working paper

15 20

A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations

Joshua C.C. Chan, Todd E. Clark, and Gary Koop

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 15-20 October 2015

A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations

Joshua C.C. Chan, Todd E. Clark, and Gary Koop

A knowledge of the level of trend inflation is key to many current policy decisions, and several methods of estimating trend inflation exist. This paper adds to the growing literature which uses survey-based long-run forecasts of inflation to estimate trend inflation. We develop a bivariate model of inflation and long-run forecasts of inflation which allows for the estimation of the link between trend inflation and the long-run forecast. Thus, our model allows for the possibilities that long-run forecasts taken from surveys can be equated with trend inflation, that the two are completely unrelated, or anything in between. By including stochastic volatility and time-variation in coefficients, it extends existing methods in empirically important ways. We use our model with a variety of inflation measures and survey-based forecasts. We find that long-run forecasts can provide substantial help in refining estimates of trend inflation over popular alternatives. But simply equating trend inflation with the long-run forecasts is not appropriate.

Keywords: trend inflation, inflation expectations, state space model, stochastic volatility.

JEL Classification: C11, C32, E31.

Suggested citation: Chan, Joshua C.C., Todd E. Clark, and Gary Koop, 2015. “A New Model of Inflation, Trend Inflation, and Long-Run Inflation Expectations,” Federal Reserve Bank of Cleveland, working paper no. 15-20.

Joshua C.C. Chan is at Australian National University, Todd E. Clark is at the Federal Reserve Bank of Cleveland (todd.clark@clev.frb.org), and Gary Koop is at the University of Strathclyde. The authors gratefully acknowledge helpful comments from Edward Knotek, Elmar Mertens, Mike West, colleagues at the Federal Reserve Bank of Cleveland, and participants at the 2015 National Bank of Poland workshop on forecasting. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Cleveland, Federal Reserve System, or any of its staff. Joshua Chan would like to acknowledge financial support by the Australian Research Council via a Discovery Early Career Researcher Award (DE150100795).

1 Introduction

As is evident in public commentary (see, e.g., Bernanke 2007 and Mishkin 2007), central bankers and other policymakers pay considerable attention to measures of long-run inflation expectations. These expectations are viewed as shedding light on the credibility of monetary policy. Monetary policy tools work differently if long-run inflation expectations are firmly anchored than if they are not. In general, monetary policy is thought to be most effective when long-run inflation expectations are stable.

These considerations have contributed to the development of a large literature on the measurement of long-run inflation expectations. One simple approach is to rely on direct estimates of inflation expectations from surveys of professionals or consumers.[1] For example, Federal Reserve commentary such as Mishkin (2007) includes long-run expectations based on the Survey of Professional Forecasters' (SPF) projection of average inflation 1 to 10 years ahead.

Other approaches rely on econometric estimates of trend inflation; under some assumptions, trend inflation should correspond to long-run inflation expectations. A large literature uses econometric methods to estimate inflation trends and forecast inflation (see, among many others, Stock and Watson, 2007, Chan, Koop and Potter, 2013, and Clark and Doh, 2014).[2] A smaller strand of the literature combines econometric models of trend with the information in surveys (see, among others, Kozicki and Tinsley, 2012, Wright, 2013 and Nason and Smith, 2014).[3]

In recent years, some countries have experienced extended periods of inflation running below survey-based estimates of long-run inflation expectations. For example, Fuhrer, Olivei, and Tootell (2012) show that actual inflation in Japan consistently ran below (survey-based) long-run inflation expectations in their sample, from the early 1990s to 2010. More recently, in the United States, for each year between 2008 and 2014, inflation in the core PCE price index has run below the SPF long-run forecast of 2 percent (which coincides with the Federal Reserve's official goal for inflation).[4] Even though survey-based inflation expectations have been stable, actual inflation has been low enough for long enough to pull some common econometric estimates of trend inflation well below 2 percent (see, e.g., Bednar and Clark 2014). These experiences raise the question of whether it is possible for survey-based inflation expectations to become disconnected from actual inflation. Such a disconnect would make such expectations less useful for

[1] Direct estimates of inflation expectations can also be obtained based on the relationship between real and nominal bonds. However, estimates of break-even inflation calculated using these are usually available only for a short time span. And there are reasons to expect that break-even inflation might reflect factors other than just long run inflation expectations (e.g. if the risk premium is time-varying). Faust and Wright (2013) find it too volatile to be a sensible forecast for long run expected inflation. For these reasons, we do not use break-even inflation data in this paper.

[2] The reader is referred to Faust and Wright (2013) for a recent survey on inflation forecasting, including a discussion of inflation surveys and methods for estimating trend inflation.

[3] Some DSGE models — developed in Del Negro and Schorfheide (2013) and references therein — treat the inflation target of the central bank as a random walk process and include survey measures of long-run inflation expectations as indicators of the target in model estimation.

[4] This statement is based on Q4/Q4 inflation rates for each year. The statement also applies to headline inflation, except that headline inflation rose above two percent for one year, 2011.

gauging the credibility of monetary policy and for forecasting inflation.

In this paper we develop a new model to examine the relationship between inflation, inflation expectations and trend inflation. We use this model to assess whether survey-based inflation expectations can become disconnected from actual inflation and, even if they do so, whether they can improve estimates of trend inflation. We build on papers such as Kozicki and Tinsley (2012) by using models which are more flexible in empirically important directions, extending recent work with unobserved components models with stochastic volatility (UCSV) such as Stock and Watson (2007, 2015), Chan, Koop and Potter (2013), Clark and Doh (2014), Garnier, Mertens, and Nelson (2015), and Mertens (2015). Papers such as Kozicki and Tinsley (2012) equate long run forecasts with trend inflation. Similarly, econometric estimates of trend inflation are sometimes calibrated to be the same as surveys. Our model breaks such links between trend inflation and long run inflation forecasts. Instead it allows us to estimate the relationship to investigate whether equating trend inflation with inflation expectations based on surveys is a sensible thing to do. Furthermore, it does so in a time-varying manner so that, e.g., trend inflation can be equal to the forecasts provided in the surveys at some points in time, but at other points in time forecasts can provide biased estimates of trend inflation. Another point of departure from the existing literature is that we only use survey data on long run inflation forecasts, allowing us to avoid the use of a subsidiary (possibly mis-specified) model linking short-run forecasts to long run inflation expectations.

An empirical application involving several measures of US inflation and long-run forecasts from two different sources shows the usefulness of our approach. We present evidence that extensions over simpler approaches such as the addition of stochastic volatility and time-varying coefficients are important in practice. Survey-based measures of inflation expectations are found to be useful for estimating trend inflation, producing smoother and more sensible estimates than the UCSV model. However, we also present evidence that the survey-based measures should not simply be equated with trend inflation as the relationship between the two is more complicated and time-varying. We conclude with an examination of out-of-sample forecasting, which shows point and density forecasts from our model to be at least as good as those from other models that have been found successful in the inflation forecasting literature.

2 Econometric Modelling of Trend Inflation

An unobserved components framework is commonly-used to model inflation, π_t, as being composed of trend (or underlying) inflation, π_t^*, and a deviation from trend, the inflation gap, c_t:

$$\pi_t = \pi_t^* + c_t. \tag{1}$$

The two components of inflation are identified by making assumptions (e.g. that trend inflation follows a random walk) that imply

$$\lim_{j\to\infty} E_t\left[\pi_{t+j}\right] = E_t[\pi^*_{t+j}] \; = \pi_t^* \tag{2}$$

and

$$\lim_{j \to \infty} E_t [c_{t+j}] = 0, \tag{3}$$

where $E_t[\bullet]$ are expectations at time t.

There are many possible econometric models consistent with this simple decomposition (which can be seen as a generalization of the Beveridge-Nelson decomposition) and we will argue for a particular modeling framework soon. But the basic justification for using surveys of long run forecasts can be clearly seen from (2). Those surveyed at time t about what inflation will be in period $t + j$ can be expected to be reporting $E_t[\pi_{t+j}]$. Thus, using (2), forecasts of long-run inflation should also provide estimates of $E_t[\pi^*_{t+j}]$ for large j and, given the random walk assumption, also for trend inflation, π^*_t. There are several ways that this relationship plus data on long-run forecasts made at time t (z_t) can be used to produce estimates of current trend inflation, with Kozicki and Tinsley (2012) being an influential recent approach.

However, there are reasons to be cautious about simply equating long run forecasts from surveys with inflation trends. For instance, surveys may produce forecasts that are biased, at least at some points in time. Surveys might also contain some noise, due to factors such as changes in participants from one survey date to another. Accordingly, we desire an econometric specification that allows us to estimate the relationship between π_t and z_t rather than imposing a particular form. In our model, a finding that long run forecasts taken from surveys can be equated with trend inflation is possible, but not assumed a priori.

Earlier work also suggests many other desirable features we want our econometric model to have. First, the inflation gap, c_t, should be stationary but may exhibit persistence. For instance, the Fed may tolerate deviations of inflation from a trend or target for a certain period of time, provided such deviations are temporary. Furthermore, the Fed's toleration for such deviations may change over time. For instance, Chan, Koop and Potter (2013) discuss how the high inflation in the 1970s may have been partly due to the combination of a large inflation gap (with only a small increase in trend inflation) with a Fed tolerant of a high degree of inflation gap persistence. When Paul Volcker subsequently became the Fed governor, this tolerance decreased and inflation gap persistence dropped. We want our model to be able to accommodate such shifts in persistence.

Second, Faust and Wright (2013) find improvements in forecast performance by using the inflation gap (as opposed to inflation itself) as a dependent variable and modeling the inflation gap as deviations of actual inflation from a slowly evolving trend. Following this recommendation, our econometric specification also has this property.

Third, a large number of papers such as Stock and Watson (2007) have found the importance of allowing for stochastic volatility, not only in the inflation equation but also in the state equations which describe the evolution of trend inflation.

Finally, a general theme of many papers on inflation modeling, including Faust and Wright (2013) and Stella and Stock (2013), is time-varying predictability. Accordingly, we want a time-varying parameter (TVP) model where coefficients can change.

All of these features are built into the following extremely flexible model which should be able to accommodate any relevant empirical properties of the data on inflation (π_t)

and the survey-based inflation expectation (z_t):

$$
\begin{aligned}
\pi_t - \pi_t^* &= b_t(\pi_{t-1} - \pi_{t-1}^*) + v_t, & (4)\\
z_t &= d_{0t} + d_{1t}\pi_t^* + \varepsilon_{z,t} + \psi\varepsilon_{z,t-1}, & (5)\\
\pi_t^* &= \pi_{t-1}^* + n_t, & (6)\\
b_t &= b_{t-1} + \varepsilon_{b,t},\ \varepsilon_{b,t} \sim TN(0, \sigma_b^2), & (7)\\
d_{it} - \mu_{di} &= \rho_{di}\left(d_{i,t-1} - \mu_{di}\right) + \varepsilon_{di,t},\ \varepsilon_{di,t} \sim\ N(0, \sigma_{di}^2),\ i = 0, 1, & (8)\\
v_t &= \lambda_{v,t}^{0.5}\varepsilon_{v,t},\ \varepsilon_{v,t} \sim N(0, 1), & (9)\\
n_t &= \lambda_{n,t}^{0.5}\varepsilon_{n,t},\ \varepsilon_{n,t} \sim N(0, 1), & (10)\\
\log(\lambda_{i,t}) &= \log(\lambda_{i,t-1}) + \nu_{i,t},\ \ \nu_{i,t} \sim\ N(0, \phi_i),\ i = v, n. & (11)
\end{aligned}
$$

All of the errors defined above are independent over time and with each other. $TN_{(0,1)}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 truncated so as to ensure $0 < b_t < 1$ at every point in time. Note that by allowing b_t to be time-varying we can find changes in the degree of persistence in the inflation gap. And truncating the errors in (7) to an appropriate interval allows us to ensure that the inflation gap is stationary at every point in time.

Variants of the model described above, excluding z_t, involving only (possibly restricted versions of) (4), (6), (7), (9), (10) and (11) have been used to estimate trend inflation by several authors. For instance, the popular UCSV model of Stock and Watson (2007) is this model with $b_t = 0$, and Chan et al (2013) use this model with bounded trend inflation but without stochastic volatility in $\varepsilon_{n,t}$. We stress that stochastic volatility is often found to be important in models of trend inflation such as these.[5] This feature allows for the possibility that the volatility of trend inflation or deviations of inflation from trend vary over time.

By adding the additional equations (5) and (8) to a conventional unobserved components model such as the one defined by (4), (6), (7), (9), (10) and (11), we can potentially improve estimates of trend inflation. That is, adding the relationship between z_t and π_t^* should provide extra information for estimating trend inflation beyond that provided in a univariate model involving inflation only.

Another important feature of our model is that we allow d_{0t} and d_{1t} to vary over time.[6] These coefficients relate to the question of whether long run inflation forecasts are unbiased estimates of trend inflation. If $d_{0t} = 0$ and $d_{1t} = 1$ they are. If $d_{1t} = 1$ but

[5] For the errors in other equations, preliminary estimates suggest that an assumption of homoskedasticity is reasonable.

[6] Although our model is less restrictive than other studies that relate inflation and survey measures of inflation expectations, our specification can be seen as consistent with the cointegration restrictions imposed in these other studies (e.g., Mertens 2015, Mertens and Nason 2015, and Nason and Smith 2014). These other studies impose stationarity of the difference between actual inflation and survey expectations. Our model is consistent with cointegration of the survey expectation z_t with trend inflation π_t^*: the innovation term of the z_t equation is a stationary MA(1) process. Although the posterior of $d_{0,t}$ and $d_{1,t}$ need not be close to 0 or 1, respectively, our prior puts the initial values of these coefficients at 0 and 1, respectively. So our prior implies cointegration of z_t with trend inflation π_t^* with a slope coefficient of 1. With π_t^* the source of integration in π_t, it follows that we can think of π_t and z_t as cointegrated as well.

$d_{0t} \neq 0$ then long run forecasts are consistently biased upwards or downwards. Under this definition, bias includes either a constant differential between trend inflation and the survey forecast or a failure of the survey to move one-for-one with trend.[7] Thus, investigating restrictions relating to d_{0t} and d_{1t} is of economic interest. To allow for persistence in a long-term inflation forecast that may not be adequately picked up by persistence in trend inflation, we add an MA(1) error term to (5). As we shall see, empirical evidence for the need for this MA error term is weak for some choices of z_t (and there is never any evidence for lag lengths greater than one), but for the Blue Chip forecasts the MA term is empirically important and we include it in our general specification. Since d_{0t} and d_{1t} are time varying, we have the potential to estimate changes in the relationship between long run forecasts and trend inflation. For instance, it is possible that long run forecasts are unbiased estimates of trend inflation at some points in time, but not others. Our model allows for this possibility, but a constant coefficient model would not.

Finally, our model excludes an economic activity indicator from the inflation gap equation (4). We do so in the interest of parsimony, motivated in part by evidence in the forecasting literature (see Faust and Wright 2013 and references therein) of the difficulty of using economic activity variables to improve predictions of inflation. However, we include in the empirical appendix results for a model supplemented to include in the inflation equation an unemployment rate gap with a time-varying coefficient. These results are very similar to those we obtain without economic activity in the model. Our specification with the unemployment gap has precedents in other recent studies, including: Stella and Stock (2013), which generalizes the UCSV formulation of Stock and Watson to relate the inflation gap to an unemployment gap; Jarocinski and Lenza (2015), which considers a model with a factor model of economic activity, for the purpose of estimating the output gap, with a structure for inflation, trend inflation, and inflation expectations that corresponds to a restricted, constant parameter version of our formulation; and Morley, Piger, and Rasche (2015), which considers a bivariate, constant parameter model relating inflation less a random walk trend to an unemployment gap.

We use Bayesian methods to estimate all the unknown parameters of our model, including latent variables such as trend inflation. The Markov Chain Monte Carlo (MCMC) algorithm used for estimation is similar to that used in previous work (e.g. Chan et al, 2015) and, hence, we say no more of it here (see the Technical Appendix for details). For model comparison, we calculate posterior model probabilities using methods that are less familiar and, accordingly, we briefly describe here. Bayesian model comparison is typically done using posterior model probabilities. That is, if the researcher has a set of models, M_r for $r = 1, .., R$, model comparison can be done by calculating $P(M_r|Data)$,

[7] Conceptually, the distinction between the infinite horizon forecast that constitutes trend inflation and the 10-year horizon of the survey expectation could cause $d_{0,t}$ to differ from 0 and $d_{1,t}$ to differ from 1. In practice, though, for professional forecasters, it seems likely that the 10-year ahead survey forecast is equivalent to an infinite horizon forecast. For example, since the Federal Reserve established its longer-run inflation objective of 2 percent, the 10 year-ahead forecast of PCE inflation from the Survey of Professional Forecasters has been anchored at 2 percent. Moreover, in a cross-country analysis, Mehrotra and Yetman (2014) find that survey forecasts at just a 24-month ahead horizon tend to cluster around central bank inflation targets.

the posterior model probability (where $P(\bullet|Data)$ is our general notation for a posterior probability, i.e. the probability of $\bullet$ conditional on the data). For model comparison involving models nested within an unrestricted specification, $P(M_r|Data)$ can often be obtained in a particularly simple way through the use of the Savage-Dickey Density Ratio (SDDR, see, e.g., Verdinelli and Wasserman, 1995). The SDDR provides a direct estimate of the Bayes factor comparing an unrestricted to restricted model (e.g. comparing our full model to a variant which restricts the MA coefficient in (5) to be zero). If equal prior weight is attached to each model, then the Bayes factor is simply the ratio of posterior model probabilities. Recently, methods for calculating the SDDR in a time-varying manner in state space models (such as the one used in this paper), have been developed by Koop, Leon-Gonzalez and Strachan (2010). In this paper, we use such methods to calculate posterior model probabilities in a time varying fashion. For instance, we can calculate $P(d_{0t} = 0, d_{1t} = 1|Data)$ for $t = 1, .., T$ to see if long run inflation forecasts provide unbiased estimates of trend inflation at some points in time but not others. The Technical Appendix provides additional details on our Bayesian econometric methods.

3 Data

Policymakers are interested in a range of different measures of inflation, and the research literature considers a range of measures. Accordingly, we provide results for a number of combinations of different measures of inflation and inflation expectations. In the interest of brevity, in the text we focus on three combinations. In the Empirical Appendix we provide results for three additional combinations; these results are very similar to those we report in the text.

In total, we provide results for four different measures of quarterly inflation (π_t in the model): i) inflation based on the consumer price index (CPI inflation), ii) inflation based on the consumer price index excluding food and energy (core CPI inflation), iii) inflation based on the price index for personal consumption expenditures (PCE inflation) and iv) inflation based on the GDP deflator (GDP deflator inflation). Inflation rates are computed as annualized log percent changes ($\pi_t = 400 \ln (P_t/P_{t-1})$ where P_t is a price index).

