation, the local level-SV model best fits the full sample, but the PTR trend-SV model fits the data almost as well. In our estimates for GDP inflation, the TVP-SV model fits the data best, and several others fit better than the local level-SV model. Perhaps even more importantly, the rankings of models have changed quite a bit over time. Accordingly, it seems difficult to say with any generality that a single model of trend inflation best fits the data. Some, although not all, of the differences among the Stock-Watson and Faust-Wright findings and ours may be attributable to the broader set of trend specifications considered in this paper. Other differences could be due to this section's use of Bayesian, rather than frequentist, or point forecast-based, concepts of model fit. We take up forecast accuracy later in this section.

5.3 Trend estimates

Figures 5 and 6 present, for core PCE inflation, estimates of trend inflation obtained from each of the models considered (specifically, posterior medians along with 70 percent credible sets). In the interest of chart readability, most of the chart panels provide just a single trend

¹⁰Using 5-year rolling windows of predictive likelihoods yielded similar results.

¹¹The Blue Chip and SPF long-run inflation expectations are quite similar. In a check of our findings, we obtained very similar results for models using Blue Chip instead of SPF to measure trend inflation.

series (median estimate) along with its credible set. The upper left panel, for the constant trend specification, reports the trend estimate from that model along with actual inflation and the PTR trend, the survey-based measure of long-run inflation expectations. The PTR trend is repeated in the upper right panel, which also provides the trend defined as last period's inflation rate.

Those models that estimate a time-varying trend (local level, random walk trend, and TVP) imply considerable variation over time in trend inflation, with non-trivial differences across models. The local level model yields a trend that is considerably more variable than the PTR measure of trend. Indeed, the local level model attributes much of the movement in actual inflation to trend shifts. However, allowing time-variation in volatilities results in a local level trend that, since the mid-1980s, is considerably smoother than the trend estimated from the local level model with constant volatility. Compared to the local level model, the random walk trend specification yields a broadly similar trend. At least in the model with constant volatility, the trend from the random walk trend model is smoother than the estimate from the local level model. Broadly, the random walk trend estimate is comparable to PTR, but with some considerable divergences over time. For example, in the late 1960s, the random walk trend estimate rose faster than PTR did and then remained at a higher level than PTR. But in the late 1970s, the PTR measure rose to a higher level than the random walk trend estimate, and remained at a higher level until about 1990. In the past few years of the sample, the random walk trend edged up (more so in the model with stochastic volatility than constant volatility), while PTR remained steady. The TVP models yield trend time series that are smoother than the other econometric estimates and PTR, but with similar contours. Of course, the estimates of trend from these models are considerably different from the trends assumed for the π_{t-1} trend model and estimated in the constant trend specification.

Figure 7 presents the results of accounting for model uncertainty by averaging core PCE trend estimates across models, based on the full-sample predictive likelihoods (i.e., on an ex post basis). As noted above, we obtain the density of model-average forecasts by sampling from the individual model densities, with a mixture approach based on the model weights. Given the large weights the PTR trend-SV and local level-SV models receive for the full sample, averaging all models based on the full-sample likelihoods yields a trend estimate that looks like a weighted average of the trends from these two specifications. Consequently,

prior to the mid-1980s stabilization of inflation, trend inflation is less variable than actual inflation but more variable than the PTR trend. But based on the changes in model fit over time described above, some might prefer to give more models some weight. For instance, in the forecasting literature (point and density forecasts), equal model weights often perform as well as or better than score-based weights (e.g., Kascha and Ravazzolo 2010 and Mazzi, Mitchell, and Montana 2010). Accordingly, the lower panel of Figure 7 presents a trend estimate obtained by applying equal weights to all of the models with stochastic volatility. This approach yields a broadly similar, but modestly more variable, trend. For example, while the likelihood-based average shows trend inflation to have been largely constant in recent years, the simple average shows some upward drift in trend late in the sample, as well as a very recent dip.

In the interest of brevity, Figures 8 and 9 present a reduced set of results for GDP inflation. Qualitatively, the estimates of trends from individual models with stochastic volatility in Figure 8 are similar to those for core PCE inflation (Figure 6). Those models that estimate a time-varying trend (local level, random walk trend, and TVP) imply considerable variation over time in trend inflation, with non-trivial differences across models. The local level model yields a trend that is considerably more variable than the PTR measure of trend. Moreover, the local level-SV estimate of trend in GDP inflation is more variable than the local level-SV estimate of trend in core PCE inflation. Again, the TVP models yield trend time series that are smoother than the other econometric estimates and PTR, but with similar contours. The model-average estimates of trend GDP inflation in Figure 9 generally show less of a rise in trend than do the average estimates of trend for core PCE inflation in Figure 7, because the GDP inflation estimate gives more weight to trend specifications that yield less of a rise in trend than does the local level-SV estimate that gets the largest weight in the core PCE inflation results.

5.4 Forecast results: core PCE inflation

Tables 4 and 5 present results for point forecasts of core PCE inflation, in the form of mean errors (Table 4) and RMSEs (Table 5).¹² The tables include results for each individual model and for three averages: one average that weights the model forecasts by rolling 10-year predictive likelihoods, a simple average of all forecasts, and a simple average of just

¹²As noted above, point forecasts are defined as posterior means of the posterior distribution of forecasts.

the models with SV.¹³ As noted above, the forecasts are combined by sampling from the appropriate mixtures of distributions. To facilitate comparisons, Table 5 reports ratios of RMSEs for each forecast relative to the RMSE of the local level-SV model; for the local level-SV model, the table provides (in the top row) the levels of the RMSEs.

To gauge statistical significance, we provide in Table 4 the results of a test of zero mean errors and in Table 5 the results of the Diebold and Mariano (1995) test for equal mean square forecast error (MSE), using asterisks to denote statistical significance based on standard normal critical values. The test of equal MSE takes the forecast of the local level-SV forecast as the benchmark. Following the recommendation of Clark and McCracken (2011c), to reduce the chances of spurious rejections at longer forecast horizons, we compute the t -tests with the Harvey, Leybourne, and Newbold (1997) small-sample adjustment of the variance that enters the test statistic.¹⁴ Our use of the Diebold-Mariano test with forecasts that are, in a few cases, nested is a deliberate choice.¹⁵ Monte Carlo evidence in Clark and McCracken (2011a,b) indicates that, with nested models, the Diebold-Mariano test compared against normal critical values can be viewed as a somewhat conservative (in the sense of tending to have size modestly below nominal size) test for equal accuracy in the finite sample. In this sense, for the limited cases in which the model of interest nests the local level-SV specification, these measures of statistical significance should be viewed as a rough guide.

Consistent with results in such studies as Clark (2011), the mean forecast errors in Table 4 show that every model consistently overstated inflation over the 1985-2011:Q2 sample. The bias increases with the forecast horizon. For most of the models, the bias in point forecasts is similar with constant volatility and stochastic volatility. The exception is the constant trend model, for which the longer horizon bias is considerably smaller with the stochastic volatility specification than with constant volatility. In many cases, the biases appear to be statistically significant. For example, with constant volatility models and the 3-, 4-, and 8-step horizons, the null of mean zero forecast errors is rejected (at least at a 10% significance level) for the constant trend, PTR trend, random walk trend, and TVP

¹³A forecast obtained from selecting the model with the highest predictive likelihood for each 10-year window didn't perform any better than those shown.

¹⁴The variance in the t -test is computed with a rectangular kernel and $h - 1$ lags. The Harvey, Leybourne, and Newbold (1997) adjustment is a multiplicative adjustment of the variance.

¹⁵We have also made a deliberate choice to treat the tests as two-sided, even though, with nested models, it is more natural to conduct a one-sided test, in which the benchmark model is taken as best unless the t -statistic implies a rejection in the positive tail of the normal distribution.

specifications.

At a high level, the RMSE results in Table 5 accord with Stock and Watson (2007): in data since 1985, a local level model with stochastic volatility is difficult to beat. In our sample, the local level-SV model is more accurate than every other model at every horizon (evidenced by RMSE ratios that are above 1), with a handful of exceptions that occur mostly at longer forecast horizons. In comparison, the constant trend model with constant volatility stands out as least accurate. That said, apart from the constant trend model, the differences in RMSE accuracy across all models are small to modest, peaking at about 13 percent. Be that as it may, at forecast horizons of 2 to 4 quarters, for a number of models — such as the random walk trend-SV and TVP-SV specifications — the accuracy advantage of the local level-SV model is statistically significant. The models using the PTR measure of trend are the only ones for which no rejection of equal accuracy occurs. In this sense, our results also accord with Faust and Wright (2011): the PTR trend-based models are also relatively good, statistically indistinguishable from the local level-SV specification.

In these point forecasting results, combining forecasts doesn't seem to offer much benefit, perhaps because the underlying model forecasts aren't sufficiently different. That said, some of the combination forecasts may be seen as providing some robustness: the equally weighted average of forecasts from models with stochastic volatility and predictive likelihood-weighted average forecasts are always about as good as the best model, even though the best model tends to vary across forecast horizons. For example, at a forecast horizon of 4 quarters, the local level-SV model is most accurate, with the PTR trend-SV model next most accurate (RMSE ratio of 1.012). At the 8-step horizon, the PTR trend-SV model is most accurate (RMSE ratio of 0.980). The predictive likelihood-weighted average of forecasts from the models with stochastic volatility yield corresponding RMSE ratios of 1.019 (4-step) and 0.982 (8-step). Another aspect of this robustness is that the Diebold-Mariano tests do not reject equal accuracy for the equally weighted average of forecasts from models with stochastic volatility and predictive likelihood-weighted average of forecasts versus the local level-SV benchmark.

