Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

Regime-Switching Models for Estimating Inflation Uncertainty

Jeremy J. Nalewaik

2015-093

Please cite this paper as:
Nalewaik, Jeremy J. (2015). "Regime-Switching Models for Estimating Inflation Uncertainty," Finance and Economics Discussion Series 2015-093. Washington: Board of Governors of the Federal Reserve System, http://dx.doi.org/10.17016/FEDS.2015.093.

Regime-Switching Models for Estimating Inflation Uncertainty

Jeremy Nalewaik*
Federal Reserve Board

September 2015

Abstract

This paper constructs regime-switching models for estimating the probability of inflation returning to its relatively high levels of variability and persistence in the 1970s and 1980s. Forecasts and probabilities of extreme events from the models are evaluated against comparable estimates from other statistical models, from surveys, and from financial markets. The paper then uses the models to construct prediction intervals around Federal Reserve Board staff forecasts of PCE price inflation, combining the recent non-parametric forecast error distribution with parametric information from the model. The outer tails of the prediction intervals depend importantly on the probability inflation is in its high-variance, high-persistence regime.

*Board of Governors of the Federal Reserve System, 20th Street and Constitution Avenue NW, Stop 80, Washington, DC 20551 (e-mail: jeremy.j.nalewaik@frb.gov). This work was completed on sabbatical at the Johns Hopkins University. Thanks to Todd Clark, Evan Koenig, David Lebow and Simon Potter for comments. The views expressed in this paper are those of the author and do not necessarily reflect the views of the Board of Governors of the Federal Reserve System or the rest of its staff.

1 Introduction

The decline in inflation variability and persistence in recent decades—see Stock and Watson (2007)—has important implications for constructing measures of forecast uncertainty. Taking on board that decline in variability and persistence is reasonable when estimating uncertainty around forecasts that are conditioned on such an environment continuing over the forecast horizon. That might be the case for the forecasts that members of the Federal Open Market Committee (FOMC) submit in advance of their meetings once per quarter, since those forecasts are conditioned on "appropriate monetary policy".[1]

However, constructing measures of *unconditional* forecast uncertainty requires estimates of the probability that the conditioning assumptions break down over the forecast horizon, and the distribution from which the data would be generated if the conditioning assumptions do break down. This paper constructs models that provide such estimates, Markov-switching models where inflation switches from a low-variance regime with a stable mean to a high-variance, random-walk regime. Probabilities from the model show the second regime governed the behavior of inflation from the late 1960s to the early 1990s, while the first regime has governed the behavior of inflation from the early 1990s to the end of the sample used in the paper.

The small number of regime transitions in the sample poses some challenges, discussed at length in the paper. Model fit is evaluated by comparing forecasts and probabilities of extreme events from the models with an array of comparable estimates, from other statistical models to surveys to financial markets. For example, the recent low-variance, low-persistence

[1]Information about these forecasts is released to the public after every second FOMC meeting in a Summary of Economic Projections (SEP). For more information, see: http://www.federalreserve.gov/monetarypolicy/fomccalendars.htm.

inflation regime is often associated with the relative stability of survey measures of long-run inflation expectations (see Bernanke (2010)), so the paper compares those survey measures with long-run inflation forecasts from the Markov-switching model. In the 1990s, the model recognized the inflation regime transition and its implications for long-run forecasting years ahead of the surveys, suggesting that proper use of the model may provide inferences about the inflation regime that are more timely and accurate than those based on surveys. It goes without saying that accurate inferences about the inflation regime are critically important for the conduct of monetary policy.

Overall, model fit seems adequate to proceed with construction of prediction intervals estimating the uncertainty around Federal Reserve Board staff forecasts of PCE price inflation. In this case, the usable history of staff forecasts does not extend back far enough to include any observations from the period of high inflation variability and persistence, so estimation of a model is necessary to account for uncertainty arising from a possible reversion to that type of behavior. However, even if a longer sample including such observations were available, the standard approach to computing prediction intervals, which gives equal weight to each observation in the sample, is inadequate when there is predictable variation in the regime probability.[2] Such an approach abdicates responsibility for estimating what could be the most important unobserved variable in the process, the regime probability.