In our text results, we focus on the (headline) CPI and PCE measures. The CPI has the advantage of being widely familiar to the public, and for much of our sample, the available inflation expectations data refer to it. However, changes over time in the methodology used to construct the CPI — such as the 1983 change in the treatment of housing costs to use rental equivalence — may create structural instabilities, because the historical data are not revised to reflect methodology changes. One reason we also consider PCE inflation is that its historical data has been revised to reflect methodology changes, reducing concerns with instabilities created by methodology changes. Another reason is that the Federal Reserve's preferred inflation measure is PCE inflation; its longer-run inflation objective is stated in terms of PCE inflation.

Reflecting data availability, our results draw on a few different sources of long-run inflation expectations. Surveys of professional forecasters have long included projections of CPI inflation or the GNP/GDP price deflator/price index, but only recently has any survey included PCE inflation.[8] However, in light of the many similarities in CPI and PCE inflation (many of the detailed price indexes used to construct the PCE measure come from the CPI), some policy work and some research measures historical expectations for PCE inflation with expectations for CPI inflation, subject to an adjustment for the difference in their average inflation rates. We follow that practice in this paper.

More specifically, our main source of long-run inflation expectations (z_t in the model) is the Federal Reserve Board of Governor's FRB/US econometric model, which includes inflation expectations as a variable denoted PTR. Defined in CPI terms, the PTR series in the Board's model splices (1) econometric estimates of inflation expectations from Kozicki and Tinsley (2001) early in the sample to (2) 5- to 10-year-ahead survey measures compiled by Richard Hoey to (3) 1- to 10-year ahead expectations from the Survey of Professional Forecasters. Defined in the PCE terms actually used in the FRB/US model, the series uses the same sources, but from 1960 through 2006, the source data are adjusted to a PCE basis by subtracting 50 basis points from the inflation expectations measured in CPI terms. We refer to these long run forecast series for CPI and PCE inflation as PTR-CPI and PTR-PCE, respectively. From 1960 through 2006, PTR-PCE is just PTR-CPI less 50 basis points.

We also use the Blue Chip Consensus as a source of long-run inflation expectations. Blue Chip has been publishing long run (6-10 year) forecasts of CPI inflation and GNP or GDP deflator inflation since 1979 in the latter case and 1983 in the former case. To extend the CPI forecast survey back to 1979, we fill in data for 1979 to 1983 using deflator forecasts from Blue Chip.[9] The forecasts are only published twice a year; we construct quarterly values using interpolation.

In the interest of text brevity, we present results for three combinations of inflation with corresponding inflation expectations: i) CPI inflation plus PTR-CPI long run forecasts, ii) PCE inflation plus PTR-PCE long run forecasts, and iii) CPI inflation plus Blue Chip CPI inflation forecasts. This set addresses robustness to different inflation measures and to different measures of inflation expectations. We provide results for other combinations in the empirical appendix: iv) core CPI inflation plus PTR-CPI long run forecasts, v) core CPI inflation plus Blue Chip CPI inflation forecasts and vi) GDP deflator inflation plus Blue Chip GDP deflator inflation forecasts.

Finally, in results based on the PTR measures of inflation expectations, we estimate the model using data from 1960:Q1 to 2014:Q3. In results based on Blue Chip expectations, the sample period is 1979:Q4 to 2014:Q3.

4 Empirical Results

In this section, we present results for three different data combinations (i.e. an inflation with appropriate long run inflation expectations measure). Empirical results are mostly presented using figures. In each case, the first set of figures plots posterior

[8] The Blue Chip consensus tracks expectations of inflation in both the CPI and GDP price index. The Survey of Professional Forecasters tracks expectations of CPI inflation and, since 2007, PCE inflation.

[9] For the next several years following 1983, Blue Chip's long-run forecasts of CPI and GDP inflation are very similar.

means (along with an interval estimate) of all the latent variables in the model (i.e. $\pi_t^*, b_t, \lambda_{v,t}, \lambda_{n,t}, d_{0t}, d_{1t}$). The figure for π_t^* also plots actual inflation (π_t) along with long-run forecasts taken from the surveys (z_t). The second set of figures relate to the question of whether long-run survey forecasts can be equated with trend inflation or not. In particular, they plot $P(d_{0t} = 0, d_{1t} = 1|Data)$, $P(d_{0t} = 0|Data)$ and $P(d_{1t} = 1|Data)$ for $t = 1, .., T$. Our model allows for these probabilities to vary over time. But it is possible that d_{0t} and/or d_{1t} are constant over time. To shed light on this, we present figures plotting, e.g., $P(d_{0t} = d_{0s}|Data)$ for $t = 1, .., T$ and a given choice of s. This addresses the question: what is the probability a parameter has changed since time s? To present evidence on whether parameters are constant over time, we present figures of this form for d_{0t}, d_{1t} (individually and jointly) and b_t.

For comparison, we also present trend inflation for a more general version of the popular UCSV model of Stock and Watson (2007) which we take as a restricted version of our model which excludes the equations involving z_t, d_{0t} and d_{1t} but allowing b_t to be non-zero, while all other choices (including the prior hyperparameters) are identical to our model. This model, which corresponds to the univariate model of Cogley, Primiceri, and Sargent (2010), is denoted as the UCSV-AR model.[10]

In the empirical appendix, we provide plots of the prior and posterior of the MA coefficient for each data combination. This is used to calculate the SDDR and help decide whether to include an MA process in (5). The priors used in this paper are informative, but not dogmatically so. In models such as ours, involving many unobserved latent variables, use of informative priors is typically necessary.[11] In the empirical appendix, we present results from a prior sensitivity analysis showing our results are fairly robust to changes in our prior.

4.1 Results Using CPI Inflation and PTR-CPI Forecasts

Using CPI inflation and the PTR-CPI long-run inflation expectations, Figures 1 through 4 present estimates of $\pi_t^*, b_t, \lambda_{v,t}, \lambda_{n,t}, d_{0t}$ and d_{1t}. Trend inflation estimates can be seen to be much smoother than actual inflation and track long-run survey-based forecasts fairly well. However, there are some differences between π_t^* and z_t, particularly around 1980. This was a time of high inflation and the professionals were forecasting long run inflation to be somewhat higher than our estimate of trend inflation. A finding that the professionals' forecasts are often slightly above our estimates of trend inflation (particularly around 1980) can also be seen in the estimates of d_{0t} and d_{1t}. Remember that $d_{0t} = 0$ and $d_{1t} = 1$ implies long run forecasts are unbiased estimates of trend inflation. Our estimate of d_{0t} is positive and d_{1t} is above one (and increases to a value well-above one around 1980). These values jointly imply that our trend inflation estimates are slightly below those of the professional throughout the sample and this difference increases around 1980. However, the lower bound of the interval estimate for d_{0t} (d_{1t}) tends to be near zero (one). Estimates

[10] However, for simplicity, we depart from Cogley, Primiceri, and Sargent (2010) by treating the variance of innovations to the AR(1) coefficient as constant instead of time-varying.

[11] Indeed, in the UC-SV model of Stock and Watson (2007), the stochastic volatility equations equivalent to our (11) are assumed to have a common error variance and this common variance is fixed at a specific value. Our prior is much less restrictive than this.

of b_t tend to be quite high until around 1980, but decrease steadily thereafter, indicating that the degree of persistence in the inflation gap has dropped over time. Such a finding is consistent with the Fed become increasingly intolerant of inflation being above implicit targets for long periods of time, particularly after Paul Volcker became Chairman of the Federal Reserve in 1979. There is also strong evidence of stochastic volatility, both in the inflation equation and in the one for trend inflation. This is consistent with the findings of Stock and Watson (2007) in their univariate model for inflation. It is interesting to note that, as in Stock and Watson (2007), both types of stochastic volatility increased throughout the 1970's, peaking around 1980, and falling subsequently. However, with our longer sample span, we are finding that the recent financial crisis was associated with a large increase in inflation volatility, but no increase in the volatility of trend inflation.

The previous discussion was based on an examination of point estimates and, with interval estimates being fairly wide, the question of how statistically important these findings are naturally arises. More formal conclusions about the consistency of long-run survey based inflation expectations with trend inflation can be based on the posterior model probabilities presented in Figure 5. With regards to d_{0t} and d_{1t} individually, evidence is somewhat ambiguous, with $P(d_{0t} = 0|Data)$ and $P(d_{1t} = 1|Data)$ both being around a half for much of the sample, although there is a substantial drop in the latter in the late 1970s and early 1980s. However, the joint probability, $P(d_{0t} = 0, d_{1t} = 1|Data)$, is more definitive. It indicates that between the mid-1970s and 1990, the probability of this joint restriction holding was near zero. Even though there is some imprecision in our estimates of the time-varying coefficients, the near-zero joint probability points to significant information in the data on changes over time in these coefficients. For this period at least, there is found to be a disconnect between the econometric estimate of trend inflation with the survey-based long-run inflation forecasts. Given this finding, it is not surprising that Figure 6, which calculates $P(d_{0t} = d_{0,1980}, d_{1t} = d_{1,1980}|Data)$ provides some evidence of time variation in these parameters, with the beginning and end of the sample indicating the most divergence from 1980s values. It is also worth noting that most of the evidence against the hypotheses that $d_{0t} = 0, d_{1t} = 1$ and that they are constant over time arises through the behavior of d_{1t}. This coefficient wanders farthest from the restriction of interest ($d_{1t} = 1$), whereas there is less evidence that d_{0t} wanders far from zero. Finally, Figure 2 provides support for the conclusion that b_t is time-varying and, in particular, decreases markedly as part of the Great Moderation of the business cycle in the early to mid 1980s.

Figure 7 plots trend inflation using the UCSV-AR model. A comparison to Figure 1 indicates the UCSV-AR estimate of trend inflation does not track the long-run survey forecasts as well as our model. In addition, the interval estimate is quite wide; estimating trend inflation without the information available in the long-run inflation expectation included in our proposed model yields a much less precise estimate. The right hand panel of Figure 7 calculates the dynamic probabilities that the trend inflation estimates equal the long-run forecasts. After the late 1980s these are very close to one, although prior to this there were periods that this probability becomes much lower. The fact that these probabilities tend to be higher than those in Figure 5 reflects the imprecision of trend inflation estimates provided by the UCSV models. The long-run forecasts are more likely to fall within the wider interval estimates produced by the UCSV-AR model than with our model.

It is also worth noting that the Bayes factor comparing the unrestricted version of our model against a restricted version without an MA term (5) is 1.5. This provides us with some weak support in favor of its inclusion. In the empirical appendix, it can be seen that that most of the posterior probability is associated with positive but small values for the MA coefficient.

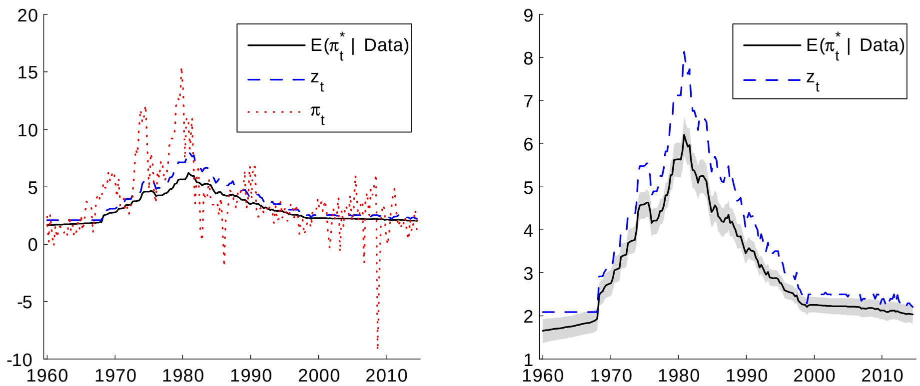

Figure 1: Posterior means and quantiles (16% and 84%) of π_t^*.

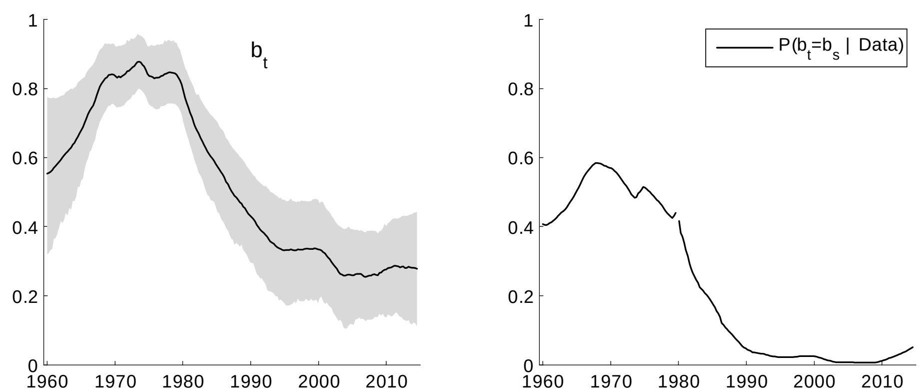

Figure 2: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s =$ 1980Q1.

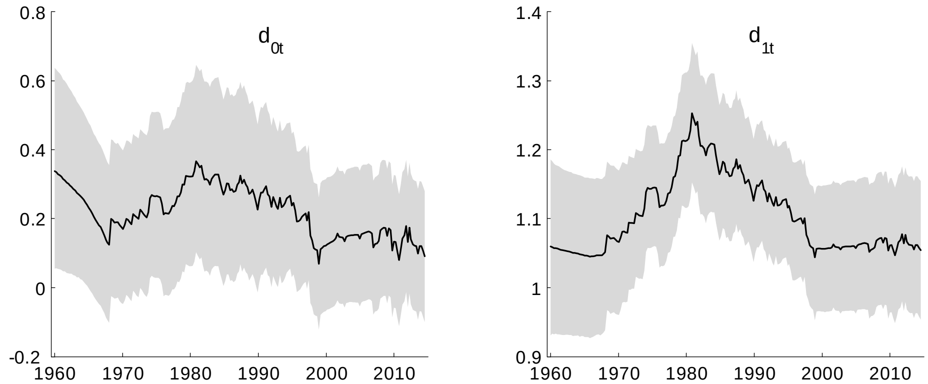

Figure 3: Posterior means and quantiles (16% and 84%) of d_{it}.

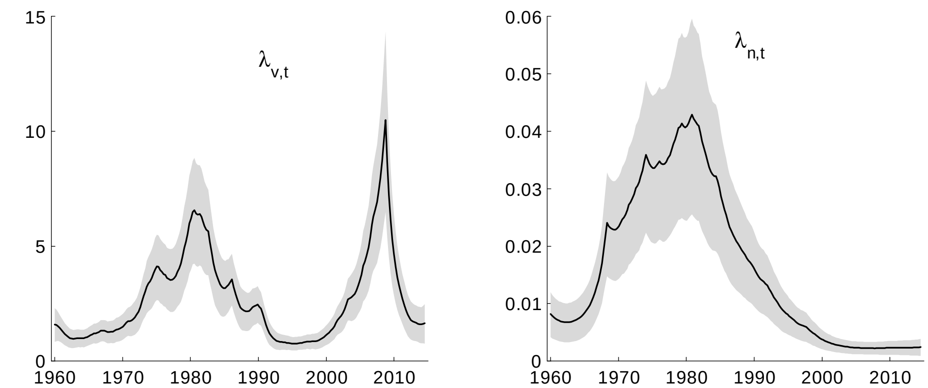

Figure 4: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

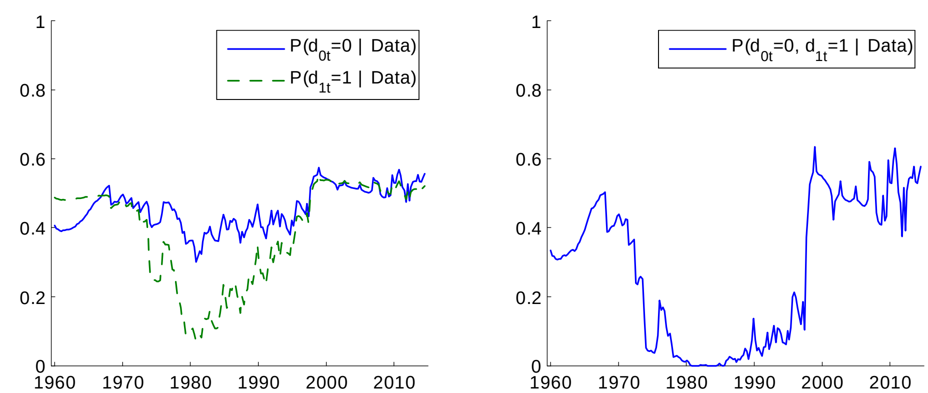

Figure 5: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}

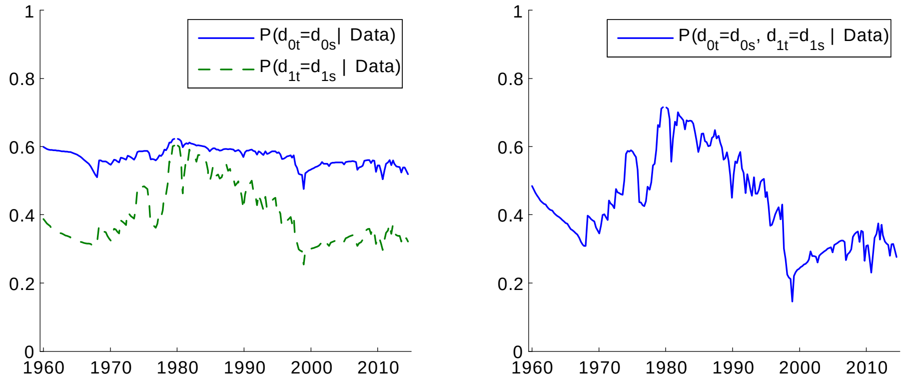

Figure 6: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

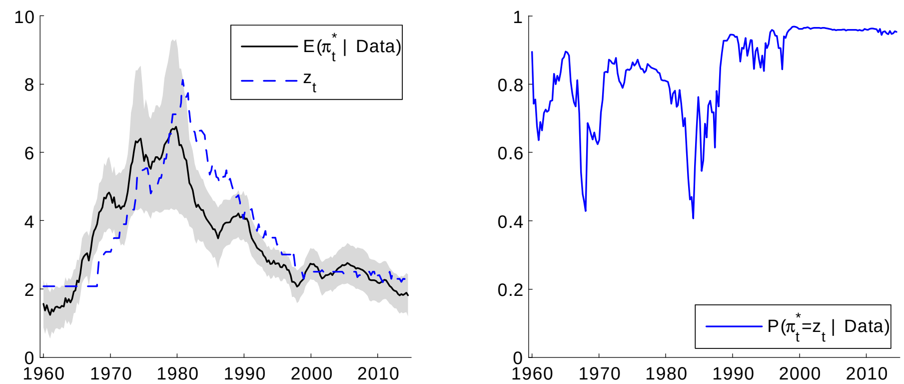

Figure 7: Posterior means and quantiles (16% and 84%) of π_t^* and dynamic probabilities $P(\pi_t^* = z_t|Data)$ for the UCSV-AR model.