For the purpose of assessing forecast performance in the broadest way possible — that is, assessing the calibration of density forecasts — Table 6 reports average log predictive scores. These average scores are the same as averages of log predictive likelihoods computed from simulated predictive densities. To facilitate comparisons across models, the top row

of each panel reports the average log score for the benchmark local level-SV model, and the remaining rows provide each model’s average log score relative to the benchmark (such that a positive number indicates a better score, while a negative number indicates a better score by the local level-SV model). To help provide a rough gauge of the significance of score differences, we provide in the table the results (using asterisks to denote statistical significance) of a test of the equality of mean scores, compared against standard normal critical values. Each test of equal mean score takes the local level-SV model as the benchmark. The test provides only a rough gauge in the sense that Amisano and Giacomini (2007) show the test to be normally distributed when applied to forecasts obtained by estimating models under the rolling scheme (constant sample size), while we estimated our models with the recursive scheme (allowing the sample to expand as forecasting moves forward in time). To reduce the chances of spurious rejections at longer forecast horizons, we compute the t -tests with the Harvey, Leybourne, and Newbold (1997) small-sample adjustment of the variance that enters the test statistic.¹⁶

The results in Table 6 indicate that, consistent with the findings of Clark (2011), the models with stochastic volatility score better than the corresponding models with constant volatility. At most forecast horizons, the inferiority of the scores of the constant volatility models relative to the local level-SV model is statistically significant. Among the models with stochastic volatility, at most horizons, the local level-SV model scores best, and the PTR trend-SV specification scores next best. While the PTR trend model’s average log scores are not quite as good as the scores of the local level-SV model, the differences do not appear to be statistically significant. The other models of trend are not as good as the two leading models. At horizons between 2 and 4 quarters, the null of equal average log score (versus the local level-SV benchmark) is rejected for the stochastic volatility versions of the constant trend, π_{t-1} trend, random walk trend, and TVP models. Overall, these density forecast results suggest that, in contrast to the point forecasting results of Stock and Watson (2007) and Faust and Wright (2011), neither the local level nor PTR measures of trend are superior; rather, the models are about equally good.

As in the case of the RMSE results, combining forecasts doesn’t seem to offer much help with average log scores, again perhaps because the underlying model forecasts are not sufficiently different to yield a benefit to pooling. However, by the average score metric,

¹⁶The variance in the t -test is computed with a rectangular kernel and $h - 1$ lags. The Harvey, Leybourne, and Newbold (1997) adjustment is a multiplicative adjustment of the variance.

too, averaging may be seen as helping with robustness, in the sense that the scores of the equally weighted average of forecasts from models with stochastic volatility and predictive likelihood-weighted average forecasts are often (not always) comparable to the scores of the better-performing individual forecasts.

Relatively large differences in model forecasts emerge with deflation probabilities (as indicated above, deflation is defined as annual average inflation less than 0), reported in Figures 10 and 11. The models with constant volatilities (top panel of Figure 10) consistently yield higher probabilities of deflation than the corresponding models with stochastic volatilities (lower panel). The stochastic volatility estimates for the PTR trend model shown in Figure 1 suggest this pattern in deflation probabilities occurs because, following the Great Moderation, the variances of shocks to inflation were smaller than they were on average across the whole sample.

The probability of deflation is generally lower in the two models with the best-calibrated forecast densities (the models with the best average log scores), the local level-SV and PTR trend-SV models, than in the others. The probability of deflation is generally highest in the π_{t-1} trend model. Broadly, across all estimates (with both constant volatility and stochastic volatility), deflation probabilities tend to be lower in models for which the trend estimates tend to be smoother. A more variable trend estimate means a model is more likely to imply low inflation, which increases the chances of deflation. These results highlight the importance of careful modeling of trend inflation in assessing the risks of deflation associated with periods of low inflation, such as 2003 and 2009.

The role of the inflation trend becomes even more important when the inflation forecasting model includes the unemployment gap. Figure 11 presents deflation probabilities for the π_{t-1} trend and PTR trend models, in their AR forms and in the bivariate forms that also include the unemployment gap.¹⁷ Not surprisingly, in the past two recessions, the models including unemployment yield greater probabilities of deflation — dramatically greater for much of the 2009-2011 period. But even with unemployment in the model, the forecast remains sensitive to the specification of trend inflation. For example, in 2011:Q1, the bivariate version of the π_{t-1} trend model yielded a deflation probability of 61%, while the bivariate version of the PTR trend model yielded a deflation probability of 25%. How-

¹⁷The predictive likelihoods in Table 3 show that, in most cases, the addition of the unemployment gap causes a modest reduction in model fit. However, Dotsey, Fujita, and Stark (2011) find some evidence that unemployment has more predictive content when the economy is weak (not when it is strong, or unconditionally).

ever, the probability estimate from the PTR-based model should probably be viewed as superior to the estimate from the π_{t-1} trend model, given that the predictive likelihoods in Table 3 indicate the former fits the data much better than the latter.¹⁸

5.5 Forecast results: GDP inflation

In the interest of brevity, for GDP inflation we present only RMSEs and log predictive scores, in Tables 7 and 8, respectively. Broadly, the results for GDP inflation are similar to those for core PCE inflation: for the GDP measure of inflation, models with stochastic volatility dominate constant volatility models in the broadest measure of forecast performance (average log scores), and it is difficult to say that one model of trend inflation is best.

In the RMSE results for GDP inflation provided in Table 7, the constant trend-SV and PTR trend-SV models are generally the most accurate. While the local level-SV model of Stock and Watson (2007) has the lowest RMSE at the 1-step horizon, other model forecasts are very nearly as accurate, and at horizons of 3 quarters and more, a number of other model forecasts beat the accuracy of the local level-SV benchmark. However, the differences across models are fairly modest — modest enough that the null of equal accuracy against the local level-SV benchmark is rejected less often in the GDP inflation results than in the core PCE inflation estimates. Averaging of forecasts seems slightly more helpful for GDP inflation than for core PCE inflation. The predictive likelihood-based average and equally weighted average of forecasts from models with stochastic volatility consistently improve on the accuracy of the benchmark local level-SV forecast, and the accuracy of these averages is comparable to the accuracy of the best-performing model in each sample and horizon.

In the average log score results of Table 8, the models with stochastic volatility uniformly score better than the corresponding models with constant volatility. At most forecast horizons, the inferiority of the scores of the constant volatility models relative to the local level-SV model is statistically significant. That said, the local level-SV model has slightly lower scores — but not by a statistically significant margin — than most of the other trend specifications that include stochastic volatility. The π_{t-1} trend model scores best at horizons of 4 quarters and less, while the TVP-SV model scores best at most longer horizons. Accordingly, in terms of density forecasts of GDP inflation, several of the trend

¹⁸Across all forecast horizons, the bivariate model with the PTR trend also yields slightly better average log scores (and better RMSEs) than the bivariate model with the π_{t-1} trend.

specifications are about equally good. From this density perspective, averaging forecasts may be seen as helping with robustness, in the sense that the scores of the equally weighted average of forecasts from models with stochastic volatility and predictive likelihood-weighted average forecasts are generally comparable to the scores of the better-performing individual forecasts.

6 Conclusions

This paper uses Bayesian methods to assess alternative models of trend inflation. Our set of models incorporates significant differences in the trend specification: AR with constant trend; AR with trend equal to last period’s inflation rate; local level model; AR with random walk trend; AR with trend equal to the long-run expectation from the Survey of Professional Forecasters (SPF); and AR with time-varying parameters. Finally, we consider versions of the models with constant shock variances and with time-varying shock variances (stochastic volatility). To further highlight the importance of trend specification to inflation inferences, in assessing deflation probabilities we also consider some specifications that include an unemployment gap.

We first use Bayesian metrics to compare the fits of these alternative models of inflation. We then use Bayesian methods of model averaging to account for uncertainty surrounding the model of trend inflation, to obtain an alternative estimate of trend inflation in the U.S. and to generate medium-term, model-average forecasts of inflation. We use predictive likelihoods to weight each model and forecast and construct probability-weighted average estimates of trend inflation and inflation forecasts. For forecasting, we consider not only point predictions but also density forecasts (specifically, average log predictive scores and deflation probabilities).