While various time-varying parameter models have been used to analyze the inflation process, the Markov-switching approach here is appealing because forecast error distributions computed non-parametrically can be combined easily in a mixture with model-based forecast error distributions that are, out of necessity, parametric. For the low-variance regime with a

[2]For example, if half the observations in the sample are from one regime and half are from the other, such an approach implicitly assumes a 50 percent probability of each regime in the relevant forecast period, possibly far off the mark in a predictable direction.

stable mean, the paper recommends the non-parametric approach, where forecast error percentiles are estimated directly from the percentiles of the sorted sample of available forecast error observations. Such an approach allows asymmetries and other non-standard features of the sample show through to prediction intervals, and the sample of Board staff forecast errors does show marked asymmetries. In particular, the forecasts at one end of the forecast error distribution—those that were too low—missed by much more than the forecasts at the other end of the distribution that were too high. The paper discusses possible reasons for this asymmetry.

The unconditional forecast error distribution is a mixture distribution, a linear combination of the non-parametric distribution in the low-variance regime and a parametric distribution in the high-variance regime, using the estimated regime probabilities as the weights on these two distributions—see Hamilton (1989). Prediction intervals computed from this mixture do not assume that inflation remains in the low-variance regime with certainty over the forecast horizon, but when the probability of the low-variance regime is very high as in recent years, the prediction intervals around one- and two-year ahead forecasts are well approximated by those computed from the low-variance regime distribution. A sizable increase in the probability of the high-variance regime would change that, increasing the mass in the tails of the distribution noticeably. Since the process in the high-variance regime contains a unit root, the width of prediction intervals would then widen more with the forecast horizon, and the location of the intervals would depends more on the current level of inflation.

Several papers estimate regime-switching models of the inflation process, such as Kim (1993), Evans and Wachtel (1993), Lanne (2006), and Davig and Doh (2014). Those papers also find time-varying persistence and uncertainty in the quarterly or monthly inflation process. This paper uses annual observations, matching the frequency of the forecasts submitted by FOMC members ahead of each meeting. Annual observations also have the benefit of averaging away some of the higher-frequency noise in the inflation process, focusing on fluctuations that are more economically interesting. The closest antecedent to this paper is Evans and Wachtel (1993), who showed that regime uncertainty is an important part of overall inflation uncertainty, a result mirrored here.

2 Data and Preliminary Analysis

The paper studies fourth-quarter over fourth-quarter (Q4/Q4) growth rates of the consumer price index (CPI) and the personal consumption expenditures (PCE) price index from 1947-2014, the longest sample available. The two inflation measures are plotted in Figure 1. The formal inflation target adopted by the FOMC in early 2012 is stated in terms of PCE price inflation, and the aforementioned forecasts of FOMC members target Q4/Q4 PCE price inflation. CPI inflation has been more widely followed traditionally, and is used in pricing inflation-indexed government bonds and other financial instruments—see section 4.2.

Most recent academic research on inflation, exemplified by Stock and Watson (2007) and Cogley, Primiceri and Sargent (2010) studies quarterly inflation time series, modelling the process as the sum of permanent and transitory components. While model parameters vary over time, at any point in time, the Stock and Watson (2007) model implies inflation has a unit root with the first difference of inflation following an MA(1) process. However, the autocorrelations and partial autocorrelations of the annual inflation measures studied here strongly favor adding a second moving average parameter. The top panel of table 1 reports maximum likelihood parameter estimates and standard errors for MA(2) processes estimated on the first differences of the two inflation measures over the full 1948-2014 sample. The t-statistics on θ_2, the second moving average parameter, are quite large, between 4 and 5.