4.2 Results Using PCE Inflation and PTR-PCE Forecasts

In this sub-section, the inflation measure is PCE inflation, and the long-run inflation expectations measure is PTR-PCE. Apart from some minor differences, results are the same as for the two previous cases. One more substantial difference is that now we are finding stronger evidence in favor of the MA error in (5) in that the relevant Bayes factor is 3.8 with a point estimate of nearly 0.4. As with the previous cases, we are finding sensible smooth trend inflation estimates which match up fairly well with long-run forecasts. But this matching is not perfect, so that equating trend inflation with survey-based measures looks questionable. For instance, $P(d_{0t} = 0, d_{1t} = 1|Data)$ is near zero for most of the 1980s and 1990s.

The trend inflation estimates provided by the UCSV-AR model are mostly sensible. However, they do not match up with z_t that closely and, even allowing for the wide interval estimate, there are many periods where the probability that trend inflation equals the long-run forecasts is very small.

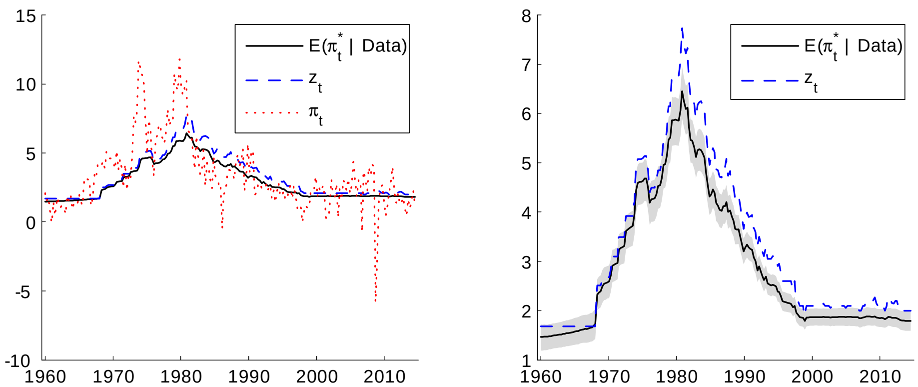

Figure 8: Posterior means and quantiles (16% and 84%) of π^*_t.

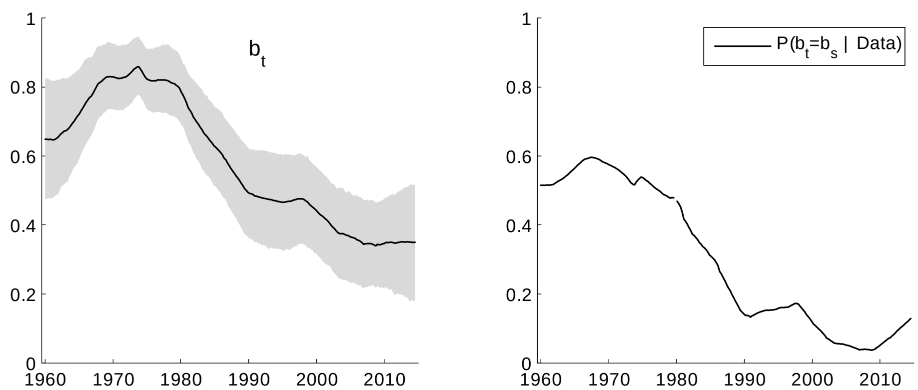

Figure 9: Posterior means and quantiles (16% and 84%) of b^*_t, and the dynamic probabilities that $b_t = b_s$ with $s = 1980Q1$.

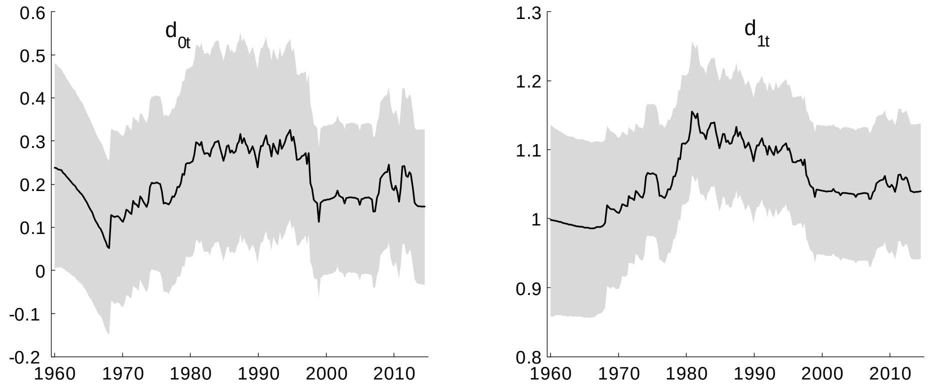

Figure 10: Posterior means and quantiles (16% and 84%) of d_{it}.

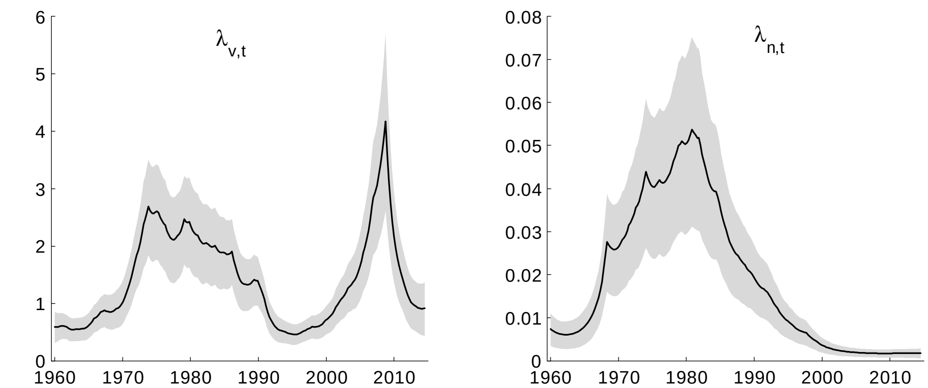

Figure 11: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

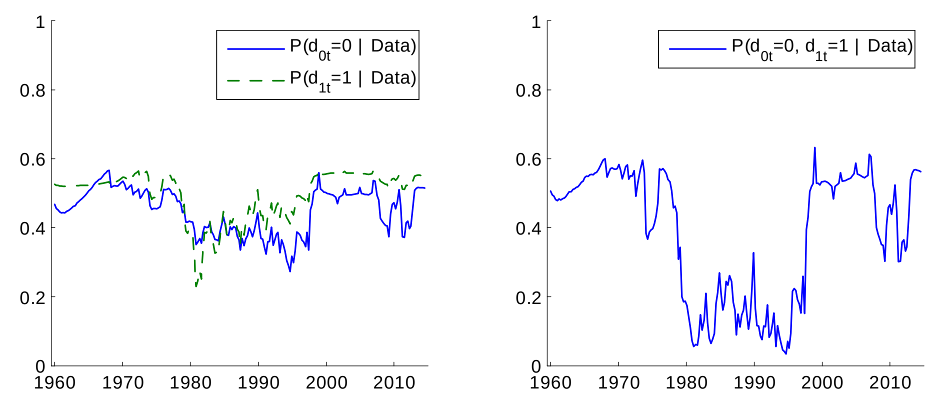

Figure 12: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

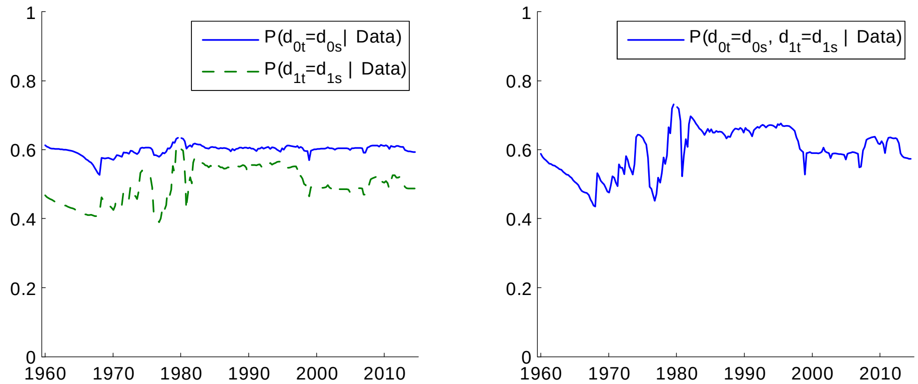

Figure 13: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

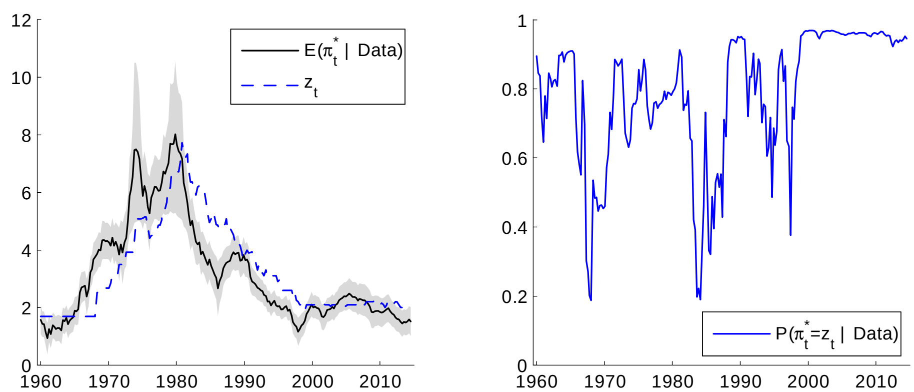

Figure 14: Posterior means and quantiles (16% and 84%) of π_t^* and dynamic probabilities $P(\pi_t^* = z_t|Data)$ for UCSV-AR model.

4.3 Results Using CPI inflation and Blue Chip CPI Forecasts

This section of empirical results uses the Blue Chip 6-10 year forecast as the measure of long run inflation expectations. Remember that these forecasts are available for a shorter time span and only begin in 1979Q4. This sub-section contains results where the inflation measure is CPI-based and the Blue Chip forecasts are of CPI inflation.

A general pattern we are finding with the Blue Chip forecasts is that there is much more evidence in favor of an MA process. In this sub-section, we are finding a Bayes factor in favor of its inclusion to be 31.8. We are also finding that the probability that $d_{0t} = 0$ and $d_{1t} = 1$ increases over time, which contrasts with results found using PTR-based long-run surveys. And we are also finding less evidence in favor of time variation in the coefficients. Although this reduced evidence of time variation could partly be due to the shorter sample span, it seems to be more due to the difference in inflation expectations.

In unreported results, we found that estimating the model using the PTR-CPI measure of expectations over the shorter sample yields results very similar to those for the full sample. Overall, though, our main results are found to be robust to changes in data. In particular, our model, by incorporating survey-based information, is producing more sensible estimates of trend inflation than a UCSV-AR model, but that simply equating trend inflation with the survey-based forecasts is not a sensible thing to do.

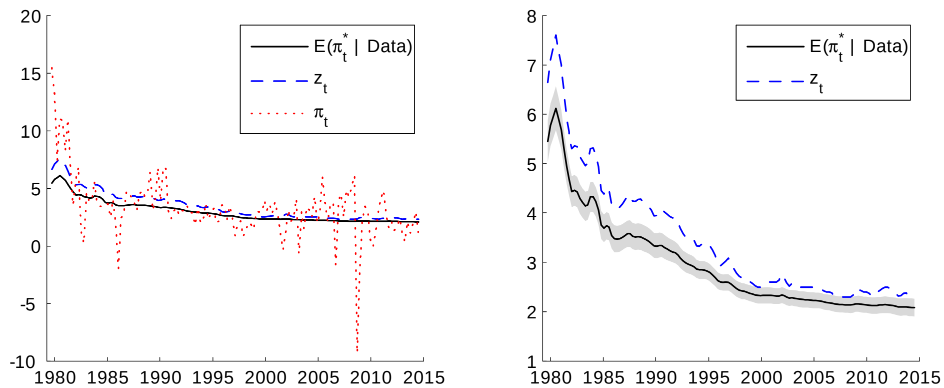

Figure 15: Posterior means and quantiles (16% and 84%) of π_t^*.

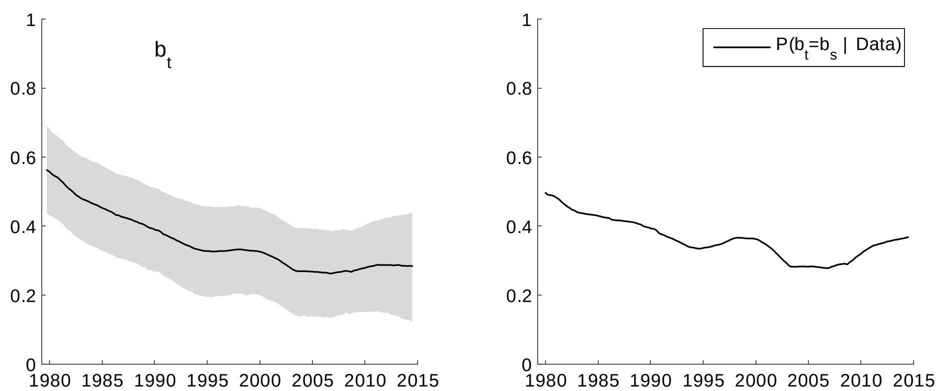

Figure 16: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = $ 1980Q1.

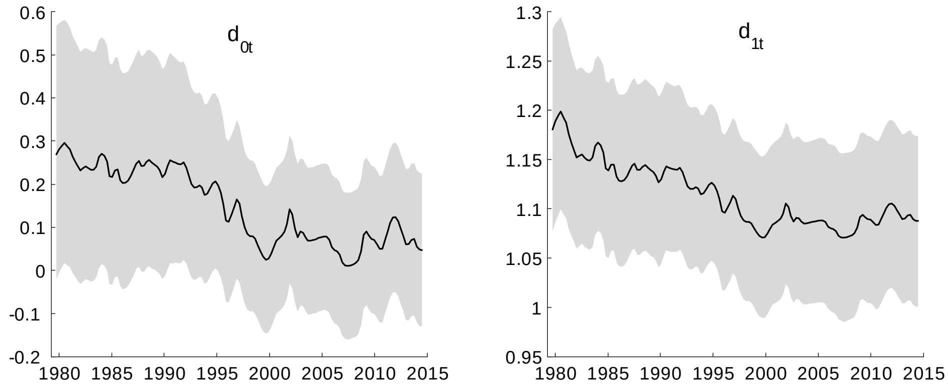

Figure 17: Posterior means and quantiles (16% and 84%) of d_{it}.

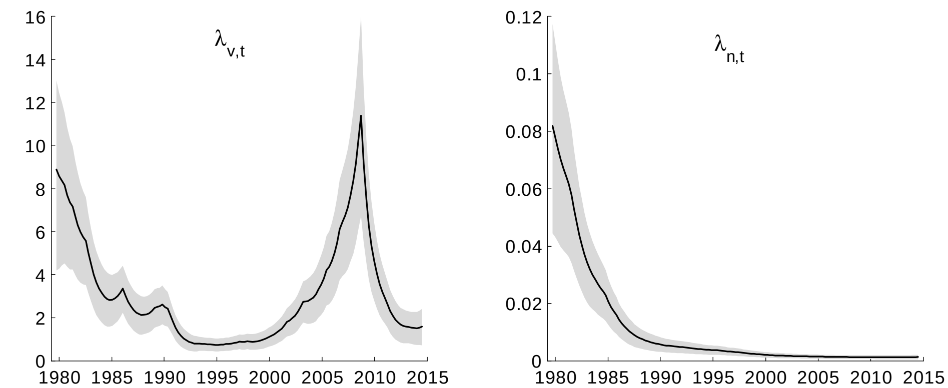

Figure 18: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

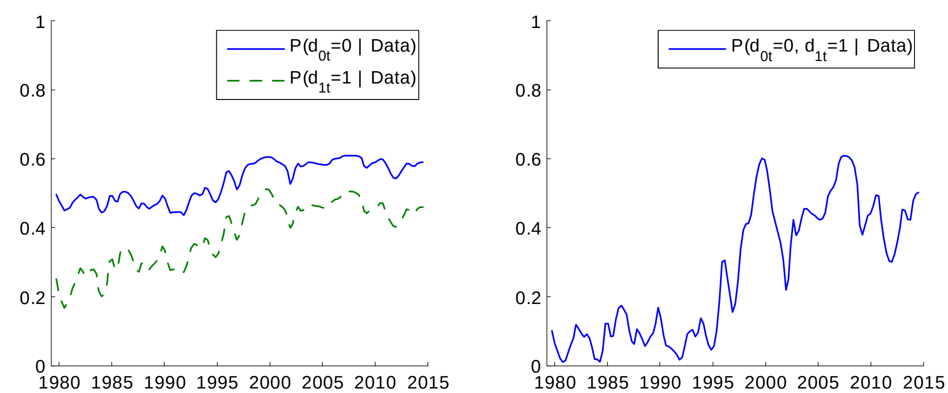

Figure 19: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

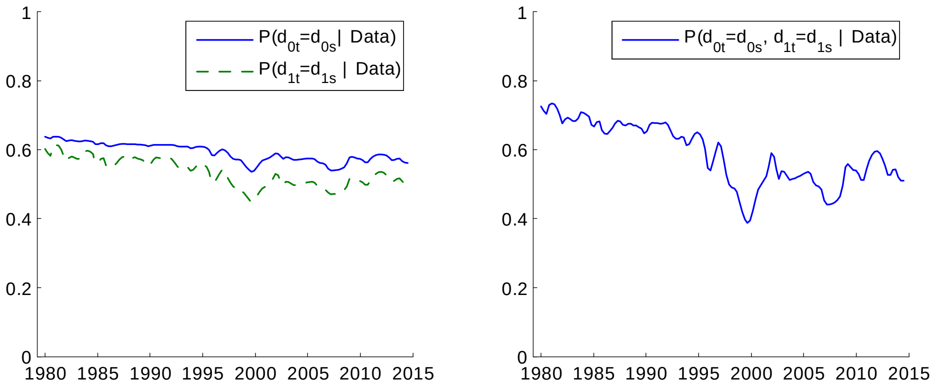

Figure 20: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

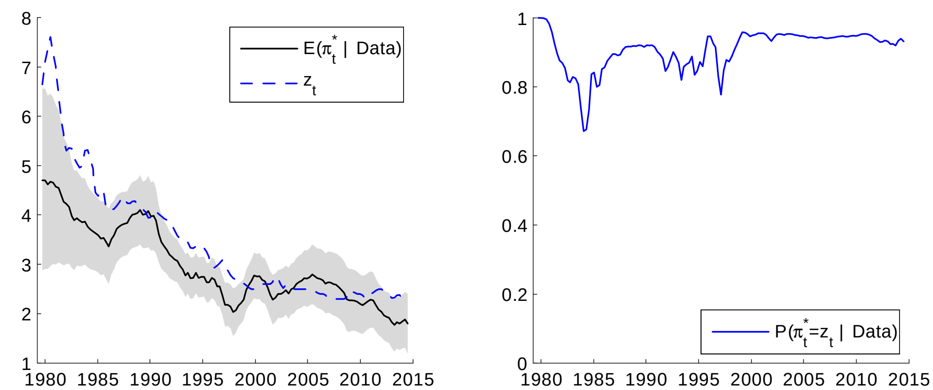

Figure 21: Posterior means and quantiles (16% and 84%) of π_t^* and dynamic probabilities $P(\pi_t^* = z_t|Data)$ for UCSV-AR model.