Our analysis yields two broad results. First, in model fit and density forecast accuracy, models with stochastic volatility consistently dominate those with constant volatility. Second, for the specification of trend inflation, it is difficult to say that one model of (time-varying) trend inflation is the best. Among alternative models of the trend in core PCE inflation, the local level specification of Stock and Watson (2007) and the SPF-based trend specification are about equally good. In full-sample fit, the local level specification of trend is best, with the SPF-based trend specification next best and reasonably close. However, model fit has evolved considerably over the sample; the local level model hasn’t always been

best. In out-of-sample forecasting, the local level model is slightly more accurate than the SPF-based trend model, but not by a significant margin. Among competing models of trend GDP inflation, several trend specifications seem to be about equally good. In full-sample fit, a model with time-varying parameters is best, with the lagged inflation trend and SPF-based trend models next and reasonably close. In forecasting, many of the models perform similarly in both point and density prediction of GDP inflation. Overall, we interpret these results as highlighting the importance of considering predictions from a range of models that fit the historical data about as well as one another, particularly in evaluating risks such as the chance of deflation.

Our findings on alternative trend models provide a somewhat different and more complicated picture than do the findings of Stock and Watson (2007) and Faust and Wright (2011), which suggest particular models to be clearly best (local level with stochastic volatility in the case of Stock-Watson and survey-based trend in the case of Faust-Wright). Some of the differences among the Stock-Watson and Faust-Wright findings and ours may be attributable to the broader set of trend specifications considered in this paper. Other differences could be due to our use of Bayesian (rather than frequentist) concepts of model fit and our analysis of density (not just point) forecasts.

7 Appendix

The first section of this appendix details the numbers of burn-in draws and skip intervals used to estimate each model and forecast. The next five sections provide the MCMC algorithms used to estimate and forecast with the following models, all with stochastic volatility (with the models ordered to simplify presentation of the algorithms): (1) the AR model with time-varying parameters (TVP); (2) the AR model with constant parameters (π_{t-1} trend, PTR trend); (3) the AR model with a steady state prior (constant trend); (4) local level; and (5) random walk trend. In the interest of brevity, we omit details for the simpler algorithms used to estimate models with constant volatility (in which inverse Gamma priors and posteriors are used for the innovation to the inflation equation). Section 7.7 details the priors used for each model.

For generality and simplicity in presentation, in the appendix we use y_t to refer to the endogenous variable (inflation, or inflation less trend in some specifications), and we use X_t to denote the vector of variables on the right-hand side (lags of inflation and, if called for in a given model, a constant).

7.1 Details of MCMC Samples

Burn and Skip Settings in MCMC Algorithms		
	burn sample (# draws)	skip interval (between draws)
constant trend	5000	5
π_{t-1} trend	5000	5
PTR trend	5000	5
local level	10000	15
random walk trend	10000	15
TVP	10000	15
constant trend-SV	10000	15
π_{t-1} trend-SV	10000	15
PTR trend-SV	10000	15
local level-SV	10000	15
random walk trend-SV	10000	25
TVP-SV	10000	15

7.2 AR Model with TVP and Stochastic Volatility

We estimate the model with a four-step Metropolis-within-Gibbs MCMC algorithm, following studies such as Cogley and Sargent (2005) and Primiceri (2005). The Metropolis step is used for the estimation of stochastic volatility, following Cogley and Sargent (2005) in their use of the Jacquier, Polson, and Rossi (1994) algorithm. In this case, y_t denotes the level of inflation.

Step 1: Draw the time series of AR coefficients b_t conditional on the history of λ_t , Q , and ϕ .

As detailed in Primiceri (2005), drawing the AR coefficients involves using the Kalman filter to move forward in time, a backward smoother to obtain posterior means and variances of the coefficients at each point in time, and then drawing coefficients from the posterior normal distribution. For the backward smoothing, we use the algorithm of Durbin and Koopman (2002) instead of the Carter and Kohn (1994) algorithm used by Primiceri. Koop and Korobilis (2010) note that either algorithm can be used for VARs with time-varying parameters; the software we use makes the Durbin and Koopman (2002) algorithm faster. The mean and variance of the period 0 coefficient vector used in the smoother is fixed at the prior mean and variance described below.

Step 2: Draw the variance λ_t conditional on the history of b_t , Q , and ϕ .

Following Cogley and Sargent (2005) and Primiceri (2005), the AR model can be rewritten as

$$y_t - X_t' b_t \equiv \tilde{y}_t = \lambda_t^{0.5} \epsilon_t, \quad (12)$$

where $\epsilon_t \sim N(0, 1)$. Taking logs of the squares yields

$$\log \tilde{y}_t^2 = \log \lambda_t + \log \epsilon_t^2. \quad (13)$$

The conditional volatility process is

$$\log(\lambda_t) = \log(\lambda_{t-1}) + \nu_t, \quad \nu_t \sim N(0, \phi). \quad (14)$$

The estimation of the time series of λ_t uses the measured $\log \tilde{y}_t^2$ and Cogley and Sargent's (2005) Metropolis algorithm; see Cogley and Sargent for further detail.

Step 3: Draw the variance matrix Q conditional on the history of b_t , the history of λ_t , and ϕ .

Following Cogley and Sargent (2005) and Primiceri (2005), the sampling of Q , the variance-covariance matrix of innovations to the AR coefficients, is based on inverse Wishart

priors and posteriors. The scale matrix of the posterior distribution is the sum of the prior mean \times the prior degrees of freedom and $\sum_{t=1}^T \hat{n}_t \hat{n}_t'$, where \hat{n}_t denotes the innovations to the posterior draws of coefficients obtained in step 1.

Step 4: Draw the variance ϕ conditional on the history of b_t , the history of λ_t , and Q .

Following Cogley and Sargent (2005), the sampling of ϕ , the variance of the innovation to the log variance associated with the inflation equation, is based on inverse Gamma priors and posteriors. The scale factor of the posterior distribution is the sum of the prior mean \times the prior degrees of freedom and $\sum_{t=1}^T \hat{\nu}_t^2$, where $\hat{\nu}_t$ denotes the innovations to the posterior draw of the volatility obtained in step 2.

Forecast density:

To generate draws of forecasts, we follow Cogley, Morozov, and Sargent's (2005) approach to simulating the predictive density. Let H denote the maximum forecast horizon considered. From a forecast origin of period T , for each retained draw of the time series of b_t up through T , λ_t up through T , Q , and ϕ , we: (1) draw innovations to coefficients for periods $T+1$ through $T+H$ from a normal distribution with variance-covariance matrix Q and use the random walk structure to compute b_{T+1}, \dots, b_{T+H} ; (2) draw innovations to log volatility for periods $T+1$ through $T+H$ from a normal distribution with variance ϕ and use the random walk model of $\log \lambda_{t+h}$ to compute $\lambda_{T+1}, \dots, \lambda_{T+H}$; (3) draw innovations to the equation for y_{T+h} , $h = 1, \dots, H$, from a normal distribution with variance λ_{T+h} , and use the autoregressive structure of the model along with the time series of coefficients b_{T+h} to obtain draws of y_{T+h} , $h = 1, \dots, H$. The resulting draws of y_{T+h} are used to compute the forecast statistics of interest (mean errors and RMSE based on the posterior mean, etc.).

7.3 AR Model with Steady State Prior and Stochastic Volatility

This model, used to estimate the constant trend specification, takes the form

$$y_t - \psi d_t = b(L)(y_{t-1} - \psi d_{t-1}) + v_t, \quad (15)$$

where d_t denotes a vector containing the deterministic mean variables, which in this application is just a constant. In this case, y_t denotes the level of inflation. We estimate the model with a four-step Metropolis-within-Gibbs MCMC algorithm, first used (with a VAR) in Clark (2011).

Step 1: Draw the slope coefficients b conditional on ψ , the history of λ_t , and ϕ .

For this step, the AR model is recast in demeaned form, using $Y_t = y_t - \psi d_t$ and letting X_t denote the vector containing the relevant lags of $y_t - \psi d_t$:

$$Y_t = X_t' b + v_t, \text{ var}(v_t) = \lambda_t. \quad (16)$$

The vector of coefficients is sampled from a normal posterior distribution with mean $\bar{\mu}_b$ and variance $\bar{\Omega}_b$, based on prior mean μ_b and Ω_b , where:

$$\bar{\Omega}_b^{-1} = \Omega_b^{-1} + \sum_{t=1}^T (\lambda_t^{-1} X_t X_t') \quad (17)$$

$$\bar{\mu}_b = \bar{\Omega}_b \left\{ \text{vec} \left(\sum_{t=1}^T \lambda_t^{-1} Y_t X_t' \right) + \Omega_b^{-1} \mu_b \right\}. \quad (18)$$

Step 2: Draw the steady state coefficients ψ conditional on b , the history of λ_t , and ϕ .