Since $\theta_2 \approx -0.5$, about half of an innovation to annual inflation reverses two years later, on average. Examining the top 20 innovations to PCE price inflation in the sample, not all showed substantial reversals after two years; those that did not include the positive innovations in 1966 and 1968 and the negative innovations in 1991 and from 1981-1984. But many large innovations did show substantial reversals. For example, negative innovations to inflation in 1949 and 1954, likely caused by recessions, reversed after expansions resumed.[3] Conversely, positive innovations to inflation in 1973-1974, 1979-1980, and 2007 reversed after the economy fell into recession. While most of these reversals are tied to expansion-recession dynamics, not all of them are: negative innovations to inflation in 1986 and 1997-1998, in part due to oil price declines, reversed in 1987-1988 and in 1999-2000.

A number of factors likely contributed to the reversals in these innovations since 1973, which were all tied to substantial swings in oil and other commodity prices. First, market-specific demand and supply responses to the commodity price shocks may have contributed to the reversals, with consumer and producer responses likely becoming more elastic as time passes after the shock and scope for adjustments increases. Second, for the oil price shocks in particular since the U.S. is a net importer of oil, aggregate economic activity ultimately moves in the opposite direction of the price innovation, driven by the effect of the price change on consumption through household purchasing power and consumer confidence.[4] These movements in aggregate economic activity are sometimes reinforced by monetary policy reactions to inflation, as in the mid-1970s and early 1980s. However, these forces take

[3]The rebound in 1950 was unusually quick, however, probably due to the onset of the Korean War.

[4]See Kilian and Vigfusson (2013), for example, and Bachmann, Berg, and Sims (2015), who show changes in gasoline prices are negatively correlated with readiness to spend on durables.

time to work through to economic activity, which in turn takes time to feed back through to inflation and reverse part of the initial innovation.[5] Based on the large negative θ_2, it appears that these processes tend to occur over a two-year period, on average, over the full sample.

The time variation in the inflation process is illustrated by the second panel of Table 1, which shows MA(2) estimates over the last 20 years of the sample. The innovation variances are much smaller over this subsample and $\theta_1 \approx -0.8$, so inflation innovations have tended to mostly reverse after a year. Indeed, each $\theta_1 + \theta_2$ is close to minus one, consistent with stationarity. Table 2 shows AR(1) estimates on the level of inflation over the same subsample, and while standard errors are large, the estimates are consistent with little to no serial correlation in the inflation process over the last two decades. These in-sample time series specifications are the first category of evidence that motivates the decision in the next section to model the inflation process as switching between stationary and non-stationary regimes.

The second category of evidence suggesting the inflation process has been stationary recently is based on out-of-sample forecast errors. Table 3 reports root mean square forecast errors (RMSEs) using 1995 to 2014 October Blue Chip long-range consensus forecasts of annual average CPI inflation for the current year and from one to four years ahead. If annual CPI inflation contained a non-stationary component, these RMSEs should grow with the forecast horizon.[6] Although one-year ahead inflation appears somewhat predictable with the one-year ahead RMSE less than the two-year ahead RMSE, this likely stems from

[5]These inflation reversals working through economic activity could occur in the oil and commodity prices themselves, but also through other prices in the consumption bundle.

[6]For example, assume for the moment that inflation could be decomposed into a permanent component

a number of factors unrelated to whether there is a random walk component to inflation, including the mechanical effect of known monthly growth rates in the second half of year T entering directly into the year $T+1$ annual average growth rate calculation. The test of the existence of a random walk component is whether the RMSEs continue to increase with the forecast horizon, which they do not as the horizon lengthens from two to four years ahead. These RMSEs are consistent with a stationary inflation process over this time period.[7]

The third category of evidence favoring stationarity in the U.S. inflation process in recent years is the extensive cross-country historical analysis in Benati (2008) relating the persistence of inflation to the monetary policy regime. That paper finds that inflation-targeting regimes tend to be associated with zero serial correlation in the inflation process, and since the U.S. has adopted such an inflation-targeting regime formally in early 2012 and gradually in a number of steps before then, this evidence supports not only stationarity but also zero serial correlation over the most recent part of the sample. Furthermore, assuming this inflation-targeting regime persists, as currently appears likely, allowing for a stationary regime in a time-varying parameter model of inflation seems likely to become all the more important going forward.