4.4 Forecasting

The primary purpose of this paper is to develop an appropriate model for investigating the relationship between inflation, trend inflation and inflation expectations. However, it is also of interest to see whether it forecasts better than plausible alternatives. To this end, we carry out a pseudo out-of-sample forecasting exercise. For the sake of brevity, we present results for CPI and PCE inflation using the PTR long-run forecasts. The evaluation period is from 1975Q1 to 2014Q3.[12] In addition to our proposed method, we consider two benchmarks. One is the UCSV model of Stock and Watson (2007), implemented as described in Section 2 (for comparability to Stock and Watson, in the forecast comparison we use the UCSV model and not the UCSV-AR model included in our

[12] We repeated the analysis with a shorter forecast evaluation period beginning in 1985Q1 (after the Great Moderation) and found results to be qualitatively similar.

full-sample model estimates reported above). The other is an AR(1) model in "gap form" similar to that used in Faust and Wright (2013), which they describe as "amazingly hard to beat by much". We call this the Faust and Wright model below.[13] We add stochastic volatility to this model to aid in comparability with our own. Specifically, we define the gap as $g_t = \pi_t - z_t$ and use the model:

$$g_t = \beta_0 + \beta_1 g_{t-1} + \epsilon_{g,t}, \quad \epsilon_{g,t} \sim N(0, \lambda_{g,t}),$$
$$\log(\lambda_{g,t}) = \log(\lambda_{g,t-1}) + \nu_{g,t}, \quad \nu_{g,t} \sim N(0, \phi_g),$$

where we assume $|\beta_1| < 1$. The forecast for π_{t+k} given data till time t is computed by adding z_t to a forecast for g_{t+k}. All other modeling choices, including priors, are the same as for our model.

Tables 1 and 2 evaluate forecast performance for CPI and PCE inflation, respectively, using root mean squared forecast errors (RMSFEs) and sums of log predictive likelihoods relative to the UCSV model. For CPI inflation, we tend to find some small improvements in RMSFE relative to UCSV, particularly at longer horizons. However, using predictive likelihoods as measures of forecast performance, our model is beating the UCSV model by a substantial amount, particularly at longer forecast horizons. Our forecasts do not beat the AR(1) model in gap form, but our model is at least competitive with one which the inflation forecasting literature has found to be among the top forecasting models.

For PCE inflation, neither our model nor Faust and Wright can beat UCSV when RMSFEs are used to measure forecast performance. However, when using predictive likelihoods, we are clearly beating UCSV at virtually every forecast horizon whereas Faust and Wright is inferior to UCSV at virtually every horizon. Hence, for PCE inflation ours is, overall, the best forecasting model.

Table 1: RMSFEs and log predictive likelihood for forecasting CPI, 1975Q1 to 2014Q3.

Relative RMSFE					
	1Q	2Q	4Q	8Q	12Q
UCSV	1.00	1.00	1.00	1.00	1.00
Faust-Wright	1.00	1.01	0.98	0.92	0.90
New model	0.98	1.00	1.03	0.98	0.96
Relative log predictive likelihood					
	1Q	2Q	4Q	8Q	12Q
UCSV	0.00	0.00	0.00	0.00	0.00
Faust-Wright	0.14	2.91	8.95	14.13	17.56
New model	3.81	4.26	5.71	9.48	14.30

Hence, for both inflation measures, our model has been found to be as good or better than popular and successful alternatives in terms of forecast performance, particularly

[13] Our specification generalizes their "fixed ρ" model by including an intercept and estimated coefficients. Accordingly, our model takes the same form as their "AR-gap" model, except that, at all horizons, we use the 1-step ahead form of the model and iterated forecasts, whereas they use a direct multi-step form of the model.

Table 2: RMSFEs and log predictive likelihood for forecasting PCE, 1975Q1 to 2014Q3.

Relative RMSFE					
	1Q	2Q	4Q	8Q	12Q
UCSV	1.00	1.00	1.00	1.00	1.00
Faust-Wright	1.01	1.03	1.00	0.98	0.99
New model	0.99	1.02	1.03	1.01	1.02
Relative log predictive likelihood					
	1Q	2Q	4Q	8Q	12Q
UCSV	0.00	0.00	0.00	0.00	0.00
Faust-Wright	-1.40	-0.88	1.90	-4.12	-3.86
New model	2.19	1.47	2.81	-0.96	2.56

as measured by predictive likelihoods. The fact that RMSFEs rarely differ much for the three approaches we are comparing indicates that most of the benefits of our model arise not from improving point forecasts, but from more accurate estimation of higher moments of the predictive distribution.

5 Summary and Conclusion

In this paper, we have developed a bivariate model of inflation and inflation expectations that incorporates empirically-important features such as time-varying parameters and stochastic volatility. In a broad sense, we have used our model to investigate the relationship between these two variables. In a narrower sense, we have investigated the degree to which survey-based long-run inflation forecasts can be used to inform estimates of trend inflation. In an extensive empirical exercise involving three different measures of inflation and two different sources for long-run inflation forecasts we find a consistent story: Long-run inflation forecasts do provide useful additional information in informing estimates of trend inflation. However, the forecasts themselves cannot simply be equated with trend inflation. In out-of-sample forecasting, our model yields point and density forecasts that are at least as good as those from other models that have been found successful in the inflation forecasting literature.

The history captured by our estimates indicates the distinction between trend inflation and long-run inflation expectations captured by surveys is practically important. For example, as noted in the introduction, for most of the period since 2008, inflation in the PCE price index has run below the Federal Reserve's longer-run inflation objective of 2 percent. Over the past couple of years, inflation has declined to very low levels. Yet, for several years before the recession that began in 2007, inflation ran steadily above target. Some estimates of trend inflation based entirely on inflation — as in the UCSV specification of Stock and Watson (2007) — have moved around with inflation, rising in the early to mid-2000s and declining markedly as of late 2014. At the other extreme, long-run inflation expectations measured from the Survey of Professional Forecasters have remained steady around 2 percent (with occasional up-ticks and down-ticks). Drawing on the information in both inflation and the survey's long-run expectation, our model's estimate of trend is much smoother than the estimate from a univariate UCSV specification, implying the trend to be stable in the face of both the rise of inflation in the years before the recession and the fall since the recession. In fact, our model estimates show trend inflation to be even more stable than the survey expectation. However, in keeping with a historical bias in the survey forecast, our estimate of trend inflation has for some time been stable, slightly below the survey expectation.

References

Bednar, W. and Clark, T. (2014). Methods for evaluating recent trend inflation. *Economic Trends*, Federal Reserve Bank of Cleveland, March 28.

Bernanke, B. (2007). Inflation expectations and inflation forecasting. Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, July 10.

Chan, J. (2013). Moving average stochastic volatility models with application to inflation forecast. *Journal of Econometrics*, 176, 162–172.

Chan, J. (2015). The stochastic volatility in mean model with time-varying parameters: An application to inflation modeling. *Journal of Business and Economic Statistics*, forthcoming.

Chan, J. and Jeliazkov, I. (2009). Efficient simulation and integrated likelihood estimation in state space models. *International Journal of Mathematical Modelling and Numerical Optimisation*, 1, 101-120.

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

Chan, J., Koop, G. and Potter, S. M. (2015). A bounded model of time variation in trend inflation, NAIRU and the Phillips curve. *Journal of Applied Econometrics*, forthcoming.

Clark, T. and Doh, T. (2014). Evaluating alternative models of trend inflation. *International Journal of Forecasting*, 30, 426-448.

Cogley, T., Primiceri, G. and Sargent, T. (2010). Inflation-gap persistence in the US. *American Economic Journal: Macroeconomics*, 2, 43-69.

Del Negro, M. and Schorfheide, F. (2013). DSGE model-based forecasting. In G. Elliott & A. Timmermann (Eds.), Handbook of economic forecasting, volume 2. Amsterdam: North Holland.

Faust, J. and Wright, J. (2013). Forecasting inflation. In G. Elliott & A. Timmermann (Eds.), Handbook of economic forecasting, volume 2. Amsterdam: North Holland.

Fuhrer, J., Olivei, G. and Tootell, G. (2012). Inflation dynamics when inflation is near zero. *Journal of Money, Credit and Banking,* 44, 83-122.

Garnier, C., Mertens, E. and Nelson, E. (2015), Trend Inflation in advanced economies, *International Journal of Central Banking*, forthcoming.

Jarocinski, M. and Lenza, M. (2015). Output gap and inflation forecasts in a Bayesian dynamic factor model of the euro area, manuscript, European Central Bank.

Kim, S., Shepherd, N. and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65, 361-393.

Koop, G. (2003). Bayesian Econometrics. Wiley & Sons, New York.

Koop, G. and Korobilis, D. (2010). Bayesian multivariate time series methods for empirical macroeconomics. *Foundations and Trends in Econometrics*, 3, 267-358.

Koop, G., Leon-Gonzalez, R. and Strachan, R. (2010). Dynamic probabilities of restrictions in state space models: An application to the Phillips curve. *Journal of Business and Economic Statistics*, 28, 370-379

Kozicki, S. and Tinsley, P. (2001). Shifting endpoints in the term structure of interest rates. *Journal of Monetary Economics*, 47, 613-652.

Kozicki, S. and Tinsley, P. (2012). Effective use of survey information in estimating the evolution of expected inflation. *Journal of Money, Credit and Banking*, 44, 145-169.

Mertens, E. (2015). Measuring the level and uncertainty of trend inflation, *Review of Economics and Statistics*, forthcoming.

Mehrotra, A. and Yetman, J. (2014). Decaying expectations: what inflation forecasts tell us about the anchoring of inflation expectations, BIS Working Papers No. 464.

Mishkin, F. (2007). Inflation dynamics. Speech at the Annual Macro Conference, Federal Reserve Bank of San Francisco, March 23.

Morley, J., Piger, J. and Rasche, R. (2015). Inflation in the G7: mind the gap(s)? *Macroeconomic Dynamics* 19, 883-912.

Nason, J. and Smith, J. (2014). Measuring the slowly evolving trend in US inflation with professional forecasts. Centre for Applied Macroeconomic Analysis, Australian National University, Working Paper 7/2014.

Rudd, J. and Peneva, E. (2015). The passthrough of labor costs to price inflation, Finance and Economics Discussion Series Working Paper 2015-042, Board of Governors of the Federal Reserve System.

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

Stock, J. and Watson, M. (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking* 39, 3-33.

Stock, J. and Watson, M. (2015). Core inflation and trend inflation, NBER Working Paper No. 21282.

Verdinelli, I. and Wasserman, L. (1995). Computing Bayes factors using a generalization of the Savage-Dickey density ratio, *Journal of the American Statistical Association*, 90, 614-618.

Wright, J. (2013). Evaluating real time VAR forecasts with an informative democratic prior. *Journal of Applied Econometrics*, 28, 762-776.

Technical Appendix

In this appendix, we specify the prior and MCMC algorithm used in this paper. We also provide additional details on the calculation of dynamic model probabilities using the SDDR.

The model is given in (4), (5), (6), (7), (8), (9), (10) and (11). We initialize the state equations (6), (8), (7) and (11) by $\pi_1^* \sim N(\pi_0^*, \lambda_{n,1}V_{\pi^*})$, $b_1 \sim N(b_0, V_b)$, $d_{i1} \sim N(\mu_{di}, \sigma_{di}^2/(1-\rho_{di}^2))$, $i = 0, 1$, and $\log(\lambda_{i,1}) \sim N(\log(\lambda_{i,0}), V_{\lambda_i})$, $i = v, n$, with $\lambda_{i,0} = 1$, $b_0 = \pi_0^* = 0$ and $V_{\lambda_i} = V_b = V_{\pi^*} = 100$. These are relatively non-informative choices.

For later reference, let $\pi = (\pi_1, \ldots, \pi_T)'$ and $d = (d_{01}, d_{11}, \ldots, d_{0T}, d_{1T})'$, and similarly define z, π^*, b, λ_v and λ_n. In addition, let θ denote the model parameters, i.e., $\theta = (\psi, \mu_{d0}, \mu_{d1}, \rho_{d0}, \rho_{d1}, \sigma_{d0}^2, \sigma_{d1}^2, \sigma_b^2, \sigma_w^2, \phi_v, \phi_n)'$.

We assume independent priors for elements of the parameter vector θ which are proper and weakly informative. The priors for μ_{di} and ρ_{di} are:

$$\mu_{d0} \sim N(a_0, V_\mu), \quad \mu_{d1} \sim N(a_1, V_\mu), \quad \rho_{di} \sim TN_{(c_1,c_2)}(a_2, V_\rho),$$

where the $TN_{(c_1,c_2)}(a_1, a_2)$ denotes the $N(a_1, a_2)$ distribution truncated to the interval (c_1, c_2) and we set $a_0 = 0$, $a_1 = 1$, $a_2 = 0.95$, $V_\mu = 0.1^2$ and $V_\rho = 0.1^2$. These choices imply relatively informative priors centered at the values which imply trend inflation is equal to long-run inflation forecasts (apart from a mean zero error). For the MA(1) coefficient, we consider the relatively non-informative prior which restricts the MA process to be invertible: $\psi \sim TN_{(-1,1)}(0, V_\psi)$ with $V_\psi = 0.25^2$. Finally, we assume independent inverse gamma priors for the variance parameters. In particular, the degree of freedom parameters are all set to the relatively non-informative value of 5, and the scale parameters are set such that $E(\sigma_{d0}^2) = E(\sigma_w^2) = E(\phi_v) = E(\phi_n) = 0.01$ and $E(\sigma_{d1}^2) = E(\sigma_b^2) = 0.001$. These values are chosen to reflect the desired smoothness of the corresponding state transition. For example, the prior mean for σ_{d0}^2 implies that with high probability the difference between consecutive d_{0t} lies within the values -0.2 and 0.2.

To estimate the model in (4), (5), (6), (7), (8), (9), (10) and (11), we extend the MCMC sampler developed in Chan, Koop and Potter (2013) which was used for a univariate bounded inflation trend model. Moreover, we also incorporate the sampler in Chan (2013) for handling the MA innovations with stochastic volatility. Specifically, we sequentially draw from the following densities:

1. $p(\pi^* \,|\, Data, b, d, \lambda_v, \lambda_n, \theta)$;
2. $p(b \,|\, Data, \pi^*, d, \lambda_v, \lambda_n, \theta)$;
3. $p(d \,|\, Data, \pi^*, b, \lambda_v, \lambda_n, \theta)$;
4. $p(\lambda_v, \lambda_n \,|\, Data, \pi^*, b, d, \theta)$;
5. $p(\mu_{d0}, \mu_{d1} \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\mu_{d0},\mu_{d1}\}})$;
6. $p(\sigma_{d0}^2, \sigma_{d1}^2 \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\sigma_{d0}^2,\sigma_{d1}^2\}})$;
7. $p(\rho_{d0}, \rho_{d1} \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\rho_{d0},\rho_{d1}\}})$;

8. $p(\psi \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\psi\}})$;

9. $p(\sigma_b^2, \sigma_w^2, \phi_v, \phi_n \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\sigma_b^2, \sigma_w^2, \phi_v, \phi_n\}})$.

Step 1: To implement Step 1, note that information about π^* comes from three sources: the two measurement equations (4) and (5), and the state equation (6). We derive an expression for each component in turn. First, write (4) as

$$H_b\pi = H_b\pi^* + \tilde{\alpha}_{\pi^*} + v, \quad v \sim N(0, \Lambda_v),$$

where $\tilde{\alpha}_{\pi^*} = (b_1(\pi_0 - \pi_0^*), 0, \ldots, 0)'$, $\Lambda_v = \mathrm{diag}(\lambda_{v,1}, \ldots, \lambda_{v,T})$ and

$$H_b = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ -b_2 & 1 & 0 & \cdots & 0 \\ 0 & -b_3 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & \cdots & -b_T & 1 \end{pmatrix}.$$

Since $|H_b| = 1$ for any b, H_b is invertible. Therefore, we have

$$(\pi \,|\, \pi^*, b, \lambda_v) \sim N(\pi^* + \alpha_{\pi^*}, (H_b'\Lambda_v^{-1}H_b)^{-1}),$$

with log density

$$\log p(\pi \,|\, \pi^*, b, \lambda_v) \propto -\frac{1}{2}(\pi - \pi^* - \alpha_{\pi^*})'H_b'\Lambda_v^{-1}H_b(\pi - \pi^* - \alpha_{\pi^*}), \tag{12}$$

where $\alpha_{\pi^*} = H_b^{-1}\tilde{\alpha}_{\pi^*}$. Note that H_b is a band matrix and α_{π^*} can be obtained quickly by solving the band system $H_b x = \tilde{\alpha}_{\pi^*}$ for x without computing the inverse H_b^{-1}.

The second component comes from (5) which can be written as:

$$z = d_0 + X_{\pi^*}\pi^* + H_\psi\epsilon_z, \quad \epsilon_z \sim N(0, \sigma_w^2 I_T),$$

where $d_0 = (d_{01}, \ldots, d_{0T})'$, $X_{\pi^*} = \mathrm{diag}(d_{11}, \ldots, d_{1T})$ and

$$H_\psi = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ \psi & 1 & 0 & \cdots & 0 \\ 0 & \psi & 1 & \cdots & 0 \\ \vdots & & \ddots & \ddots & \vdots \\ 0 & 0 & \cdots & \psi & 1 \end{pmatrix}.$$

Thus, ignoring any terms not involving π^*, we have

$$\begin{aligned} \log p(z \,|\, \pi^*, d, \sigma_w^2) &\propto -\frac{1}{2\sigma_w^2}(z - d_0 - X_{\pi^*}\pi^*)'(H_\psi H_\psi'^{-1}(z - d_0 - X_{\pi^*}\pi^*), \\ &= -\frac{1}{2\sigma_w^2}(\tilde{z} - \tilde{X}_{\pi^*}\pi^*)'(\tilde{z} - \tilde{X}_{\pi^*}\pi^*), \end{aligned} \tag{13}$$

where $\tilde{z} = H_\psi^{-1}(z - d_0)$ and $\tilde{X}_{\pi^*} = H_\psi^{-1} X_{\pi^*}$. Since H_ψ is a band matrix, $\tilde{z}$ can be computed quickly by solving a linear system of equations without finding the inverse H_ψ^{-1}. The matrix $\tilde{X}_{\pi^*}$ is lower triangular that is in general not banded. However, most of the elements away from the main diagonal band are close to zero. In our implementation we construct a band approximation by replacing all elements below 10^{-6} with 0. Since the cut-off point is so small, it has no impact on the results, but it substantially speeds up the computation.