For this step, the AR model is rewritten as

$$q_t = (1 - b(L))\psi d_t + v_t, \text{ where } q_t \equiv (1 - b(L))y_t. \quad (19)$$

The dependent variable q_t is obtained by applying to y_t the lag polynomial estimated with the draw of the b coefficients obtained in step 1. Letting k denote the number of lags in the model, the right-hand side term $(1 - b(L))\psi d_t$ simplifies to $\Theta \bar{d}_t$, where, as in Villani (2009) with some modifications, \bar{d}_t contains current and lagged values of the elements of d_t , and Θ is defined such that $\text{vec}(\Theta) = U \text{vec}(\psi)$:

$$\bar{d}_t = (d_t', -d_{t-1}', -d_{t-2}', \dots, -d_{t-k}')' = (1, -1, \dots, -1)' \quad (20)$$

$$U = \begin{pmatrix} 1 \\ b_1 \\ b_2 \\ \vdots \\ b_k \end{pmatrix}. \quad (21)$$

The vector of coefficients ψ is sampled from a normal posterior distribution with mean $\bar{\mu}_\psi$ and variance $\bar{\Omega}_\psi$, based on prior mean μ_ψ and Ω_ψ , where:

$$\bar{\Omega}_\psi^{-1} = \Omega_\psi^{-1} + U' \left\{ \sum_{t=1}^T (\bar{d}_t \bar{d}_t' \lambda_t^{-1}) \right\} U \quad (22)$$

$$\bar{\mu}_\psi = \bar{\Omega}_\psi \left\{ U' \text{vec} \left(\sum_{t=1}^T \lambda_t^{-1} q_t \bar{d}_t' \right) + \Omega_\psi^{-1} \mu_\psi \right\}. \quad (23)$$

Step 3: Draw the time series λ_t conditional on b , ψ , and ϕ .

Following Cogley and Sargent (2005), the AR model can be rewritten as

$$(1 - b(L))(y_t - \psi d_t) \equiv \tilde{y}_t = \lambda_t^{0.5} \epsilon_t, \quad (24)$$

where $\epsilon_t \sim N(0, 1)$. Taking logs of the squares yields

$$\log \tilde{y}_t^2 = \log \lambda_t + \log \epsilon_t^2. \quad (25)$$

Using the measured $\log \tilde{y}_t^2$, this step proceeds as with step 2 of the TVP-stochastic volatility algorithm, except that the AR coefficients are constant.

Step 4: Draw the innovation variance conditional on b , ψ , and the history of λ_t .

This step proceeds as with step 4 of the TVP-stochastic volatility algorithm, except that the AR coefficients are constant.

Forecast density:

The simulation of the predictive density follows the steps described above for the TVP model, except that the steps for simulating time series of the AR coefficients are eliminated.

7.4 AR Model with Stochastic Volatility

We estimate the model with a three-step Metropolis-within-Gibbs MCMC algorithm. In the AR-SV model formulation, y_t denotes the level of inflation less last period's inflation rate in the case of the π_{t-1} trend model and the level of inflation less PTR in the case of the PTR trend model. For the bivariate models used in the paper (the π_{t-1} trend and PTR trend models that include an unemployment gap), see Clark and Davig (2011) for the VAR version of the AR algorithm described in this section.

Step 1: Draw the AR coefficient vector b conditional on the history of λ_t and ϕ .

This step relies on the GLS solution also used for the AR model with a steady state prior. The vector of coefficients is sampled from a multivariate normal posterior distribution with mean $\bar{\mu}_B$ and variance $\bar{\Omega}_B$, based on prior mean μ_B and Ω_B . The posterior mean and variance are:

$$\bar{\Omega}_B^{-1} = \Omega_B^{-1} + \sum_{t=1}^T (\lambda_t^{-1} X_t X_t') \quad (26)$$

$$\bar{\mu}_B = \bar{\Omega}_B \left\{ \text{vec} \left(\sum_{t=1}^T \lambda_t^{-1} y_t X_t' \right) + \Omega_B^{-1} \mu_B \right\}. \quad (27)$$

Step 2: Draw the variance time series λ_t conditional on b and ϕ .

This step proceeds as with step 2 of the TVP-stochastic volatility algorithm, except that the AR coefficients are constant.

Step 3: Draw the variance ϕ , conditional on b and the history of λ_t .

This step proceeds as with step 4 of the TVP-stochastic volatility algorithm, except that the AR coefficients are constant.

Forecast density:

The simulation of the predictive density follows the steps described above for the TVP model, except that the steps for simulating time series of the AR coefficients are eliminated.

7.5 Local Level Model with Stochastic Volatility

We estimate the model with a five-step Metropolis-within-Gibbs MCMC algorithm.

Step 1: Using the state-space representation of the model, draw the time series of trend (π_t^*) and noise components (v_t) conditional on the history of $\lambda_{v,t}$, history of $\lambda_{n,t}$, ϕ_v , and ϕ_n .

In this step we use the conventional Kalman forward filter and the Durbin and Koopman (2002) backward smoother to draw the trend and noise time series. The mean and variance of the period 0 trend used in the smoother is fixed at the prior mean and variance described below.

Step 2: Draw the (trend) innovation variance time series $\lambda_{n,t}$ conditional on the time series of π_t^* and v_t , ϕ_v , and ϕ_n

This step uses logs of the squares of the time series of trend innovations n_t , $t = 1, \dots, T$ obtained from step 1's smoother:

$$\log n_t^2 = \log \lambda_{n,t} + \log \epsilon_{n,t}^2. \quad (28)$$

The estimation of the time series of $\lambda_{n,t}$ uses the measured $\log n_t^2$ and Cogley and Sargent's (2005) version of the Metropolis algorithm of Jacquier, Polson, and Rossi (1994).

Step 3: Draw the (noise) innovation variance time series $\lambda_{v,t}$ conditional on the time series of π_t^* and v_t , ϕ_v , and ϕ_n

Using the time series of innovations v_t , $t = 1, \dots, T$ obtained from step 1's smoother, this step relies on

$$\log v_t^2 = \log \lambda_{v,t} + \log \epsilon_{v,t}^2. \quad (29)$$

The estimation of the time series of $\lambda_{v,t}$ uses the measured $\log v_t^2$ and Cogley and Sargent's (2005) version of the Metropolis algorithm of Jacquier, Polson, and Rossi (1994).

Step 4: Draw the variance ϕ_n conditional on the latent states and other parameters:

The sampling of ϕ_n , the variance of the innovation to the log variance associated with the trend component, is based on inverse Gamma priors and posteriors. The scale factor of the posterior distribution is the sum of the prior mean \times the prior degrees of freedom and $\sum_{t=1}^T \hat{\nu}_{n,t}^2$, where $\hat{\nu}_{n,t}$ denotes the innovations to the posterior draw of the volatility obtained in step 2.

Step 5: Draw the variance ϕ_v conditional on the latent states and other parameters:

The sampling of ϕ_v , the variance of the innovation to the log variance associated with the noise component, is based on inverse Gamma priors and posteriors. The scale factor of the posterior distribution is the sum of the prior mean \times the prior degrees of freedom and $\sum_{t=1}^T \hat{\nu}_{v,t}^2$, where $\hat{\nu}_{v,t}$ denotes the innovations to the posterior draw of the volatility obtained in step 3.

Forecast density:

To generate draws of forecasts, we use an appropriately modified version of Cogley, Morozov, and Sargent's (2005) approach to simulating the predictive density. Let H denote the maximum forecast horizon considered. From a forecast origin of period T , for each retained draw of the time series of π_t^* , $\lambda_{v,t}$, and $\lambda_{n,t}$ up through T , ϕ_v , and ϕ_n , we: (1) draw innovations to noise volatility for periods $T+1$ through $T+H$ from a normal distribution with variance ϕ_v and use the random walk model of $\log \lambda_{v,t+h}$ to compute $\lambda_{v,T+1}, \dots, \lambda_{v,T+H}$; (2) draw innovations to trend volatility for periods $T+1$ through $T+H$ from a normal distribution with variance ϕ_n and use the random walk model of $\log \lambda_{n,t+h}$ to compute $\lambda_{n,T+1}, \dots, \lambda_{n,T+H}$; (3) draw innovations to π_{T+h}^* , $h = 1, \dots, H$, from a normal distribution with variance $\lambda_{n,T+h}$, and compute the trend π_{T+h}^* , $h = 1, \dots, H$ from the random walk structure of the trend; (4) draw noise innovations v_{T+h}^* , $h = 1, \dots, H$, from a normal distribution with variance $\lambda_{v,T+h}$; and (5) form a draw of y_{T+h} , $h = 1, \dots, H$ from the sum of the draw of trend (π_{T+h}^*) and noise (v_{T+h}^*). The resulting draws of y_{T+h} are used to compute the forecast statistics of interest (mean errors and RMSE based on the posterior mean, etc.).

7.6 Random Walk Trend Model with Stochastic Volatility

We estimate the model with a five-step Metropolis-within-Gibbs MCMC algorithm.

Step 1: Draw the AR coefficients b conditional on the history of π_t^* and λ_t , σ_n^2 , and ϕ .

For this step, y_t denotes detrended inflation, $\pi_t - \pi_t^*$. Using the time series of y_t , drawing

the AR coefficients relies on the same GLS solution used above for the AR-SV model.

Step 2: Using the state-space representation of the model, draw the time series of trend (π_t^*) conditional on b , the history of λ_t , σ_n^2 , and ϕ .

In this step, the measurement equation relates $y_t - b(L)y_{t-1}$ to a state vector containing π_t^* and the appropriate lags. We use the conventional Kalman forward filter and the Durbin and Koopman (2002) backward smoother to draw the trend time series. The mean and variance of the period 0 trend used in the smoother is fixed at the prior mean and variance described below.

Step 3: Using innovations to trend obtained in step 2, draw the variance of innovations to trend σ_n^2 conditional on b , the history of π_t^* and λ_t , and ϕ .