Given all these empirical results over the last 20 years favoring stationarity, it is a bit

---

and a transitory component, as in Stock and Watson (2007) excluding parameter variation for simplicity, so:

$$\begin{aligned} \pi_t &= \tau_t + \nu_t \quad \mathrm{var}(\nu_t) = \sigma_\nu^2 \\ \tau_t &= \tau_{t-1} + \varepsilon_t \quad \mathrm{var}(\varepsilon_t) = \sigma_\varepsilon^2. \end{aligned}$$

Then the variance of the forecast error $\mathrm{var}(\pi_{t+k} - E(\pi_{t+k}|\pi_t)) = \sigma_\nu^2 + k\sigma_\varepsilon^2$ increases linearly with k because the permanent innovations cumulate with the length of the forcast horizon.

[7]For similar results using average RMSEs from a range of private and government sector forecasters, see the RMSE tables released with each FOMC SEP; for example: http://www.federalreserve.gov/monetarypolicy/files/fomcminutes20141217.pdf. For background, see Reifschneider and Tulip (2007).

surprising that the Atkeson-Ohanian (2001) random walk model for four-quarter inflation remains a benchmark in many forecast evaluations. A better benchmark over the past 20 years is its polar opposite, a model where inflation has zero rather than full persistence.

3 Model Estimation and Results

3.1 Estimation and Current-Period Probabilities

Inflation follows a two regime or two state ($S_t = 1$ or $S_t = 2$) Markov-switching process where:

$$\begin{aligned} \pi_t | S_t = 1 \qquad & \sim N\left(\mu_{S1}, \sigma_{s1}^2\right) \\ \pi_t | S_t = 2, \pi_{t-1}, \varepsilon_{t-1}, \varepsilon_{t-2} \quad & \sim N\left(\pi_{t-1} + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2}, \sigma_{s2}^2\right). \end{aligned}$$

Let $\mathcal{H}_{t-1} = \{1, \pi_{t-1}, \pi_{t-2}, \ldots, \pi_1, \varepsilon_0 = 0, \varepsilon_{-1} = 0\}$ denote the history of the inflation, and:

$$\begin{aligned} \operatorname{prob}\left(S_{t-1} = 1 \quad | \quad \mathcal{H}_{t-1}\right) &= p_{t-1|t-1}, \qquad \text{and:} \\ \operatorname{prob}\left(S_{t-1} = 2 \quad | \quad \mathcal{H}_{t-1}\right) &= 1 - p_{t-1|t-1}. \end{aligned}$$

The probabilities and likelihood function, a weighted average of the state-contingent likelihood functions with these probabilities as weights, is computed as in Kim (1994), with the probability of each state persisting from one year to the next governed by the transition matrix:

$$\begin{pmatrix} p_{t+1|t} \\ 1 - p_{t+1|t} \end{pmatrix} = \begin{bmatrix} p_{11} & 1 - p_{22} \\ 1 - p_{11} & p_{22} \end{bmatrix} \begin{pmatrix} p_{t|t} \\ 1 - p_{t|t} \end{pmatrix}.$$

The initial probability is treated as an estimated parameter.

Results from estimating the model on each of the two inflation measures from 1948-2014 are reported in Table 4, with one- and two-sided (smoothed) probabilities plotted in Figures 2 and 3. The innovation standard deviations in the random walk state 2 are between two and three times as large as the innovation standard deviations in the stationary state, so it is natural to label state 2 the high-variance, high-persistence state. The two inflation measures produce very similar smoothed probabilities, the most notable difference being that the CPI probabilities show a bit less certainty that inflation exited the high-variance, high-persistence state in the 1950s. Both smoothed probability series show inflation likely returned to the high-variance, high-persistence state in 1967, where it remained until 1991. μ_{S1} is determined by mean inflation from the mid-1950s to mid-1960s and from 1991 to present, and the estimate for PCE price inflation is close to the current FOMC target of 2 percent.