The third component is contributed by the state equation (6):

$$\log p(\pi^* \,|\, \lambda_n) \propto -\frac{1}{2}(\pi^* - \delta_{\pi^*})' H' \Lambda_n^{-1} H (\pi^* - \delta_{\pi^*}), \tag{14}$$

where H is the $T \times T$ first difference matrix, $\Lambda_n = \text{diag}(\lambda_{n,1} V_{\pi^*}, \lambda_{n,2}, \ldots, \lambda_{n,T})$ and $\delta_{\pi^*} = H^{-1}(\pi_0^*, 0, \ldots, 0)'$. Then, combining (12), (13) and (14), we finally obtain

$$\begin{aligned}
&\log p(\pi^* \,|\, Data, b, d, \lambda_v, \lambda_n, \theta) \\
&\propto -\frac{1}{2}(\pi - \pi^* - \alpha_{\pi^*})' H_b' \Lambda_v^{-1} H_b (\pi - \pi^* - \alpha_{\pi^*}) - \frac{1}{2\sigma_w^2}(\tilde{z} - \tilde{X}_{\pi^*}\pi^*)'(\tilde{z} - X_{\pi^*}\pi^*) \\
&\quad - \frac{1}{2}(\pi^* - \delta_{\pi^*})' H' \Lambda_n^{-1} H (\pi^* - \delta_{\pi^*}), \\
&\propto -\frac{1}{2}(\pi^* - \hat{\pi}^*)' K_{\pi^*} (\pi^* - \hat{\pi}^*),
\end{aligned}$$

which is the kernel of the $N(\hat{\pi}^*, K_{\pi^*}^{-1})$ distribution, where

$$\begin{aligned}
K_{\pi^*} &= \left(H_b' \Lambda_v^{-1} H_b + \frac{1}{\sigma_w^2} \tilde{X}_{\pi^*}' \tilde{X}_{\pi^*} + H' \Lambda_n^{-1} H \right)^{-1}, \\
\hat{\pi}^* &= K_{\pi^*}^{-1} \left(H_b' \Lambda_v^{-1} H_b (\pi - \alpha_{\pi^*}) + \frac{1}{\sigma_w^2} \tilde{X}_{\pi^*}' \tilde{z} + H' \Lambda_n^{-1} H \delta_{\pi^*} \right).
\end{aligned}$$

If we use the band approximation of $\tilde{X}_{\pi^*}$ as described above, the precision K_{π^*} is also a band matrix. Then, we use the precision sampler in Chan and Jeliazkov (2009) to sample π^* from the conditional distribution $(\pi^* \,|\, Data, b, d, \lambda_v, \lambda_n, \theta)$.

Step 2: Next, we derive the conditional density $p(b \,|\, Data, \pi^*, d, \lambda_v, \lambda_n, \theta)$. Due to the inequality restriction $0 < b_t < 1$, this joint density is non-normal. We first rewrite (4) as:

$$\tilde{\pi} = X_b b + v, \quad v \sim N(0, \Lambda_v),$$

where $\tilde{\pi} = (\pi_1 - \pi_1^*, \ldots, \pi_T - \pi_T^*)'$ and $X_b = \text{diag}(\pi_0 - \pi_0^*, \ldots, \pi_{T-1} - \pi_{T-1}^*)$. It follows that the log density of $(\pi \,|\, \pi^*, b, \lambda_v)$ can also be written as follows:

$$\log p(\pi \,|\, \pi^*, b, \lambda_v) \propto -\frac{1}{2}(\tilde{\pi} - X_b b)' \Lambda_v^{-1} (\tilde{\pi} - X_b b), \tag{15}$$

Next, write (7) as

$$Hb = \tilde{\delta}_b + \epsilon_b,$$

where $\tilde{\delta}_b = (b_0, 0, \ldots, 0)'$ and the elements of ϵ_b are independent truncated normal variables. Note that $\Pr(0 < b_1 < 1) = \Phi((1-b_0)/\sqrt{V_b}) - \Phi(b_0/\sqrt{V_b})$ and

$$\Pr(0 < b_t < 1) = \Phi\left(\frac{1-b_{t-1}}{\sigma_b}\right) - \Phi\left(\frac{-b_{t-1}}{\sigma_b}\right),$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard normal distribution. Hence, the prior density for b is given by

$$\log p(b \,|\, \sigma_b^2) \propto \frac{1}{2}(b-\delta_b)'H'\Sigma_b^{-1}H(b-\delta_b) + g_b(b, \sigma_b^2), \tag{16}$$

where $\Sigma_b = \text{diag}(V_b, \sigma_b^2, \ldots, \sigma_b^2)$, $\delta_b = H^{-1}\tilde{\delta}_b$ and

$$g_b(b, \sigma_b^2) = -\sum_{t=2}^{T} \log\left(\Phi\left(\frac{1-b_{t-1}}{\sigma_b}\right) - \Phi\left(\frac{-b_{t-1}}{\sigma_b}\right)\right).$$

Combining (15) and (16), we obtain

$$\log(b \,|\, Data, \pi^*, d, \lambda_v, \lambda_n, \theta) \propto -\frac{1}{2}(b-\hat{b})'K_b^{-1}(b-\hat{b}) + g_b(b, \sigma_b^2),$$

where

$$K_b = \left(H'\Sigma_b^{-1}H + X_b'\Lambda_v^{-1}X_b\right)^{-1}, \quad \hat{\tau}^{\pi} = K_b^{-1}(H_b'^{-1}\tilde{\delta}_b + X_b'\Lambda_v^{-1}\tilde{\pi}).$$

We follow Chan, Koop and Potter (2013) to sample b. Specifically, candidate draws are first obtained from the $N(\hat{b}, K_b^{-1})$ distribution using the precision sampler in Chan and Jeliazkov (2009), and they are accepted or rejected via an acceptance-rejection Metropolis-Hastings step.

Step 3: To sample from $p(d \,|\, Data, \pi^*, b, \lambda_v, \lambda_n, \theta)$, we first rewrite (5) and (8) as

$$\begin{aligned} z &= X_d d + H_\psi \varepsilon_z, & \varepsilon_z &\sim N(0, \sigma_w^2 I_T), \\ H_{\rho_d} d &= \tilde{\delta}_d + \varepsilon_d & \varepsilon_d &\sim N(0, \Sigma_d), \end{aligned}$$

where $\tilde{\delta}_d = (\mu_{d0}, \mu_{d1}, (1-\rho_{d0})\mu_{d0}, (1-\rho_{d1})\mu_{d1}, \ldots, (1-\rho_{d0})\mu_{d0}, (1-\rho_{d1})\mu_{d1})'$, $\Sigma_d = \text{diag}(\sigma_{d0}^2/(1-\rho_{d0}^2), \sigma_{d1}^2/(1-\rho_{d1}^2), \sigma_{d0}^2, \sigma_{d1}^2, \ldots, \sigma_{d0}^2, \sigma_{d1}^2)$,

$$X_d = \begin{pmatrix} 1 & \pi_1^* & 0 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \pi_2^* & 0 & \cdots & 0 \\ \vdots & & & & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & 1 & \pi_T^* \end{pmatrix}, \quad H_{\rho_d} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ -\rho_{d0} & 1 & 0 & 0 & \cdots & 0 \\ 0 & -\rho_{d1} & 1 & 0 & \cdots & 0 \\ \vdots & & \ddots & \ddots & & \vdots \\ 0 & 0 & \cdots & -\rho_{d0} & 1 & \\ 0 & 0 & \cdots & 0 & -\rho_{d1} & 1 \end{pmatrix}.$$

Using standard linear regression results (see, e.g., Koop, 2003, pp. 60-61), we have $(d \,|\, Data, \pi^*, b, \lambda_v, \lambda_n, \theta) \sim N(\hat{d}, K_d^{-1})$, where

$$K_d = \left(H_{\rho_d}'\Sigma_d^{-1}H_{\rho_d} + \frac{1}{\sigma_w^2}\tilde{X}_d'\tilde{X}_d\right)^{-1}, \quad \hat{d} = K_d^{-1}\left(H_{\rho_d}'\Sigma_d^{-1}\tilde{\delta}_d + \frac{1}{\sigma_w^2}\tilde{X}_d'(H_\psi^{-1}z)\right)$$

with $\tilde{X}_d = H_\psi^{-1} X_d$. As before, we construct a band approximation of $\tilde{X}_d$ by replacing all elements less than 10^{-6} with 0. Then, the precision K_d is a band matrix and the precision sampler in Chan and Jeliazkov (2009) is used to sample d.

Step 4: To implement Step 4, note that λ_v and λ_n are conditionally independent given the parameters and other states. Hence, we can draw them sequentially using the auxiliary mixture sampler of Kim, Shepherd and Chib (1998). See also Koop and Korobilis (2010), p. 308–310, for a textbook treatment. Note that in conventional implementations, a forward-filtering-backward-smoothing algorithm is used; here it is replaced by the more efficient precision sampler of Chan and Jeliazkov (2009).

Steps 5 and 6: Both the densities of (μ_{d0}, μ_{d1}) and $(\sigma^2_{d0}, \sigma^2_{d1})$ are standard. In fact, we have

$$(\mu_{di} \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\mu_{d0},\mu_{d1}\}}) \sim N(\hat{\mu}_{di}, K_{di}^{-1}),$$
$$(\sigma^2_{di} \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\sigma^2_{d0},\sigma^2_{d1}\}}) \sim IG(\nu_{di} + T/2, \tilde{S}_{di}),$$

where $K_{di} = 1/V_\mu + (1-\rho_{di}^2)/\sigma^2_{di} + (T-1)(1-\rho_{di})^2/\sigma^2_{di}$, $\hat{\mu}_{di} = K_{di}^{-1}(a_i/V_\mu + (1-\rho^2_{di})d_{i1}/\sigma^2_{di} + \sum_{t=2}^T (1-\rho_{di})(d_{it} - \rho_{di} d_{i,t-1})/\sigma^2_{di})$ and $\tilde{S}_{di} = S_{di} + ((1-\rho^2_{di})(d_{i1} - \mu_{di})^2 + \sum_{t=2}^T (d_{it} - \mu_{di}(1-\rho_{di}) - \rho_{di} d_{i,t-1})^2)/2$.

Step 7: It follows from (8) that

$$p(\rho_{di} \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\rho_{d0},\rho_{d1}\}}) \propto p(\rho_{di}) g_{\rho_{di}}(\rho_{di}) e^{-\frac{1}{2\sigma^2_{di}} \sum_{t=2}^T (d_{it} - \mu_{di} - \rho_{di}(d_{i,t-1} - \mu_{di}))^2},$$

where $p(\rho_{di})$ is the truncated normal prior for ρ_{di} and $g(\rho_{di}) = (1-\rho^2_{di})^{1/2} \exp(-\frac{1}{2\sigma^2_{di}}(1-\rho^2_{di})(d_{i1} - \mu_{di})^2)$. This conditional density is non-standard, and we implement an independence-chain Metropolis-Hastings step with proposal distribution $N(\hat{\rho}_{di}, K_{\rho_{di}}^{-1})$, where $K_{\rho_{di}} = 1/V_\rho + X'_{\rho_{di}} X_{\rho_{di}}/\sigma^2_{di}$ and $\hat{\rho}_{di} = K_{\rho_{di}}^{-1}(a_2/V_\rho + X'_{\rho_{di}} y_{\rho_{di}}/\sigma^2_{di})$, with $X_{\rho_{di}} = (d_{i1} - \mu_{di}, \ldots, d_{i,T-1} - \mu_{di})'$ and $y_{\rho_{di}} = (d_{i2} - \mu_{di}, \ldots, d_{iT} - \mu_{di})'$. Then, given the current draw ρ_{di}, a proposal ρ^*_{di} is accepted with probability $\min(1, g_{\rho_{di}}(\rho^*_{di})/g_{\rho_{di}}(\rho_{di}))$; otherwise the Markov chain stays at the current state ρ_{di}.

Step 8: To sample ψ, note that

$$\begin{aligned}\log p(\psi \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\psi\}}) &\propto \log p(z \,|\, \pi^*, d, \sigma^2_w) + \log p(\psi) \\ &\propto -\frac{1}{2\sigma^2_w}(z - d_0 - X_{\pi^*}\pi^*)'(H_\psi H_\psi')^{-1}(z - d_0 - X_{\pi^*}\pi^*) + \log p(\psi),\end{aligned}$$

where $p(\psi)$ is the prior density of ψ. Following Chan (2013), we sample ψ via an independence-chain Metropolis-Hastings step. Specifically, since this log density can be quickly evaluated using band matrix routines, we maximize it numerically to obtain the mode and negative Hessian, denoted as $\hat{\psi}$ and K_ψ, respectively. Then, we generate candidate draws from the $N(\hat{\psi}, K_\psi^{-1})$ distribution.

Step 9: To sample $\sigma^2_b, \sigma^2_w, \phi_v$ and ϕ_n, first note that these parameters are conditionally independent given the data and the states. Hence, we can sample each element one by one. The variance parameters σ^2_w, ϕ_v and ϕ_n follow inverse-Gamma distributions:

$$(\sigma^2_w \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\sigma^2_b,\sigma^2_w,\phi_v,\phi_n\}}) \sim IG\left(\nu_{\sigma^2_w} + \frac{T}{2}, S_{\sigma^2_w} + \frac{1}{2}\sum_{t=1}^{T}\epsilon^2_{z,t}\right)$$

$$(\phi_i \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\sigma^2_b,\sigma^2_w,\phi_v,\phi_n\}}) \sim IG\left(\nu_{\phi_i} + \frac{T-1}{2}, S_{\phi_i} + \frac{1}{2}\sum_{t=2}^{T}(\log(\lambda_{it}) - \log(\lambda_{i,t-1}))^2\right), i = v, n,$$

where the elements of ϵ_z can be computed as $\epsilon_z = H_\psi^{-1}(z - X_d d)$. Next, the log conditional density for σ^2_b is given by

$$\log(\sigma^2_b \,|\, Data, \pi^*, b, d, \lambda_v, \lambda_n, \theta_{-\{\sigma^2_b,\sigma^2_w,\phi_v,\phi_n\}}) \propto -(\nu_{\sigma^2_b} + 1)\log\sigma^2_b - \frac{S_{\sigma^2_b}}{\sigma^2_b} - \frac{T-1}{2}\log\sigma^2_b - \frac{1}{2\sigma^2_b}\sum_{t=2}^{T}(b_t - b_{t-1})^2 + g_b(b, \sigma^2_b),$$

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

$$IG\left(\nu_{\sigma^2_b} + \frac{T-1}{2}, S_{\sigma^2_b} + \frac{1}{2}\sum_{t=2}^{T}(b_t - b_{t-1})^2\right).$$

We adopt the approach in Koop, Leon-Gonzalez, and Strachan (2010) to compute the various dynamic posterior model probabilities reported in the paper. For concreteness, suppose we wish to compute $P(d_{0t} = 0 \,|\, Data)$. Recall that the posterior odds ratio in favor of the restriction $d_{0t} = 0$ can be obtained using the SDDR

$$\mathrm{PO}_t = \frac{p(d_{0t} = 0 \,|\, Data)}{p(d_{0t} = 0)}.$$

Therefore, it suffices to calculate the quantities $p(d_{0t} = 0 \,|\, Data)$ and $p(d_{0t} = 0)$.[14] Koop, Leon-Gonzalez, and Strachan (2010) describe how one can compute both quantities using Monte Carlo methods based on the Kalman filter. We follow a similar approach, but use a direct method based on band matrix routines as in Chan (2015). More specifically, note that the full conditional posterior density $p(d \,|\, Data, \pi^*, b, \lambda_v, \lambda_n, \theta)$ is normal as derived in Step 3 of the sampler. Since all marginals of a jointly normal density are normal, so is $p(d_{0t} \,|\, Data, \pi^*, b, \lambda_v, \lambda_n, \theta)$. Therefore, we can evaluate $p(d_{0t} = 0 \,|\, Data, \pi^*, b, \lambda_v, \lambda_n, \theta)$ exactly. It follows that $p(d_{0t} = 0 \,|\, Data)$ can be estimated using the Monte Carlo average:

$$p(d_{0t} = \hat{0} \,|\, Data) = \frac{1}{R}\sum_{i=1}^{R} p(d_{0t} = 0 \,|\, Data, \pi^{*(i)}, b^{(i)}, \lambda^{(i)}_v, \lambda^{(i)}_n, \theta^{(i)}),$$

[14]The notation $P(d_{0t} = 0|Data)$ should not be confused with $p(d_{0t} = 0|Data)$. The former is the posterior probability associated with the model which imposes the restriction $d_{0t} = 0$. The latter is the probability density function of the unrestricted posterior evaluated at the point $d_{0t} = 0$.

where $(\pi^{*(i)}, b^{(i)}, \lambda_v^{(i)}, \lambda_n^{(i)}, \theta^{(i)}), i = 1, \ldots, R$, are posterior draws. Similarly, we can estimate $p(d_{0t} = 0)$ using Monte Carlo methods since the prior density is also normal. Finally, we can calculate the posterior probabilities using $P(d_{0t} = 0 \,|\, Data) = \mathrm{PO}_t/(1+\mathrm{PO}_t)$.

Dynamic probabilities of the form $P(d_{0t} = d_{0s} \,|\, Data)$ can be computed using the same approach with slight modifications. As before, the joint density of d given the data and parameters is normal. Hence, the bivariate density (d_{0t}, d_{0s}) is normal, so is the linear transformation $d_{0t} - d_{0s}$. Therefore, $P(d_{0t} - d_{0s} = 0 \,|\, Data)$ can be computed as before.

Empirical Appendix

In this appendix, we present additional results for a specification including the unemployment gap in the inflation equation, three other data specifications, empirical evidence on a key parameter in our model, ψ, and a prior sensitivity analysis.