The sampling of σ_n^2 is based on an inverse Gamma prior and posterior. The scale factor of the posterior distribution is the sum of the prior mean \times the prior degrees of freedom and $\sum_{t=1}^T \hat{n}_t^2$, where \hat{n}_t denotes the innovations to the posterior draw of the trend obtained in step 2.

Step 4: Draw the variance λ_t conditional on b , the history of π_t^* , σ_n^2 , and ϕ .

Letting y_t denote detrended inflation and X_t denote the appropriate lags of y_t , the model can be rewritten as

$$y_t - X_t' b \equiv \tilde{y}_t = \lambda_t^{0.5} \epsilon_t, \quad (30)$$

where $\epsilon_t \sim N(0, 1)$. Taking logs of the squares yields

$$\log \tilde{y}_t^2 = \log \lambda_t + \log \epsilon_t^2. \quad (31)$$

This step proceeds as with step 2 of the TVP-stochastic volatility algorithm, except that the AR coefficients are constant.

Step 5: Draw the variance ϕ conditional on b , the history of π_t^* and λ_t , and σ_n^2 .

This step proceeds as with step 4 of the TVP-stochastic volatility algorithm, except that the AR coefficients are constant.

Forecast density:

To generate draws of forecasts, we follow an approach similar to those described above. Let H denote the maximum forecast horizon considered. From a forecast origin of period T , for each retained draw of b , the time series of π_t^* and λ_t up through T , σ_n^2 , and ϕ , we: (1) draw innovations to the trend π_t^* for periods $T+1$ through $T+H$ from a normal distribution with variance σ_n^2 and use the random walk structure to compute $\pi_{T+1}^*, \dots, \pi_{T+H}^*$; (2) draw

innovations to log volatility for periods $T + 1$ through $T + H$ from a normal distribution with variance ϕ and use the random walk model of $\log \lambda_{t+h}$ to compute $\lambda_{T+1}, \dots, \lambda_{T+H}$; (3) draw innovations to the equation for y_{T+h} , $h = 1, \dots, H$, from a normal distribution with variance λ_{T+h} , and use the autoregressive structure of the model along with the time series of trends π_{T+h}^* to obtain draws of y_{T+h} , $h = 1, \dots, H$. The resulting draws of y_{T+h} are used to compute the forecast statistics of interest (mean errors and RMSE based on the posterior mean, etc.).

7.7 Priors

As indicated in the text, in estimating models for the sample of 1960-2011:Q2 (and, for forecasting purposes, other samples that begin in 1960 and end before 2011:Q2), we use priors that are taken from the posteriors obtained from a sample of 1947-59. We use data for 1913-46 as a training sample from which we obtain some estimates that provide the priors used in the model estimates for 1947-59. We do so for each model (that is, for each model, we generate its prior by estimating the model for the 1947-59 sample). In this section, we provide more information on the priors (corresponding to posteriors for the 1947-59 sample) used to generate our estimates for the 1960-2011 sample. The priors use normal (N), inverse Gamma (IG), and inverse Wishart (IW) distributions.

7.7.1 Constant trend

For the AR(2) model with a constant rate of trend inflation, let π^* denote trend inflation, b denote the vector of AR coefficients, and σ_v^2 denote the error variance. The (independent) priors for the model are:

$$\begin{aligned}\pi^* &\sim N(\hat{\pi}_{TS}^*, 4 \cdot \hat{V}_{TS}(\pi^*)) \\ b &\sim N(\hat{b}_{TS}, 4 \cdot \hat{V}_{TS}(b)) \\ \sigma_v^2 &\sim IG(3 \cdot \hat{\sigma}_{v,TS}^2, 3),\end{aligned}$$

where $\hat{\pi}_{TS}^*$, \hat{b}_{TS} , and $\hat{\sigma}_{v,TS}^2$ denote posterior means obtained (for the same model) from the 1947-59 sample and $\hat{V}_{TS}(\pi^*)$ and $\hat{V}_{TS}(b)$ denote corresponding posterior variances.

In the specification with stochastic volatility, the priors are:

$$\begin{aligned}\pi^* &\sim N(\hat{\pi}_{TS}^*, 4 \cdot \hat{V}_{TS}(\pi^*)) \\ b &\sim N(\hat{b}_{TS}, 4 \cdot \hat{V}_{TS}(b))\end{aligned}$$

$$\begin{aligned}\phi &\sim IG(5 \cdot \widehat{\phi}_{TS}, 5) \\ \log \lambda_0 &\sim N(\log \widehat{\lambda}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\log \lambda_0)),\end{aligned}$$

where $\widehat{\pi}_{TS}^*$, \widehat{b}_{TS} , and $\widehat{\phi}_{TS}$ denote posterior means obtained (for the same model) from the 1947-59 sample and $\widehat{V}_{TS}(\pi^*)$ and $\widehat{V}_{TS}(b)$ denote corresponding posterior variances. We use λ_0 to denote the value of volatility in 1959:Q4; in the 1960-2011 estimation, our prior mean and variance for this initial value are set to the posterior mean and variance (scaled by 4) of the log volatility estimate for this quarter obtained from the estimation for 1947-59.

7.7.2 π_{t-1} trend

Using the same basic notation, the priors for the AR(2) model with the π_{t-1} trend specification are:

$$\begin{aligned}b &\sim N(\widehat{b}_{TS}, 4 \cdot \widehat{V}_{TS}(b)) \\ \sigma_v^2 &\sim IG(3 \cdot \widehat{\sigma}_{v,TS}^2, 3).\end{aligned}$$

In the specification with stochastic volatility, the priors are:

$$\begin{aligned}b &\sim N(\widehat{b}_{TS}, 4 \cdot \widehat{V}_{TS}(b)) \\ \phi &\sim IG(5 \cdot \widehat{\phi}_{TS}, 5) \\ \log \lambda_0 &\sim N(\log \widehat{\lambda}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\log \lambda_0)).\end{aligned}$$

7.7.3 PTR trend

For the AR(2) model with the PTR trend, we set the priors for b and σ_v^2 at the values used for the constant trend model. In the specification with stochastic volatility, we take the prior from the model with constant trend and stochastic volatility.

7.7.4 Local level

The (independent) priors for the local level model are:

$$\begin{aligned}\pi_0^* &\sim N(\widehat{\pi}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\pi_0^*)) \\ \sigma_v^2 &\sim IG(3 \cdot \widehat{\sigma}_{v,TS}^2, 3) \\ \sigma_n^2 &\sim IG(5 \cdot \widehat{\sigma}_{n,TS}^2, 5),\end{aligned}$$

where π_0^* denotes the value of trend inflation in 1959:Q4 that is needed as an initial value in the state-space representation of the model (for 1960-2011). In the 1960-2011 estimation,

our prior mean and variance for this initial value are set to the posterior mean and variance (scaled by 4) of the trend estimate for this quarter obtained from the estimation for 1947-59.

In the case of the local level specification with stochastic volatility, the priors are:

$$\begin{aligned}
\pi_0^* &\sim N(\widehat{\pi}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\pi_0^*)) \\
\phi_v &\sim IG(5 \cdot \widehat{\phi}_{v,TS}, 5) \\
\phi_n &\sim IG(5 \cdot \widehat{\phi}_{n,TS}, 5) \\
\log \lambda_{v,0} &\sim N(\log \widehat{\lambda}_{v,0,TS}, 4 \cdot \widehat{V}_{TS}(\log \lambda_{v,0})) \\
\log \lambda_{n,0} &\sim N(\log \widehat{\lambda}_{n,0,TS}, 4 \cdot \widehat{V}_{TS}(\log \lambda_{n,0})),
\end{aligned}$$

where $\lambda_{v,0}$ and $\lambda_{n,0}$ denote the values of noise and trend volatility in 1959:Q4. In the 1960-2011 estimation, our prior means and variances for these initial values are set to the posterior means and variances (scaled by 4) of the log volatility estimates for this quarter obtained from the estimation for 1947-59.

7.7.5 Random walk trend

The (independent) priors for the AR(2) model with random walk specification for trend inflation (π_t^*) are:

$$\begin{aligned}
b &\sim N(\widehat{b}_{TS}, 4 \cdot \widehat{V}_{TS}(b)) \\
\pi_0^* &\sim N(\widehat{\pi}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\pi_0^*)) \\
\sigma_v^2 &\sim IG(3 \cdot \widehat{\sigma}_{v,TS}^2, 3) \\
\sigma_n^2 &\sim IG(5 \cdot \widehat{\sigma}_{n,TS}^2, 5),
\end{aligned}$$

where π_0^* denotes the vector of values of trend in 1959:Q2–1959:Q4 that are needed as initial values in the state-space representation of the model (for 1960-2011). In the 1960-2011 estimation, our prior mean and variance for this vector of initial values is set to the posterior mean and variance (scaled by 4) of the estimates for these quarters obtained from the estimation for 1947-59.