Of course, the two-sided probabilities have the benefit of hindsight; the one-sided probabilities of inflation being in the high-variance, high-persistence state were low in 1967, and it took the elevated inflation readings of 1968 and 1969 to move them up close to 100 percent. Similarly, the one-sided probabilities of inflation being in the high-variance, high-persistence state remained above 50 percent until 1993 and did not drop below 10 percent until 1995 for PCE price inflation and 1996 for CPI inflation. Since 1997, the one-sided probabilities of inflation being in the high-variance, high-persistence state have been between 2 and 6 percent in all years except 2005 and 2007, when elevated inflation readings produced probabilities of 10 and 12 percent, respectively. In 2014, the last year in the sample, the one-sided probabilities are 5 percent estimated with the CPI and 3 percent estimated with the PCE price index.

Starting from such low single digit probabilities of inflation being in the high-variance, high-persistence state, it is inflation readings outside of a two standard deviation range around μ_{S1} that move the probabilities up above those levels. For CPI inflation, that two-standard deviation range is between 0.5 and 4.0 percent, and for PCE price inflation, it is between 0.4 and 3.3 percent. However, if inflation readings outside of these ranges do not persist for more than one period, the model does not infer a high probability of inflation being in the high-variance, high-persistence state, since such a state has itself shown considerable persistence when it has appeared in the past. For example, even if Q4/Q4 PCE price inflation posts a reading of 0.0 percent in 2015, possible given the negative inflation readings early in the year due to steep oil price declines, any PCE price inflation reading between 1.3 and 2.7 percent in 2016 would produce a probability of inflation being in the high-variance, high-persistence state of below 10 percent. Also, these ranges are not immutable, since the sample used to compute the state-contingent means and variances underlying them is not yet very large. Indeed, if two inflation readings such as those considered above did occur, the estimated standard deviation of PCE price inflation in the low-variance state would increase, widening the range.

3.2 Caveats to Model Probabilities

It is important to note that since the model has been constructed symmetrically, an inflation reading a certain distance below the mean μ_{S1} produces a similar probability change as an inflation reading a certain distance above μ_{S1}, even though the sample contains no cases where low inflation readings produced a transiton to the high-variance, high-persistence state. This imposition of symmetry was informed by the out-of-sample experience of the Great Depression, which produced four consecutive years of sizable negative inflation rates as the economy contracted sharply, a period that clearly warrants classification in the high-variance, high-persistence state.[8] Nevertheless, it is reasonable to be more skeptical of sizable increases in the probability of the high-variance, high-persistence state when they are generated by low inflation readings instead of high inflation readings.

In particular, it is not clear whether it is possible for low inflation readings to produce a transiton to the high-variance, high-persistence state outside of an economic contraction approaching the severity of the Great Depression, since such a transition has not occured in the U.S. data. A severe recession may very well be a prerequisite for precipitating such a transition. If that is the case, and if the Federal Reserve's nominal interest rate policy is an important determinant of whether or not the economy enters a severe recession, the notion that low inflation in an expansion can produce a transition into the high-variance, high-persistence state is not particularly plausible, since the Federal Reserve's dual mandate makes it is unlikely they would tighten policy enough to generate a severe recession while inflation is low.

Some economic theories posit that declines in inflation that reduce inflation expectations cause households to delay durable goods purchases, with the expected inflation component of the real interest rate reducing consumer demand, possibly generating economic weakness and driving inflation down further. Such a mechanism may well be operable in environments of extreme deflation like the Great Depression. However, more recent empirical evidence from less extreme environments shows the behavior of U.S. consumers is not consistent with that theory; if anything, that evidence is consistent with declining inflation expectations *increasing* willingness to spend on durable goods–see Bachmann, Berg, and Sims (2015).[9]

[8] Estimation of the model on annual data extending back to the Great Depression confirms this.