Results for a Model Augmented with Economic Activity

To examine the robustness of our results to including economic activity as a predictor of inflation, we follow common practice (e.g., Morley, Piger, and Rasche 2015, Stella and Stock 2013) and define the relevant activity variable as an unemployment gap, defined as the actual unemployment rate less the Congressional Budget Office's estimate of the natural rate of unemployment.[15] We augment the model to include the unemployment gap, denoted x_t, as follows:

$$
\begin{aligned}
\pi_t - \pi_t^* &= b_t(\pi_{t-1} - \pi_{t-1}^*) + \beta_t x_{t-1} + v_t, && (17)\\
z_t &= d_{0t} + d_{1t}\pi_t^* + \varepsilon_{z,t} + \psi\varepsilon_{z,t-1}, && (18)\\
\pi_t^* &= \pi_{t-1}^* + n_t, && (19)\\
b_t &= b_{t-1} + \varepsilon_{b,t},\ \varepsilon_{b,t} \sim TN(0,\sigma_b^2), && (20)\\
\beta_t &= \beta_{t-1} + \varepsilon_{\beta,t},\ \varepsilon_{\beta,t} \sim N(0,\sigma_\beta^2), && (21)\\
d_{it} - \mu_{di} &= \rho_{di}\left(d_{i,t-1} - \mu_{di}\right) + \varepsilon_{di,t},\ \varepsilon_{di,t} \sim\ N(0,\sigma_{di}^2),\ i = 0,1, && (22)\\
v_t &= \lambda_{v,t}^{0.5}\varepsilon_{v,t},\ \varepsilon_{v,t} \sim N(0,1), && (23)\\
n_t &= \lambda_{n,t}^{0.5}\varepsilon_{n,t},\ \varepsilon_{n,t} \sim N(0,1), && (24)\\
\log(\lambda_{i,t}) &= \log(\lambda_{i,t-1}) + \nu_{i,t},\quad \nu_{i,t} \sim\ N(0,\phi_i),\ i = v,n. && (25)
\end{aligned}
$$

For comparison to our baseline model, we estimate this activity-augmented model using CPI inflation and PTR-CPI inflation expectations. Figures A1-A5 below present key results from the model. Figure A1 (compare to Figure 1) indicates the inclusion of an economic activity measure has little effect on our estimates of trend inflation. Figures A2 through A4 (compare to Figures 2-4) show that the augmented model yields evidence on time variation in inflation persistence, volatility, and the coefficients of the expectations equation very similar to the evidence from the baseline model. Finally, Figure A5 reveals substantial evidence of a gradual, but sizable, trend in the coefficient on the unemployment gap in the inflation equation. Our estimates point to a Phillips curve that has become flatter (with inflation becoming less sensitive to economic activity) over time.

Results Using Core CPI Inflation and PTR-CPI Forecasts, 1961-2014

Figures A6 through A12 are the same as Figures 1 through 7 except that they use core CPI inflation. When using our model, results are very similar to those with CPI inflation.

[15] Following studies such as Rudd and Peneva (2015), we use the measure the CBO refers to as its short-term estimate of the natural rate, which incorporates a temporary, substantial rise in the natural rate in the period following the start of the Great Recession, attributable to structural factors such as extended unemployment insurance benefits.

There are a few minor differences (e.g. there is less evidence in favor of time-variation in d_{1t} and the Bayes factor favoring the model which includes the MA coefficient is, at 1.4, slightly lower than with CPI). The major difference is found in Figure A12 for the UCSV-AR model. With core inflation, the UCSV-AR model is producing very erratic, unreasonable estimates of trend inflation, especially prior to 1990. This is probably due to our use of a relatively non-informative prior. With our approach, the extra information in z_t allows for sensible estimation of trend inflation, even without much prior information. With the UCSV-AR model, additional prior information appears to be required to give sensible estimates.

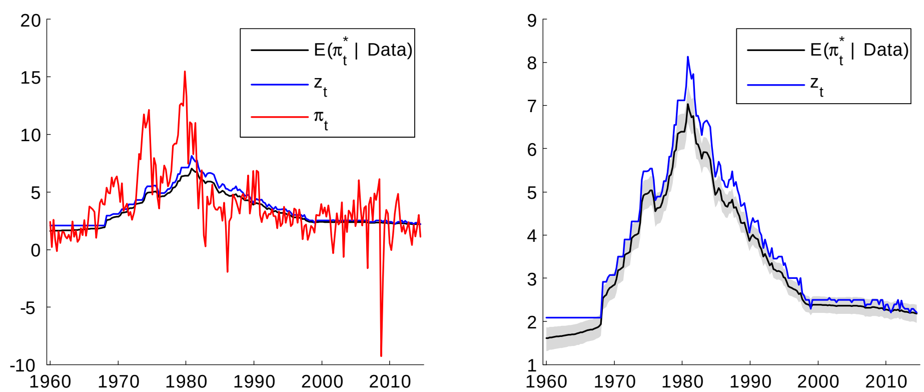

Figure A1: Posterior means and quantiles (16% and 84%) of π_t^*.

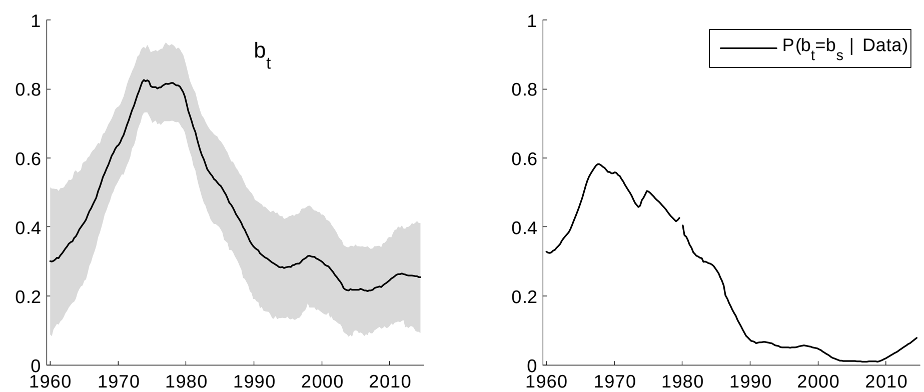

Figure A2: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = 1980Q1$.

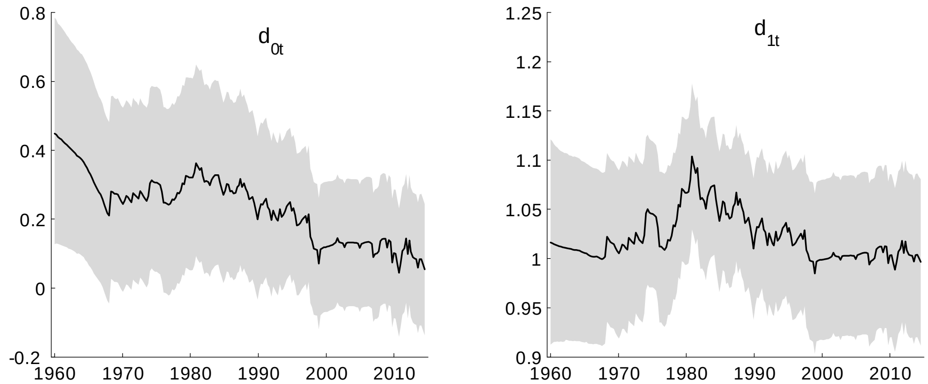

Figure A3: Posterior means and quantiles (16% and 84%) of d_{it}.

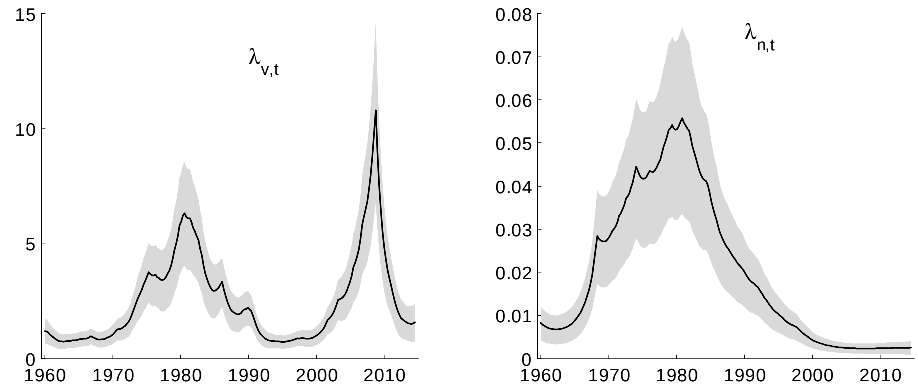

Figure A4: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

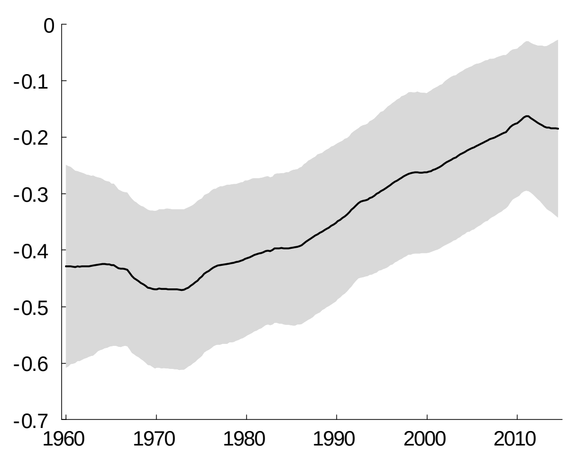

Figure A5: Posterior means and quantiles (16% and 84%) of β_t, the coefficient associated with the unemployment gap.

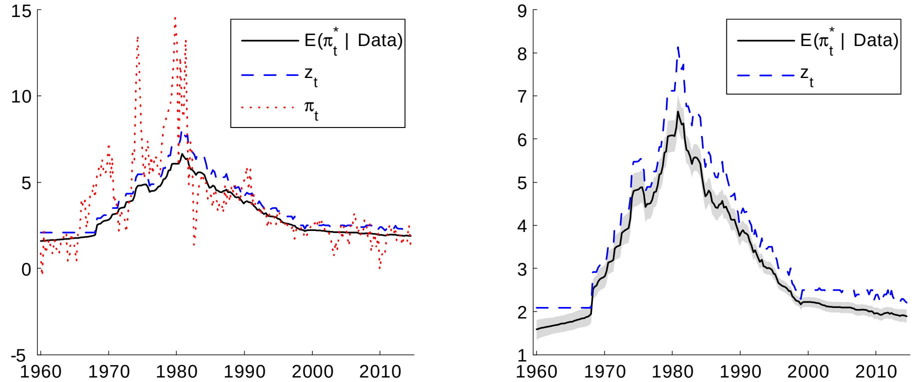

Figure A6: Posterior means and quantiles (16% and 84%) of π_t^*.

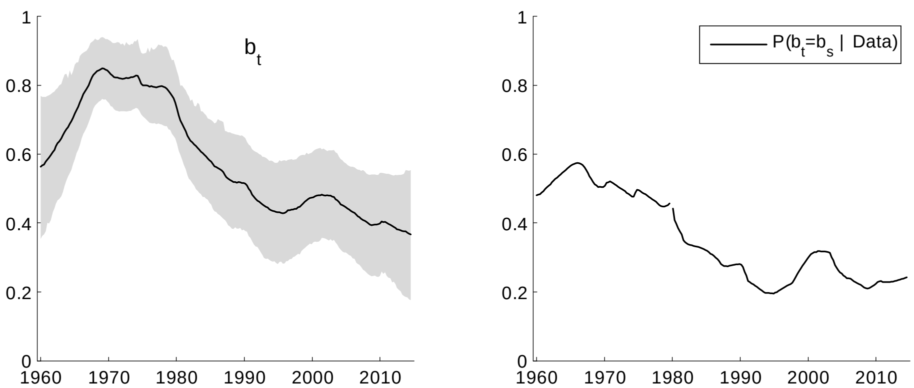

Figure A7: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = 1980Q1$.

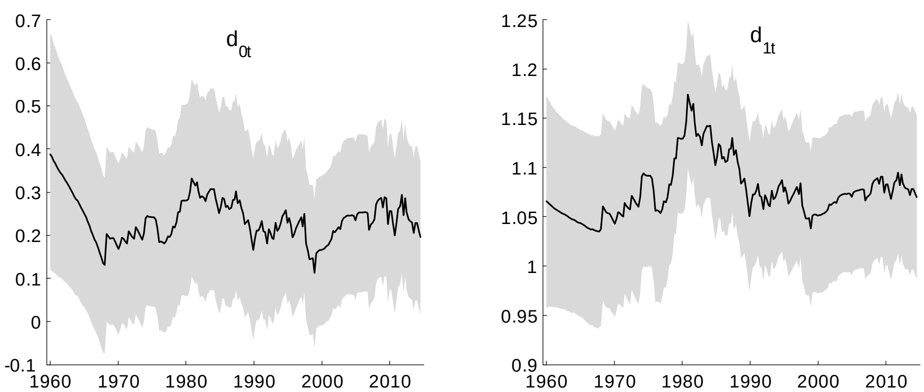

Figure A8: Posterior means and quantiles (16% and 84%) of d_{it}.

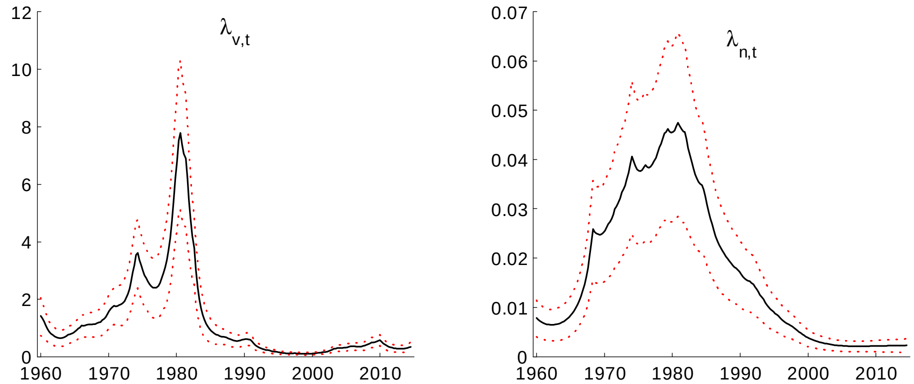

Figure A9: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

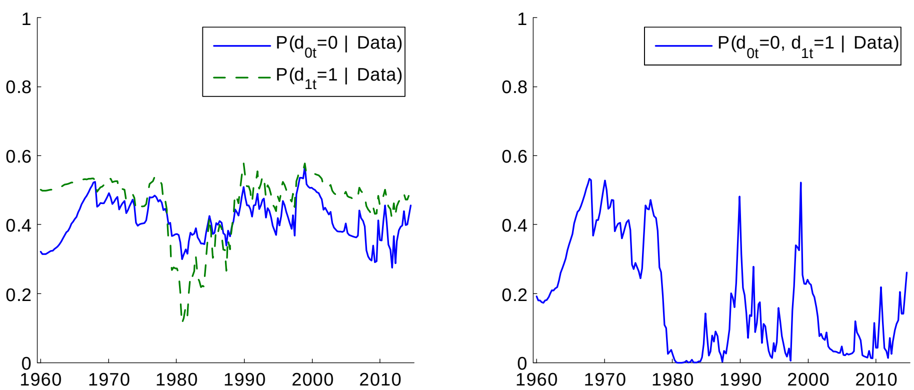

Figure A10: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

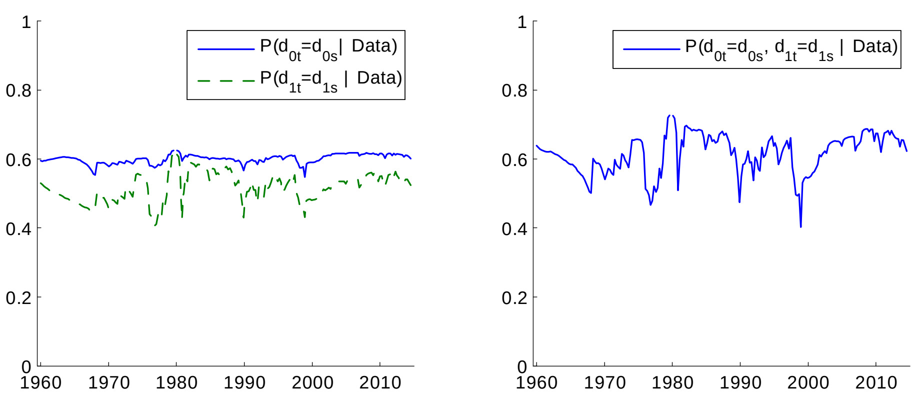

Figure A11: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

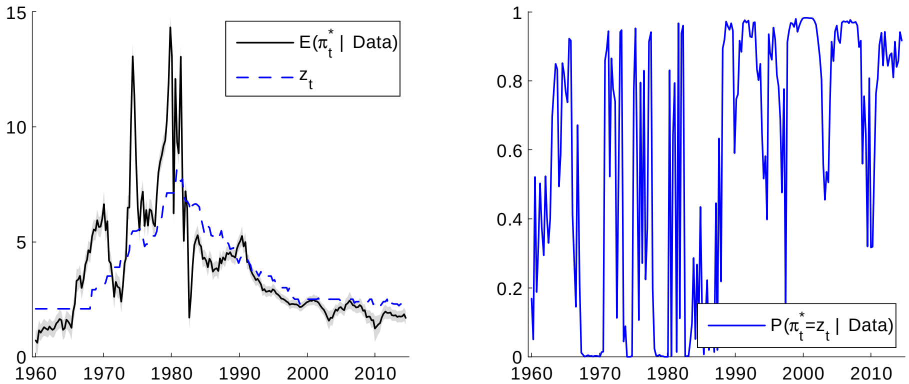

Figure A12: Posterior means and quantiles (16% and 84%) of π_t^* and dynamic probabilities $P(\pi_t^* = z_t|Data)$ for UCSV-AR model

Results Using Core CPI inflation and Blue Chip CPI Forecasts, 1979-2014

Figures A18 through A24 are based on core CPI inflation and Blue Chip forecasts of CPI inflation. There is strong evidence in favor of an MA term in the inflation expectations equation. The Bayes factor against the model with $\psi = 0$ is 29.8. Our estimate of trend inflation is consistently slightly below the suvey-based long-run inflation forecasts, but, other than this, the two series move closely together. This contrasts with the UCSV-AR estimates of trend inflation in Figure A19 which are much more erratic and do not move as closely with the long-run inflation forecasts.