With stochastic volatility, the priors are:

$$\begin{aligned}
b &\sim N(\widehat{b}_{TS}, 4 \cdot \widehat{V}_{TS}(b)) \\
\pi_0^* &\sim N(\widehat{\pi}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\pi_0^*)) \\
\sigma_n^2 &\sim IG(5 \cdot \widehat{\sigma}_{n,TS}^2, 5)
\end{aligned}$$

$$\begin{aligned}\phi &\sim IG(5 \cdot \widehat{\phi}_{TS}, 5) \\ \log \lambda_0 &\sim N(\log \widehat{\lambda}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\log \lambda_0)).\end{aligned}$$

7.7.6 TVP

The (independent) priors for the AR(2) model with TVP are:

$$\begin{aligned}b_0^* &\sim N(\widehat{b}_{0,TS}, 4 \cdot \widehat{V}_{TS}(b_0^*)) \\ Q &\sim IW(\widehat{Q}_{TS} \cdot 20, 20) \\ \sigma_v^2 &\sim IG(3 \cdot \widehat{\sigma}_{v,TS}^2, 3),\end{aligned}$$

where b_0^* denotes the value of the coefficient vector in 1959:Q4 that is needed as an initial value in the state-space representation of the model (for 1960-2011). In the 1960-2011 estimation, our prior mean and variance for this initial value are set to the posterior mean and variance (scaled by 4) of the estimate for 1959:Q4 obtained from the estimation for 1947-59.

With stochastic volatility, the priors are:

$$\begin{aligned}b_0^* &\sim N(\widehat{b}_{0,TS}, 4 \cdot \widehat{V}_{TS}(b_0^*)) \\ Q &\sim IW(\widehat{Q}_{TS} \cdot 20, 20) \\ \sigma_v^2 &\sim IG(3 \cdot \widehat{\sigma}_{v,TS}^2, 3), \\ \phi &\sim IG(5 \cdot \widehat{\phi}_{TS}, 5) \\ \log \lambda_0 &\sim N(\log \widehat{\lambda}_{0,TS}, 4 \cdot \widehat{V}_{TS}(\log \lambda_0)).\end{aligned}$$

References

- Amisano, Gianni, and Raffaella Giacomini (2007), “Comparing Density Forecasts via Weighted Likelihood Ratio Tests,” *Journal of Business and Economic Statistics* 25, 177-190.
- Atkeson, Andrew, and Lee E. Ohanian, (2001), “Are Phillips Curves Useful for Forecasting Inflation?” *Federal Reserve Bank of Minneapolis Quarterly Review* 25, 2-11.
- Bauwens, Luc, Gary Koop, Dimitris Korobilis, and Jeroen V.K. Rombouts (2011), “A Comparison of Forecasting Procedures for Macroeconomic Series: The Contribution of Structural Break Models,” manuscript, CIRANO.
- Beveridge, Stephen, and Charles R. Nelson (1981), “A New Approach to Decomposition of Economic Time Series into Permanent And Transitory Components with Particular Attention to Measurement of the Business Cycle,” *Journal of Monetary Economics* 7, 151-174
- Bjornland, Hilde, Karsten Gerdrup, Christie Smith, Anne Sofie Jore, and Leif Anders Thorsrud (2010), “Weights and Pools for a Norwegian Density Combination,” Norges Bank Working Paper 2010-06.
- Brayton, Flint, John M. Roberts, and John C. Williams (1999), “What’s Happened to the Phillips Curve?” manuscript, Board of Governors of the Federal Reserve System.
- Carter, C.K., and R. Kohn (1994), “On Gibbs Sampling for State Space Models,” *Biometrika* 81, 541-553.
- Clark, Todd E. (2011), “Real-Time Density Forecasts from BVARs with Stochastic Volatility,” *Journal of Business and Economic Statistics* 29, 327-341.
- Clark, Todd E., and Troy Davig (2011), “Decomposing the Declining Volatility of Long-Term Inflation Expectations,” *Journal of Economic Dynamics and Control* 35, 981-999.
- Clark, Todd E., and Michael W. McCracken (2011a), “Nested Forecast Model Comparisons: A New Approach to Testing Equal Accuracy,” manuscript, Federal Reserve Bank of St. Louis, January.

- Clark, Todd E., and Michael W. McCracken (2011b), “Testing for Unconditional Predictive Ability,” in *Oxford Handbook of Economic Forecasting*, Michael P. Clements and David F. Hendry, eds., Oxford: Oxford University Press.
- Clark, Todd E., and Michael W. McCracken (2011c), “Advances in Forecast Evaluation,” manuscript, Federal Reserve Bank of Cleveland.
- Clark, Todd E. and Michael W. McCracken (2010), “Averaging Forecasts from VARs with Uncertain Instabilities,” *Journal of Applied Econometrics* 25, 5-29.
- Clark, Todd E. and Michael W. McCracken (2008), “Forecasting with Small Macroeconomic VARs in the Presence of Instability,” in *Forecasting in the Presence of Structural Breaks and Model Uncertainty*, Rapach D.E., Wohar M.E. (eds), Bingley, UK: Emerald Publishing.
- Cogley, Timothy (2002), “A Simple Adaptive Measure of Core Inflation,” *Journal of Money, Credit, and Banking* 34, 94-113.
- Cogley, Timothy, Sergei Morozov, and Thomas J. Sargent (2005), “Bayesian Fan Charts for U.K. Inflation: Forecasting and Sources of Uncertainty in an Evolving Monetary System,” *Journal of Economic Dynamics and Control* 29, 1893-1925.
- Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent (2010), “Inflation-Gap Persistence in the US,” *American Economic Journal: Macroeconomics* 2, 43-69.
- Cogley, Timothy, and Thomas J. Sargent (2005), “Drifts and Volatilities: Monetary Policies and Outcomes in the Post-World War II U.S.,” *Review of Economic Dynamics* 8, 262-302.
- Cogley, Timothy, and Argia M. Sbordone (2008), “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review* 98, 2101-2126.
- Diebold, Francis X., and Roberto S. Mariano (1995), “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics* 13, 253-263.

- Dotsey, Michael, Shigeru Fujita, and Tom Stark (2011), "Do Phillips Curves Conditionally Help to Forecast Inflation?" Working Paper No. 11-40, Federal Reserve Bank of Philadelphia.
- Durbin, J., and S.J. Koopman (2002), "A Simple and Efficient Simulation Smoother for State Space Time Series Analysis," *Biometrika* 89, 603-615.
- Faust, Jon, and Jonathan Wright (2011), "Forecasting Inflation," manuscript, Johns Hopkins University.
- Geweke, John (1999), "Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication," *Econometric Reviews* 18, 1-73.
- Geweke, John, and Gianni Amisano (2010), "Comparing and Evaluating Bayesian Predictive Distributions of Asset Returns," *International Journal of Forecasting* 26, 216-230.
- Geweke, John, and Charles Whiteman (2006), "Bayesian Forecasting," *Handbook of Economic Forecasting*, G. Elliott, C. Granger, and A. Timmermann (Eds.), North Holland: Amsterdam.
- Gordon, Robert J. (1998), "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU," *Brookings Papers on Economic Activity* (No. 2), 297-333.
- Harvey, David I., Stephen J. Leybourne, and Paul Newbold (1997), "Testing the Equality of Prediction Mean Squared Errors," *International Journal of Forecasting* 13, 281-91.
- Ireland, Peter N. (2007), "Changes in the Federal Reserve's Inflation Target: Causes and Consequences," *Journal of Money, Credit, and Banking* 39, 1851-1882.
- Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi, (1994), "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics* 12, 371-418.
- Jore, Anne Sofie, James Mitchell, and Shaun P. Vahey (2010), "Combining Forecast Densities from VARs with Uncertain Instabilities," *Journal of Applied Econometrics* 25, 621-634.

- Kadiyala, K. Rao, and Sune Karlsson (1997), "Numerical Methods for Estimation and Inference in Bayesian VAR Models," *Journal of Applied Econometrics* 12, 99-132.
- Kascha, Christian, and Francesco Ravazzolo (2010), "Combining Inflation Density Forecasts," *Journal of Forecasting* 29, 231-250.
- Kiley, Michael T. (2008), "Monetary Policy Actions and Long-Run Inflation Expectations," Federal Reserve Board FEDS Working Paper 2008-03.
- Koop, Gary, and Dimitris Korobilis (2010), "Bayesian Multivariate Time Series Methods for Empirical Macroeconomics," *Foundations and Trends in Econometrics* 3, 267-358.
- Koop, Gary, and Simon M. Potter (2008), "Time-Varying VARs with Inequality Restrictions," *Journal of Applied Econometrics*, forthcoming
- Kozicki, Sharon, and Peter A. Tinsley (2006), "Survey-Based Estimates of the Term Structure of Expected U.S. Inflation," Bank of Canada Working Paper 2006-46.
- Kozicki, Sharon, and Peter A. Tinsley (2002), "Alternative Sources of the Lag Dynamics of Inflation," in *Price Adjustment and Monetary Policy*, Bank of Canada Conference Proceedings, 3-47.
- Kozicki, Sharon, and Peter A. Tinsley (2001a), "Shifting Endpoints in the Term Structure of Interest Rates," *Journal of Monetary Economics* 47, 613-652.
- Kozicki, Sharon, and Peter A. Tinsley (2001b), "Term Structure Views of Monetary Policy Under Alternative Models of Agent Expectations," *Journal of Economic Dynamics and Control* 25, 149-184.
- Kozicki, Sharon, and P.A. Tinsley (1998), "Moving Endpoints and the Internal Consistency of Agents' Ex Ante Forecasts," *Computational Economics* 11, 21-40.
- Macroeconomic Advisers (2007), "The Gravitational Pull of Inflation Expectations," *Macro Focus*, January 18.
- Mazzi, Gian Luigi, James Mitchell, and Gaetana Montana (2010), "Density Nowcasts and Model Combination: Nowcasting Euro-Area GDP Growth Over the 2008-09 Recession,"

manuscript, NIESR.