[9] Bachmann, Berg, and Sims (2015) suggest their results may stem from money illusion, leading households

Certainly, this mechanism of lower expected inflation reducing spending by raising real interest rates does not appear operable in economic expansions on anywhere close to the scale that would be required to drive inflation into its high-variance, high-persistence state. Furthermore, mechanisms like downward nominal wage rigidity, which could feed through to downward rigidities in price inflation—see Akerlof, Dickens, and Perry (1996)—might make it unlikely that inflation could follow a random walk below certain levels, as would have to be the case after a transition to the high-variance, high-persistence state.[10]

At a minimum, the nature of the shocks producing low inflation warrants careful examination in evaluating model probabilities. Low inflation readings generated by positive developments in commodity supply, as was the case in late 2014 and early 2015 as advances in tight oil extraction technologies led to sizable increases in U.S. oil production and plunging oil prices, seem particularly unlikely to generate a transition to the high-variance, high-persistence state. As mentioned earlier in the discussion of the MA(2) dynamics of inflation, low inflation readings of this type ultimately prove favorable for economic activity and tend to produce inflation reversals later on. As such, large changes in model probabilities stemming from supply-induced spells of low inflation in economic expansions should be heavily discounted, not taken seriously. Less likely to represent unrealistic, phantom threats to the

---

to respond to nominal rather than real interest rates. Shafir, Diamond, and Tversky (1997) report results from a wide range of surveys suggestive of money illusion. Their results show a strong negative correlation between willingness to spend on durables and inflation over the past year, and while they suggest that this result supports the notion that higher expected inflation increases spending, the opposite is more likely the case because expected inflation is positively correlated with inflation over the past year. For example, from 1979-2014, the SPF 10-year expected inflation measure discussed in the next section has a correlation of 0.86 with Q4/Q4 CPI inflation in the year of the survey. The older results in Juster and Wachtel (1972) also generally support a negative effect of expected inflation on durable goods spending.

[10]Since money illusion is the main motivation for downward nominal wage rigidity in Akerlof, Dickens, and Perry (1996), and since money illusion dictates that it is nominal rather than real interest rates that matter for consumer decisions, evidence of nominal rigidities is further evidence against the notion that expected inflation has a positive effect on spending outside of extreme inflation environments in the real world.

stability of the inflation process are large probability changes stemming from low inflation in severe recessions or high inflation stemming from tightening resource utilization in economic expansions.

3.3 Transition Probabilities over Longer Horizons

A key result from the models, especially for the computation of prediction intervals, is the estimated probabilities of a regime transition over various time horizons. Probabilities of regime persistence are generally quite high, since the estimation sample of 67 periods contains only one transition out of the low-variance, low-persistence state (in 1967) and two transitions out of the high-variance, high-persistence state (in the 1950s and in 1991). Conditional on starting in the low-variance regime, the probability of being in the high variance regime k years later is a simple function of the Markov transition matrix raised to the power of k; Table 5 reports these probabilities using PCE price inflation for k ranging from 1 to 10 (as well as the very long run, k=100). After starting in the low-variance regime, the probability of inflation being in the high-variance regime is about 6 percent after two years, 13 percent after five years, and 21 percent after ten years. The probability of inflation being in the high-variance regime asymptotes to 32 percent, its unconditional probability.

These transition probabilities based on historical frequencies are important for determining the properties of inflation over the long run, so the paper evaluates their plausibility in several ways. Most of these evaluations are in the next section, but table 5 begins by comparing the probabilities from the Markov-switching model with probabilities of the permanent innovation variance from the Stock-Watson (2007) UCSV model reaching its average level in the high-variance, high-persistence state. A key defining feature of inflation in the UCSV model is its relatively high permanent innovation variance in the 1970s, so this is a natural way to compare outcomes across the two different models. The permanent innovation variance starts at its relatively low estimated level in 2014, and evolves as a logarithmic random walk, so the probability that it reaches its average level in the high-variance, high-persistence state gradually increases with the time horizon.[11] These probabilities are somewhat larger than the probabilities of transitioning to the high-variance, high-persistence state from the Markov-switching model, whose properties may be closer to other models that have been developed recently that reduce the dominance of the unit root in the UCSV model in the long run—see Chan, Koop, and Potter (2012). However, for small k values, the differences between the USCV and Markov-switching models are not particularly large.