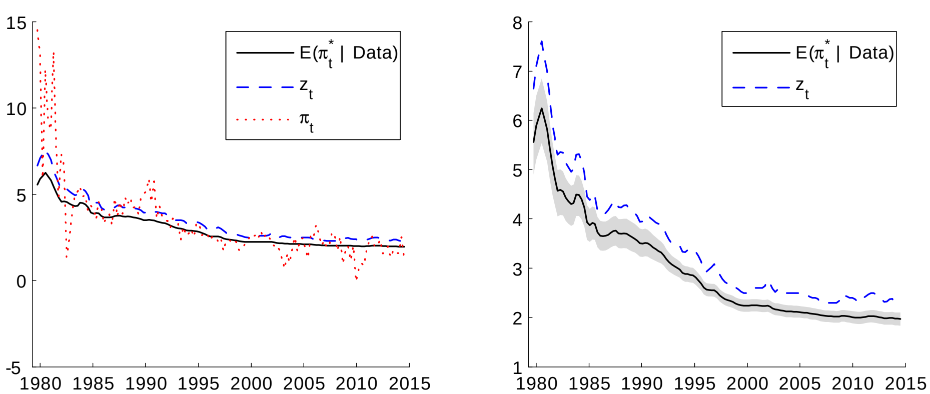

Figure A13: Posterior means and quantiles (16% and 84%) of π_t^*.

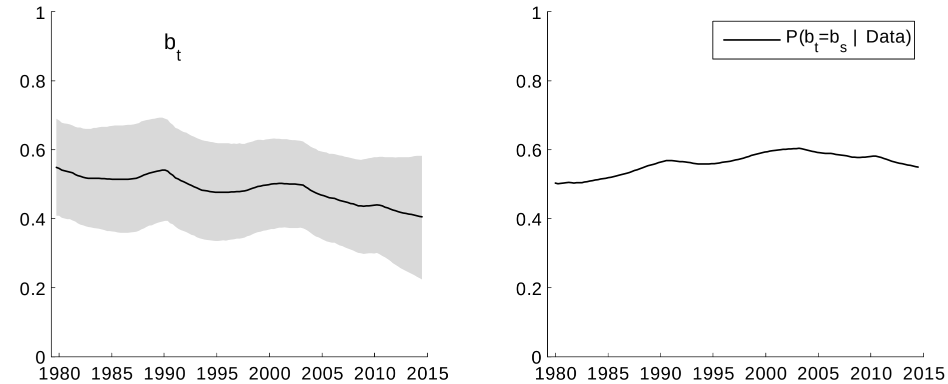

Figure A14: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = 1980Q1$.

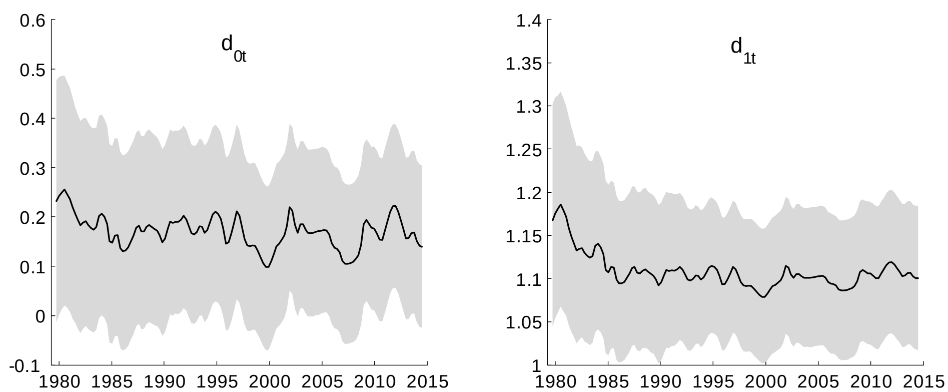

Figure A15: Posterior means and quantiles (16% and 84%) of d_{it}.

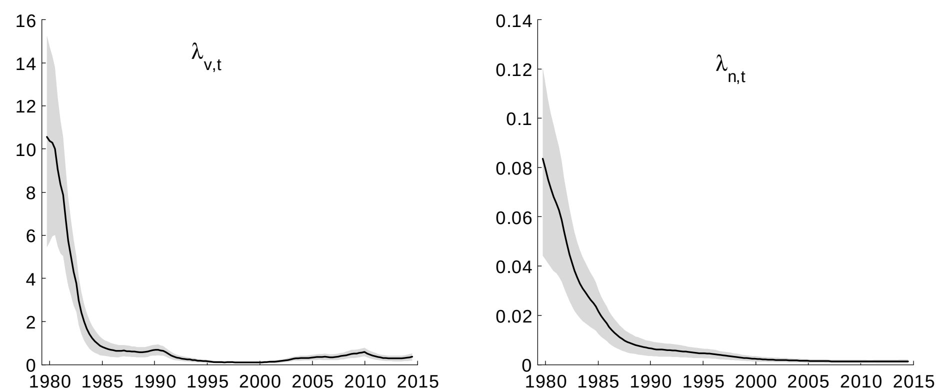

Figure A16: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

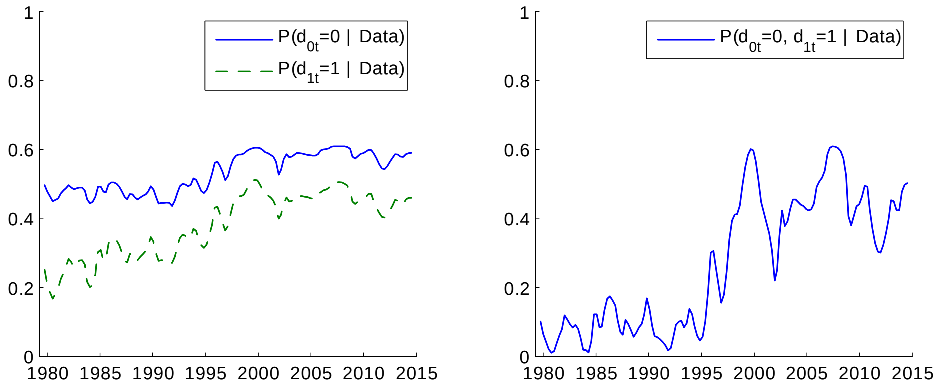

Figure A17: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

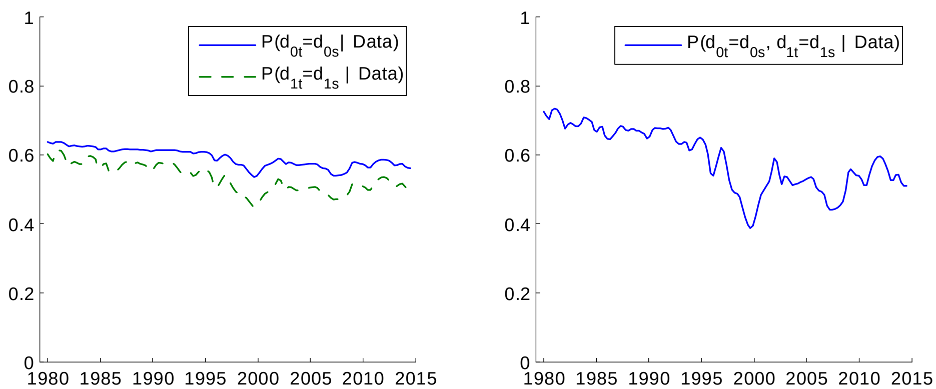

Figure A18: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

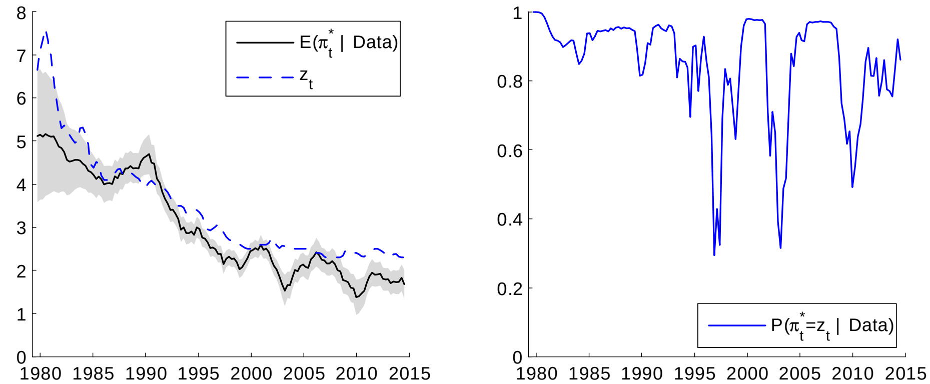

Figure A19: Posterior means and quantiles (16% and 84%) of π_t^* and dynamic probabilities $P(\pi_t^* = z_t|Data)$ for UCSV-AR model.

Results Using GDP Deflator inflation and Blue Chip GDP Deflator Forecasts, 1979-2014

In this sub-section, we again find strong evidence in favor of including the MA process in the equation for z_t with the Bayes factor in its favor being 37.9. In general, we are finding the same pattern as in previous sections where our model is producing sensible, smooth, estimates of trend inflation which match up closely, but not perfectly with long run forecasts. This reinforces our story that including survey-based inflation forecasts can improve estimation of trend inflation, but this should be done in a model based fashion rather than simply equation long run forecasts with trend inflation.

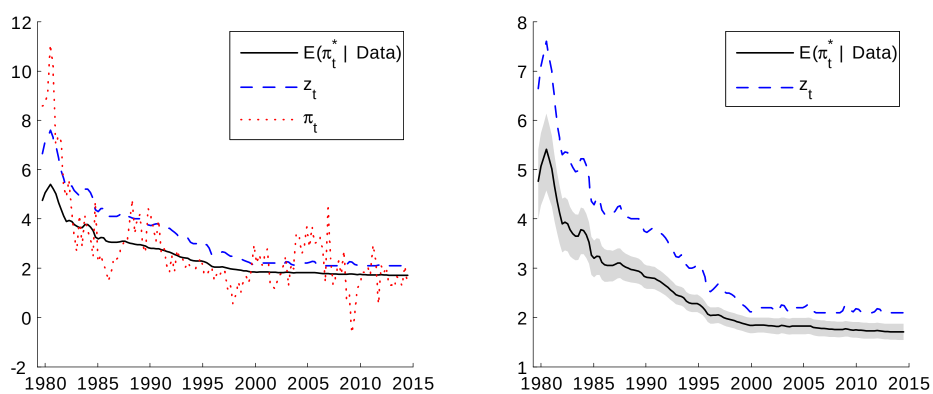

Figure A20: Posterior means and quantiles (16% and 84%) of π_t^*.

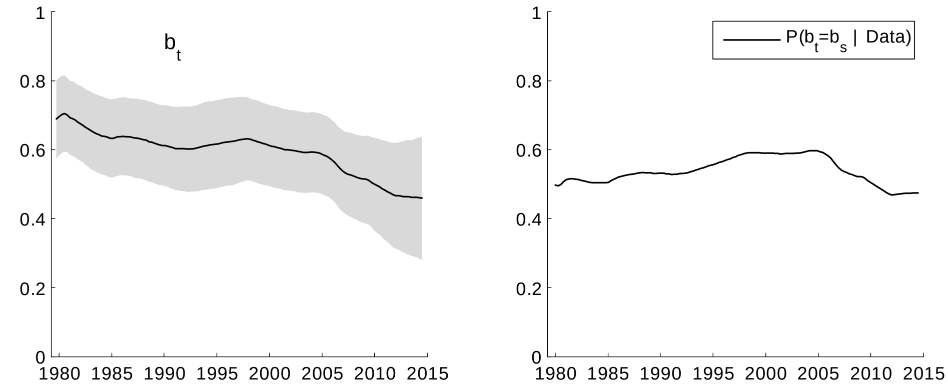

Figure A21: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = 1980Q1$.

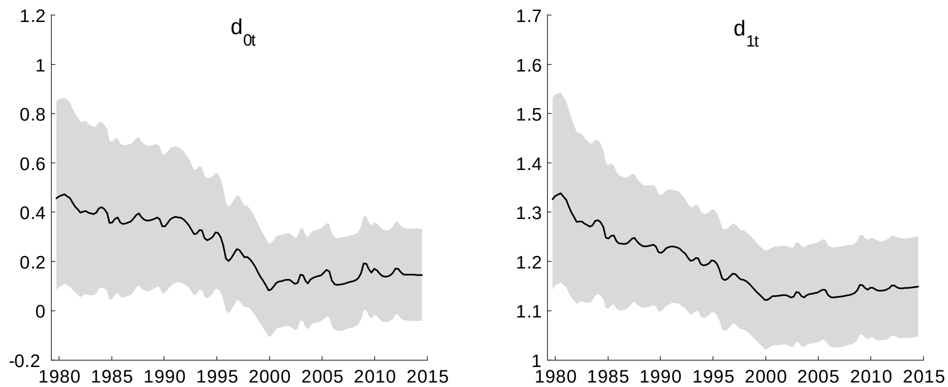

Figure A22: Posterior means and quantiles (16% and 84%) of d_{it}.

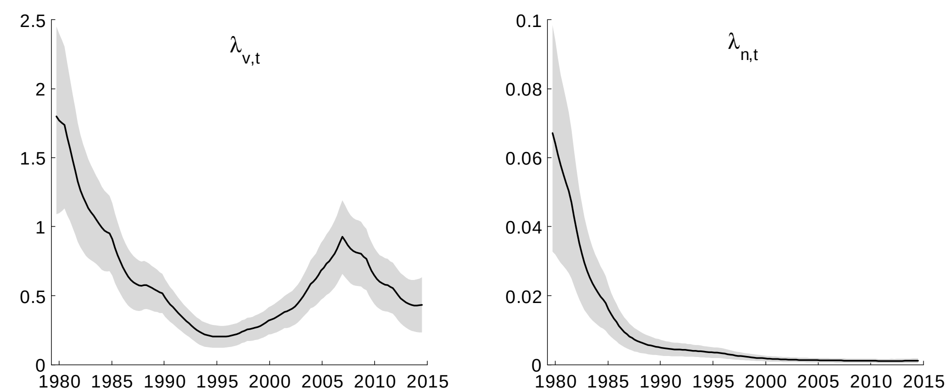

Figure A23: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

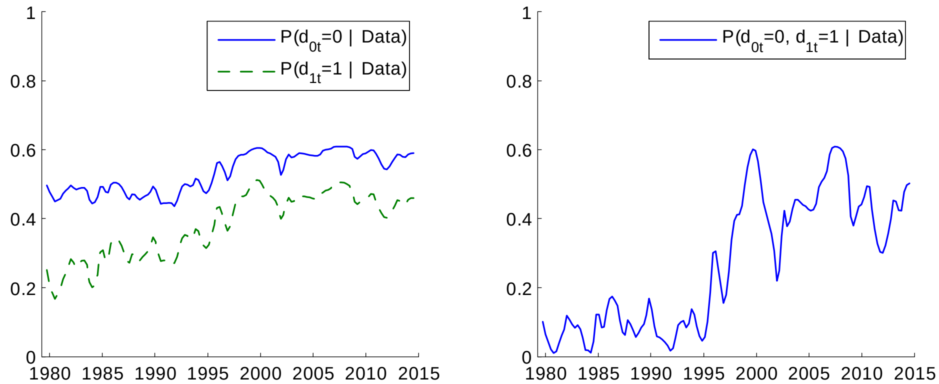

Figure A24: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

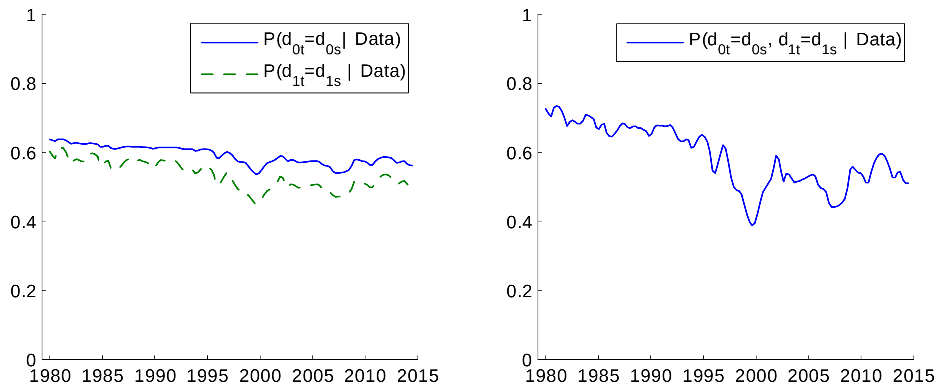

Figure A25: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

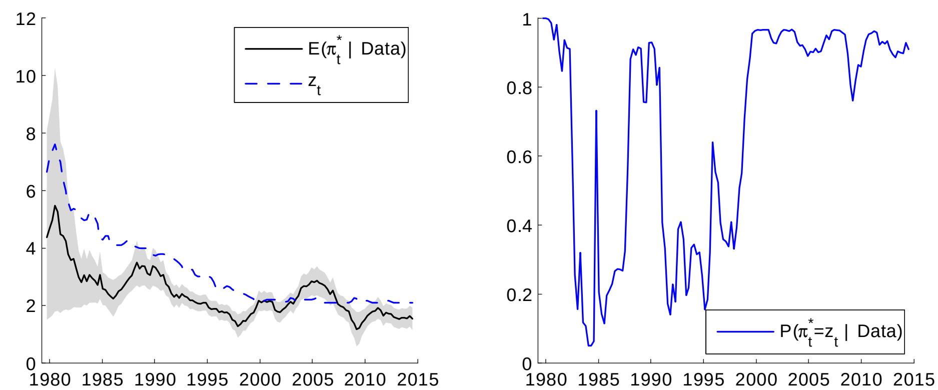

Figure A26: Posterior means and quantiles (16% and 84%) of π_t^* and dynamic probabilities $P(\pi_t^* = z_t|Data)$ for UCSV-AR model.

Moving average component of model

Figure A27 presents the priors and posteriors for our six different data configurations (i.e. a combination of an inflation measure with a long-run inflation forecast). These were used to calculate the Bayes factors presented in the body of the paper using the SDDR. It can be seen that all of them suggest a positive MA coefficient, but it is only when using Blue Chip forecasts is it the case this evidence is strong.