Mertens, Elmar (2011), “Measuring the Level and Uncertainty of Trend Inflation,” FEDS Working Paper No. 2011-42, Federal Reserve Board of Governors.

Morley, James, and Jeremy M. Piger (2010), “The Asymmetric Business Cycle,” *Review of Economics and Statistics*, forthcoming.

Piger, Jeremy M., and Robert H. Rasche (2008), “Inflation: Do Expectations Trump the Gap?” *International Journal of Central Banking* 4, 85-116.

Stock, James H., and Mark W. Watson (2010), “Modeling Inflation after the Crisis,” Federal Reserve Bank of Kansas City, Jackson Hole Symposium.

Stock, James H., and Mark W. Watson (2007), “Has U.S. Inflation Become Harder to Forecast?” *Journal of Money, Credit, and Banking* 39, 3-33.

Stock, James H., and Mark W. Watson (2003), “Forecasting Output and Inflation: The Role of Asset Prices,” *Journal of Economic Literature* 41, 788-829.

Villani, Mattias (2009), “Steady-State Priors for Vector Autoregressions,” *Journal of Applied Econometrics* 24, 630-650.

Watson, Mark W. (1986), “Univariate Detrending Methods with Stochastic Trends,” *Journal of Monetary Economics* 18, 49-75.

Whelan, Karl (2002), “A Guide to U.S. Chain Aggregated NIPA Data,” *Review of Income and Wealth* 48, 217-233.

Williams, John C. (2009), “The Risk of Deflation,” *FRBSF Economic Letter*, March 27.

Wright, Jonathan (2011), “Evaluating Real-Time VAR Forecasts with an Informative democratic Prior,” manuscript, Johns Hopkins University.

Table 1. Estimates of Model Parameters, 1960-2011:Q2 Sample
Core PCE Models with Constant Volatility
(Posterior means (standard deviations))

model	b_1	b_2	σ_v^2	σ_n^2	other
constant trend	0.727 (0.069)	0.221 (0.069)	0.728 (0.073)	NA	$\psi = 2.991 (1.203)$
π_{t-1} trend	-0.265 (0.072)	-0.061 (0.071)	0.755 (0.075)	NA	NA
PTR trend	0.613 (0.069)	0.215 (0.069)	0.688 (0.068)	NA	NA
local level	NA	NA	0.214 (0.040)	0.481 (0.070)	NA
random walk trend	0.630 (0.101)	0.192 (0.089)	0.604 (0.083)	0.115 (0.043)	NA
TVP	NA	NA	0.647 (0.068)	NA	$\text{tr}(Q) = 0.006$

Notes: The coefficients b_1 and b_2 refer to the slope coefficients of an AR(2) process. σ_v^2 refers to the variance of the error term in the AR(2) model or, in the case of the local level model, the noise component. In models with random walk trends, σ_n^2 denotes the variance of innovations to trend. ψ denotes the trend in the AR(2) with a constant trend. Q is the variance-covariance matrix of innovations to the vector of AR(2) coefficients following a random walk process in the AR model with TVP.

Table 2. Estimates of Model Parameters, 1960-2011:Q2 Sample
Core PCE Models with Stochastic Volatility
(Posterior means (standard deviations))

model	b_1	b_2	ϕ or ϕ_v	ϕ_n or σ_n^2	other
constant trend	0.683 (0.070)	0.264 (0.071)	0.083 (0.040)		$\psi = 2.317 (0.979)$
π_{t-1} trend	-0.338 (0.076)	-0.125 (0.072)	0.069 (0.037)		NA
PTR trend	0.582 (0.073)	0.241 (0.072)	0.073 (0.034)		NA
local level	NA	NA	0.031 (0.021)	0.124 (0.088)	NA
random walk trend	0.527 (0.119)	0.198 (0.099)	0.087 (0.059)	0.109 (0.039)	NA
TVP	NA	NA	0.057 (0.031)		$\text{tr}(Q) = 0.005$

Notes: The coefficients b_1 and b_2 refer to the slope coefficients of an AR(2) process. In all models except the local level-SV specification, ϕ refers to the variance of innovations to the volatility of the AR(2) equation error. In the local level model, ϕ_v is the variance of innovations to the volatility of the noise component, and ϕ_n is the variance of innovations to the volatility of the noise component. In the random walk trend model, σ_n^2 denotes the variance of innovations to trend. ψ denotes the trend in the AR(2) with a constant trend. Q is the variance-covariance matrix of innovations to the vector of AR(2) coefficients following a random walk process in the AR model with TVP.

Table 3. Full-Sample Log Predictive Likelihoods and Implied Model Probabilities

model	Core PCE		GDP price index	
	PL	model probability	PL	model probability
constant trend	-178.295	0.00000	-211.316	0.00000
π_{t-1} trend	-181.428	0.00000	-210.105	0.00000
PTR trend	-172.960	0.00000	-206.904	0.00000
local level	-181.104	0.00000	-214.505	0.00000
random walk trend	-179.499	0.00000	-210.385	0.00000
TVP	-176.980	0.00000	-209.390	0.00000
constant trend-SV	-165.252	0.00981	-192.069	0.09368
π_{t-1} trend-SV	-165.465	0.00792	-191.336	0.19491
PTR trend-SV	-161.884	0.28464	-191.246	0.21334
local level-SV	-161.191	0.56919	-192.901	0.04075
random walk trend-SV	-163.933	0.03670	-194.733	0.00653
TVP-SV	-163.017	0.09173	-190.498	0.45078
π_{t-1} trend with unemp. gap, SV	-166.816	NA	-194.173	NA
PTR trend trend with unemp. gap, SV	-161.470	NA	-196.700	NA

Note: The likelihoods are computed by summing 1-step ahead log predictive likelihoods for the period 1975:Q1-2011:Q2.

Table 4. Mean Forecast Errors, Core PCE Inflation, 1985:Q1-2011:Q2

model	$h=1$	$h=2$	$h=3$	$h=4$	$h=8$	$h=12$	$h=16$
constant trend	-0.072	-0.148 *	-0.214 **	-0.277 ***	-0.503 ***	-0.698 ***	-0.870 ***
π_{t-1} trend	-0.016	-0.048	-0.071	-0.098	-0.190	-0.287	-0.387
PTR trend	-0.069	-0.141 *	-0.197 **	-0.249 ***	-0.419 **	-0.552 **	-0.703 *
local level	-0.015	-0.049	-0.074	-0.097	-0.186	-0.286	-0.391
random walk trend	-0.060	-0.126	-0.179 **	-0.230 **	-0.398 **	-0.548 *	-0.694 *
TVP	-0.067	-0.129	-0.181 **	-0.227 **	-0.374 *	-0.504	-0.637
constant trend-SV	-0.031	-0.075	-0.110	-0.144	-0.264	-0.382	-0.508
π_{t-1} trend-SV	-0.018	-0.050	-0.072	-0.096	-0.190	-0.291	-0.393
PTR trend-SV	-0.064	-0.133 *	-0.190 **	-0.237 ***	-0.403 **	-0.536 **	-0.688 *
local level-SV	-0.038	-0.071	-0.091	-0.110	-0.196	-0.281	-0.389
random walk trend-SV	-0.040	-0.093	-0.133	-0.169 *	-0.291	-0.406	-0.522
TVP-SV	-0.050	-0.102	-0.143	-0.177 *	-0.300	-0.413	-0.537
PL-weighted average	-0.038	-0.085	-0.115	-0.145	-0.254	-0.360	-0.482
simple average, all	-0.045	-0.097	-0.138	-0.177 **	-0.308 *	-0.429	-0.554
simple average, just SV	-0.039	-0.086	-0.121	-0.155 *	-0.274	-0.385	-0.507

Notes: The forecast errors are defined as actual inflation rates less forecasts. Mean errors that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p -values are based on t -statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997). The PL-weighted average forecast is computed in pseudo-real time, using weights at each forecast origin based on the prior 10 years of log predictive likelihood values, but not future information.

Table 5. Forecast RMSEs, Core PCE Inflation, 1985:Q1-2011:Q2

*(RMSE levels for local level-SV,
ratio of RMSE to local level-SV RMSE for all other models)*

model	$h=1$	$h=2$	$h=3$	$h=4$	$h=8$	$h=12$	$h=16$
local level-SV	0.578	0.625	0.652	0.644	0.781	0.918	0.994
constant trend	1.043	1.079 **	1.101 ***	1.139 ***	1.127	1.175	1.189
π_{t-1} trend	1.043	1.075 **	1.083 **	1.088 **	1.021	1.046 *	1.011
PTR trend	0.998	1.003	0.990	1.013	0.992	0.983	1.062
local level	1.032	1.061 *	1.067 **	1.075 *	1.011	1.036	1.011
random walk trend	1.032	1.059 **	1.062 ***	1.087 ***	1.061	1.093	1.126
TVP	1.018	1.048 **	1.048 **	1.071 ***	1.053	1.082	1.119
constant trend-SV	1.037	1.063 **	1.064 *	1.070 *	0.993	1.011	0.979
π_{t-1} trend-SV	1.036	1.056 **	1.070 **	1.066 *	1.013	1.028	1.018
PTR trend-SV	1.001	1.005	0.996	1.012	0.980	0.975	1.055
random walk trend-SV	1.037	1.072 **	1.074 ***	1.084 ***	1.015	1.030	1.055
TVP-SV	1.022	1.056 **	1.056 ***	1.070 ***	1.020	1.041	1.081
PL-weighted average	1.010	1.019	1.022	1.019	0.982	0.994	0.993
simple average, all	1.019	1.035	1.037 *	1.045 *	0.989	1.003	1.021
simple average, just SV	1.014	1.034	1.030	1.035	0.986	0.990	1.005

Notes: To provide a rough gauge of whether the RMSE ratios are significantly different from 1, we use the Diebold-Mariano t -statistic for equal MSE, applied to the forecast of each model relative to the benchmark of the local level-SV forecast. Differences in accuracy that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p -values are based on t -statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997). The PL-weighted average forecast is computed in pseudo-real time, using weights at each forecast origin based on the prior 10 years of log predictive likelihood values, but not future information.

Table 6. Average Log Scores, Core PCE Inflation, 1985:Q1-2011:Q2
*(average score for local level-SV,
difference relative to local level-SV for all other models)*

model	$h=1$	$h=2$	$h=3$	$h=4$	$h=8$	$h=12$	$h=16$
local level-SV	-0.896	-0.998	-1.029	-1.067	-1.282	-1.433	-1.532
constant trend	-0.176 ***	-0.256 ***	-0.349 ***	-0.402 ***	-0.428 ***	-0.427 ***	-0.415 *
π_{t-1} trend	-0.174 ***	-0.263 ***	-0.364 ***	-0.424 ***	-0.492 ***	-0.543 ***	-0.573 ***
PTR trend	-0.131 **	-0.171 ***	-0.227 ***	-0.251 ***	-0.172	-0.098	-0.071
local level	-0.212 ***	-0.302 ***	-0.404 ***	-0.458 ***	-0.530 ***	-0.574 ***	-0.603 ***
random walk trend	-0.181 ***	-0.260 ***	-0.345 ***	-0.395 ***	-0.397 ***	-0.394 ***	-0.389 **
TVP	-0.154 ***	-0.211 ***	-0.282 ***	-0.322 ***	-0.300 ***	-0.283 ***	-0.268
constant trend-SV	-0.040	-0.083 **	-0.091 ***	-0.128 ***	-0.086	-0.093	-0.088
π_{t-1} trend-SV	-0.033	-0.086 **	-0.081 **	-0.105 ***	-0.088	-0.118	-0.127
PTR trend-SV	-0.008	-0.029	-0.053	-0.063	-0.029	-0.037	-0.083
random walk trend-SV	-0.033	-0.067 **	-0.126 ***	-0.136 ***	-0.130 *	-0.148 *	-0.158
TVP-SV	-0.023	-0.053 **	-0.073 ***	-0.084 ***	-0.046	-0.061	-0.069
PL-weighted average	-0.039 **	-0.047 **	-0.046 **	-0.061 **	-0.052	-0.057	-0.064 *
simple average, all	-0.058	-0.104 ***	-0.159 ***	-0.191 ***	-0.186 **	-0.183 **	-0.184 *
simple average, just SV	-0.011	-0.034	-0.059 ***	-0.076 **	-0.045	-0.053	-0.054

Notes: The average log scores are computed as averages of log predictive likelihoods (a larger number corresponds to a better model). To provide a rough gauge of the statistical significance of differences in average log scores, we use the Amisano-Giacomini t -test of equal means, applied to the log score for each model relative to the benchmark of the local level-SV forecast. Differences in average scores that are statistically different from zero are denoted by one, two, or three asterisks, corresponding to significance levels of 10%, 5%, and 1%, respectively. The underlying p -values are based on t -statistics computed with a serial correlation-robust variance, using a rectangular kernel, $h - 1$ lags, and the small-sample adjustment of Harvey, Leybourne, and Newbold (1997). The PL-weighted average forecast is computed in pseudo-real time, using weights at each forecast origin based on the prior 10 years of log predictive likelihood values, but not future information.

Table 7. Forecast RMSEs, GDP Inflation, 1985:Q1-2011:Q2
*(RMSE levels for local level-SV,
ratio of RMSE to local level-SV RMSE for all other models)*

model	$h=1$	$h=2$	$h=3$	$h=4$	$h=8$	$h=12$	$h=16$
local level-SV	0.778	0.832	0.863	0.925	1.107	1.254	1.285
constant trend	1.034	1.042	1.035	1.052	1.083	1.116	1.177
π_{t-1} trend	1.013	0.997	1.010	1.007	1.019 *	1.007	1.005
PTR trend	1.007	1.007	0.993	0.986	0.945	0.928	1.012
local level	1.017	1.019	1.014	1.014	1.018 **	1.017 **	1.021 *
random walk trend	1.024	1.029	1.009	1.000	0.974	0.961	0.982
TVP	1.024	1.036	1.022	1.018	0.979	0.969	1.001
constant trend-SV	1.013	1.011	0.982	0.978	0.941	0.927	0.929
π_{t-1} trend-SV	1.001	0.990	1.007	1.002	1.014 *	1.009	1.011
PTR trend-SV	1.005	1.002	0.984	0.973	0.939	0.917	0.997
random walk trend-SV	1.028	1.028	0.997	0.995	0.960 *	0.945	0.951
TVP-SV	1.017	1.019	0.992	0.988	0.952 *	0.940	0.950
PL-weighted average	1.008	1.005	0.991	0.972	0.936	0.912	0.917
simple average, all	1.007	1.004	0.989	0.979	0.948	0.939	0.955
simple average, just SV	1.004	0.998	0.981	0.975	0.950	0.932	0.939

Note: See the notes to Table 5.

Table 8. Average Log Scores, GDP Inflation, 1985:Q1-2011:Q2
*(average score for local level-SV,
difference relative to local level-SV for all other models)*

model	$h=1$	$h=2$	$h=3$	$h=4$	$h=8$	$h=12$	$h=16$
local level-SV	-1.141	-1.236	-1.299	-1.380	-1.581	-1.718	-1.778
constant trend	-0.186 ***	-0.245 ***	-0.296 ***	-0.310 ***	-0.338 ***	-0.332 *	-0.353 **
π_{t-1} trend	-0.175 ***	-0.230 ***	-0.260 ***	-0.285 ***	-0.353 ***	-0.381 **	-0.447 ***
PTR trend	-0.154 ***	-0.189 ***	-0.216 ***	-0.203 **	-0.135	-0.073	-0.072
local level	-0.214 ***	-0.275 ***	-0.332 ***	-0.355 ***	-0.426 ***	-0.460 ***	-0.521 ***
random walk trend	-0.185 ***	-0.251 ***	-0.287 ***	-0.288 ***	-0.278 ***	-0.256	-0.263 *
TVP	-0.176 ***	-0.222 ***	-0.268 ***	-0.265 ***	-0.238 **	-0.199	-0.205
constant trend-SV	-0.002	0.011	0.019	0.015	0.016	0.015	0.034
π_{t-1} trend-SV	0.019	0.030	0.032	0.039	0.029	0.012	-0.010
PTR trend-SV	0.013	0.012	0.017	0.021	0.016	0.006	-0.046
random walk trend-SV	-0.026	-0.029	-0.010	-0.032	-0.030	-0.033	-0.044
TVP-SV	0.005	-0.005	0.025	0.024	0.035	0.043	0.030
PL-weighted average	-0.010	0.002	0.014	0.004	0.026	0.028	0.024
simple average, all	-0.049	-0.074 *	-0.097 **	-0.095 **	-0.104 **	-0.096	-0.116
simple average, just SV	0.003	0.011	0.023	0.023	0.031	0.018	0.014

Note: See the notes to Table 6.




Figure 1: Posterior estimates of time-varying volatility, PTR trend model, core PCE inflation





Figure 2: Cumulative log predictive likelihoods, relative to local level-SV model, core PCE inflation





Figure 3: Cumulative log predictive likelihoods, relative to local level-SV model, GDP deflator inflation





Figure 4: Model weights based on 10-year rolling predictive likelihoods, core PCE inflation








Figure 5: Posterior estimates of trend inflation, constant volatility models, core PCE inflation








Figure 6: Posterior estimates of trend inflation, stochastic volatility models, core PCE inflation





Figure 7: Model average estimates of trend inflation, core PCE inflation








Figure 8: Posterior estimates of trend inflation, stochastic volatility models, GDP inflation





Figure 9: Model average estimates of trend inflation, GDP inflation





Figure 10: Posterior estimates of deflation probabilities, core PCE inflation




Figure 11: Deflation probabilities from models with unemployment gap, core PCE inflation