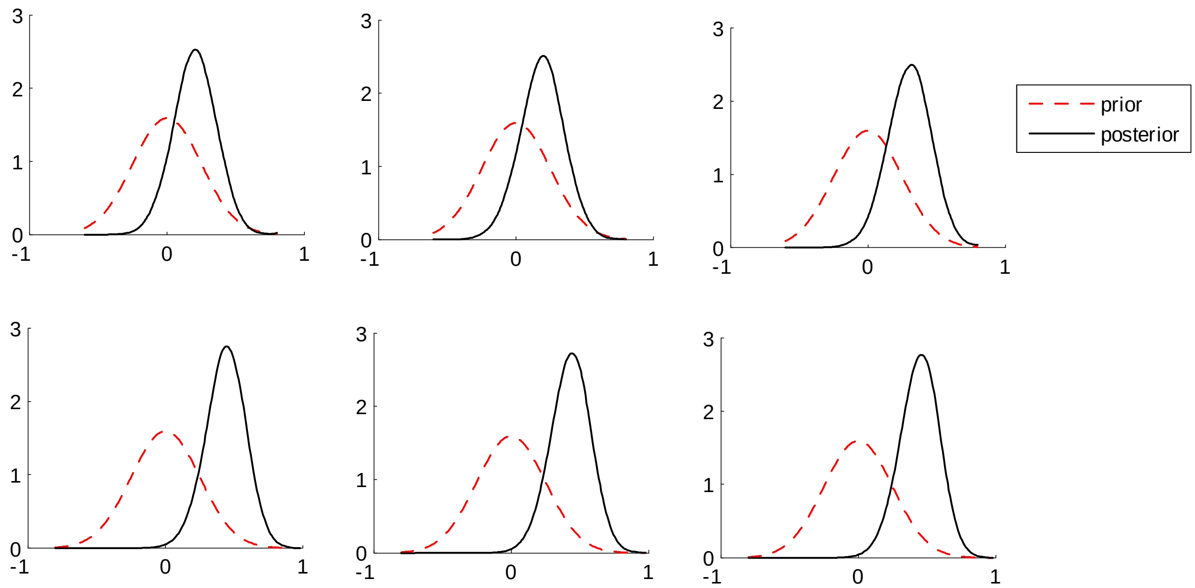

Figure A27: Prior and posterior densities of ψ for six combinations of inflation and inflation expectation: CPI and PTR-CPI (top left), core CPI and PTR-CPI (top middle), PCE and PTR (top right), CPI and Blue Chip CPI (bottom left), core CPI and Blue Chip CPI (bottom middle) and GDP deflator and Blue Chip GDP deflator (bottom right).

Sensitivity Analysis

There are many parameters in this model and, thus, many dimensions we could carry out a prior sensitivity analysis. However, many parameters are initial conditions for state equations. Empirically, researchers often use diffuse initial conditions and we have followed this practice. Hence, we do not present results relating to prior sensitivity relating to these parameters. Similarly for several parameters (e.g. the MA coefficient) we are finding sensible parameter estimates using relatively non-informative priors and, hence, do not investigate prior sensitivity to them. Instead we focus on the prior hyperparameter V_μ (remember that $\mu_{d0} \sim N(a_0, V_\mu), \quad \mu_{d1} \sim N(a_1, V_\mu)$) as this is a key parameter in the state equation for d_{it}. Note that the term $d_{1t}\pi_t^*$, involving two latent series, appears in (5). In order to allay concerns about separately identifying these two series, we present results for different choices of V_μ. Results in the paper are for $V_\mu = 0.1^2$. This is an informative, but not dogmatic choice, attaching appreciable prior weight to intervals of ± 0.2 around the theoretically-justified prior mean values. In this appendix, we present results for the more informative choice of $V_\mu = 0.025^2$ and the very non-informative choice of $V_\mu = 1$. We use CPI inflation and PTR-CPI long run forecasts.

Figures A28 through A33 present results using $V_\mu = 0.025^2$. Overall results are very similar to Figures 1 through 7. Since the prior hyperparameters we are changing are for the state equation for d_{it}, it is unsurprising that the largest impacts are seen by comparing Figure 6 to Figure A32. It can be seen that $P(d_{0t} = 0, d_{1t} = 1|Data)$ and $P(d_{1t} = 1|Data)$ are slightly higher using the tighter prior, but the difference is small.

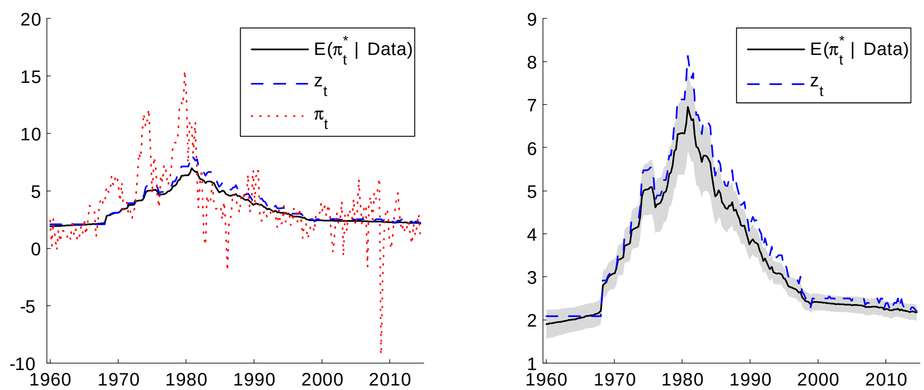

Figure A28: Posterior means and quantiles (16% and 84%) of π_t^*.

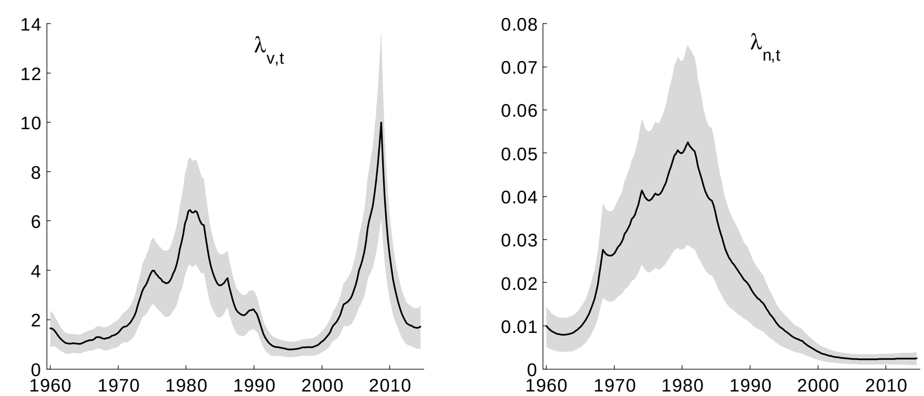

Figure A29: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

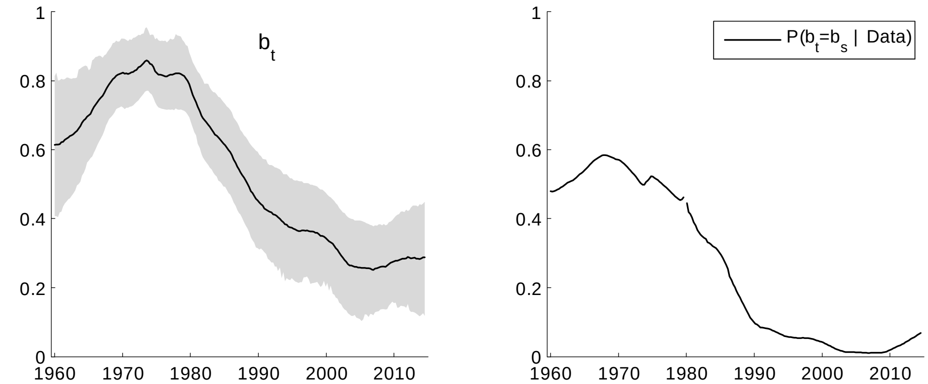

Figure A30: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = 1980Q1$.

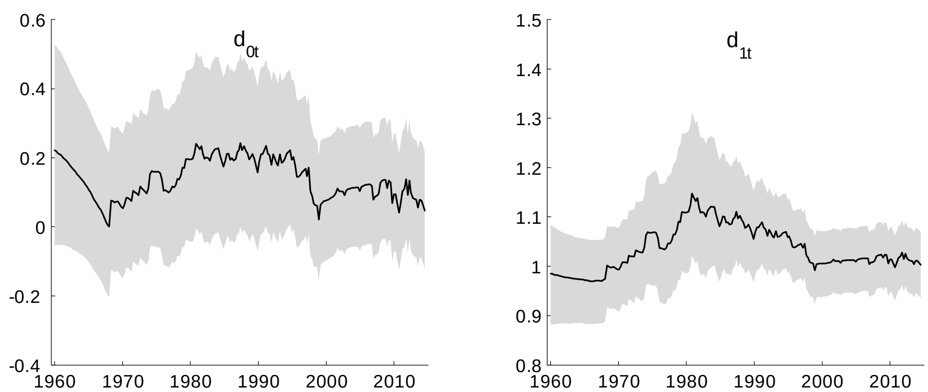

Figure A31: Posterior means and quantiles (16% and 84%) of d_{it}.

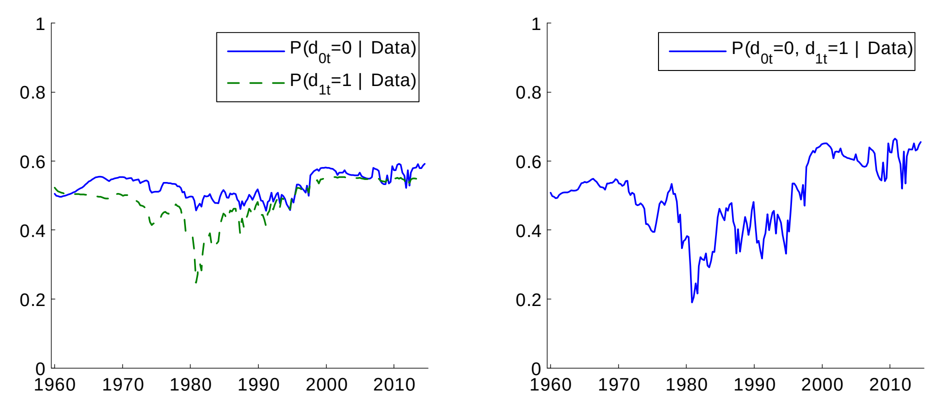

Figure A32: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

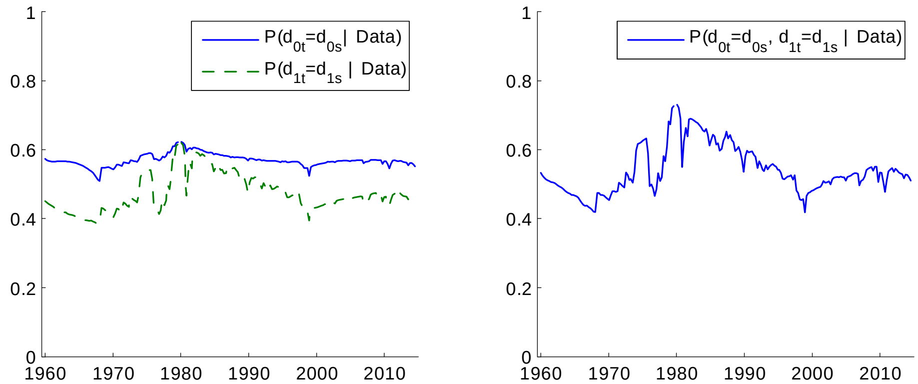

Figure A33: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.

Figures A34 through A39 present results for the very non-informative case with $V_\mu = 1$. Estimates of trend inflation are not greatly affected by this, but estimates of d_{0t} and d_{1t} are. In particular, they are now far away from theoretically-suggested values of $d_{0t} = 0$ and $d_{1t} = 1$. This reinforces a point made previously. That is, our model has the desirable feature that it does not dogmatically impose $d_{0t} = 0$ and $d_{1t} = 1$ a priori as other approaches implicitly do. However, given the great flexibility of our model, we do require some prior information in order to get reasonable results.

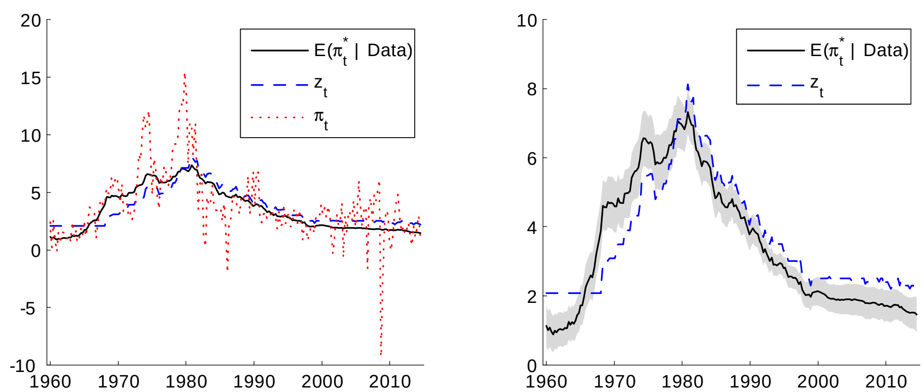

Figure A34: Posterior means and quantiles (16% and 84%) of π_t^*.

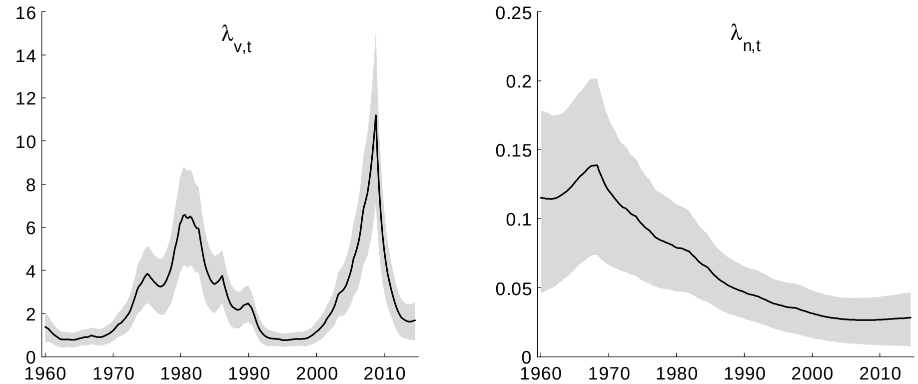

Figure A35: Posterior means and quantiles (16% and 84%) of $\lambda_{v,t}$ and $\lambda_{n,t}$.

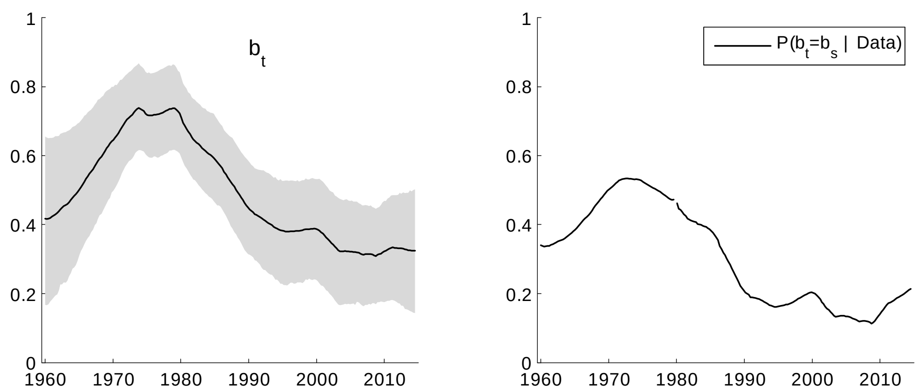

Figure A36: Posterior means and quantiles (16% and 84%) of b_t^*, and the dynamic probabilities that $b_t = b_s$ with $s = 1980\text{Q}1$.

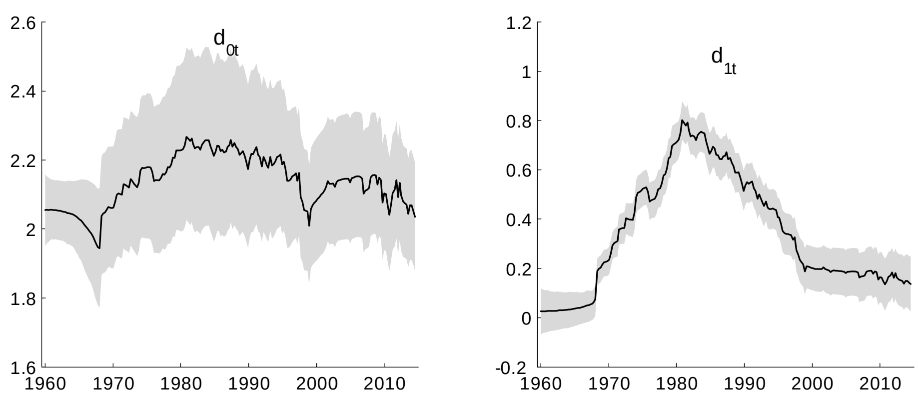

Figure A37: Posterior means and quantiles (16% and 84%) of d_{it}.

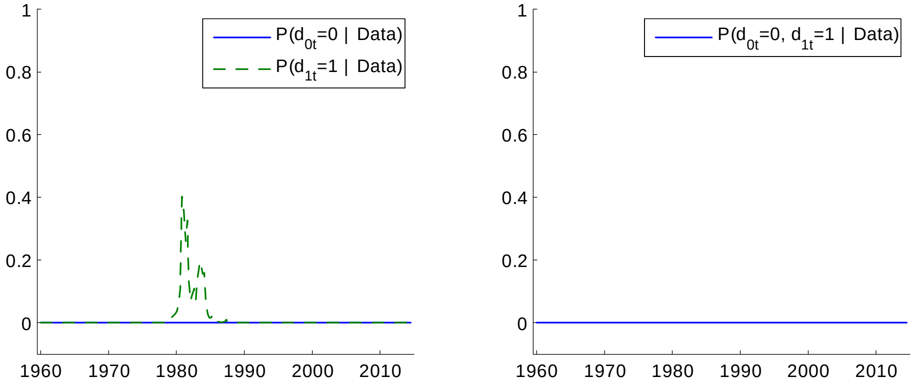

Figure A38: Marginal and joint dynamic probabilities for d_{0t} and d_{1t}.

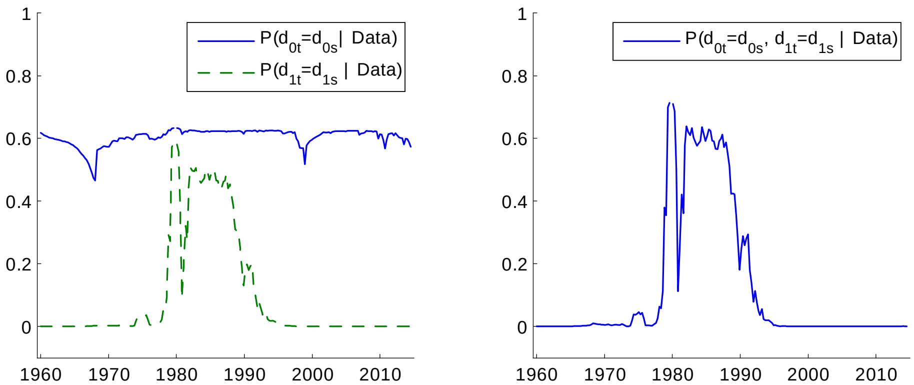

Figure A39: The marginal and joint dynamic probabilities that $d_{it} = d_{i,s}$ with $s =$ 1980Q1.