4 Further Model Evaluation

This section evaluates model fit along a number of different dimensions, comparing model-based forecasts and distributions with analogous results from surveys and financial markets. In the model:

$$\begin{aligned}
\mathrm{E}\left(\pi_{t+1}|\mathcal{H}_t, S_{t+1}=2\right) &= \pi_t + \theta_1\hat{\varepsilon}_{t|t} + \theta_2\hat{\varepsilon}_{t-1|t} \\
\mathrm{E}\left(\pi_{t+2}|\mathcal{H}_t, S_{t+1}=2, S_{t+2}=2\right) &= \pi_t + \left(\theta_1+\theta_2\right)\hat{\varepsilon}_{t|t} + \theta_2\hat{\varepsilon}_{t-1|t} \\
\mathrm{E}\left(\pi_{t+2+k}|\mathcal{H}_t, S_{t+1}=2, S_{t+2}=2, \ldots S_{t+2+k}=2\right) &= \mathrm{E}\left(\pi_{t+2}|\mathcal{H}_t, S_{t+1}=2, S_{t+2}=2\right), \forall k.
\end{aligned}$$

Here $\hat{\varepsilon}_{t|t} = \mathrm{E}\left(\varepsilon_t|\mathcal{H}_t\right)$. In any period $t+k$ where $S_{t+k}=1$ the forecast is μ_{S1}, as it is all subsequent periods $t+k+j$ even if $S_{t+k+j}=2$ since subsequent state 2 shocks are mean zero. The overall forecast is then the probability-weighted average of these different state dependent forecasts, where the probability weights are computed by iterating $p_{t|t}$ forward in

[11] See the footnotes of the table for a description of the calculations in more detail.

time with the matrix P, similar to what is done in table 5.

4.1 Comparison with Survey-Based Measures of Long-Run Inflation Expectations

Figure 4 shows median long-run inflation expectations from the Survey of Profession Forecasters (SPF), defined as expected average CPI inflation over the next ten years, along with ten-year inflation forecasts from the Markov-switching model computed using the one-sided probability in each year. The SPF readings are from the fourth quarter of each year, and the Markov-switching model forecasts jump off from the inflation reading in the current year. The subsequent realizations of ten-year average CPI inflation are not plotted to avoid clutter, but the inset box shows RMSE computed using those realizations after each forecast through 2004.

Since the probabilities of the high-variance, high-persistence regime are close to 100 percent in the early years of the sample, those forecasts depend heavily on current-year inflation adjusted for MA(2) dynamics and on the probability of a regime transition over the next 10 years. Those transition probabilities are pure functions of the matrix P, so those 10-year average forecasts provide a gauge of how well the transition probabilities from the P matrix line up with the information in survey-based inflation expectations. They line up well: from 1979 to 1985, the largest difference between the model- and survey-based forecasts is 0.6 percentage point, while the forecasts move together almost 4 percentage points. This result mirrors Evans and Wachtel (1993), and is not guaranteed by the maximum likelihood estimation, which makes no use of the 10-year SPF information and minimizes 1-year ahead forecast errors, not 10-year ahead forecast errors.

From the mid-1980s forward, the 10-year forecasts from the model generally lead the SPF

10-year forecasts down. Subsequent realizations of 10-year average inflation were lower than most of these forecasts up through the mid-1990s, producing the relatively low RMSE for the model compared to the SPF. The inset box shows that from 1990 forward, the ratio of the two RMSEs skews more heavily in favor of the model, because the model recognized the shift out of the high-variance, high-persistence regime faster than the professional forecasters, who may have needed more convincing of the stabilization of the inflation process. While the model was designed ex post, this evidence at least suggests that the Markov-switching model might recognize regime transitions quicker than survey-based measures of long-run inflation expectations, implying that the survey measures do not exploit all available information on the state of the inflation process. In particular, tracking these measures alone might not be sufficient to recognize regime transitions in a timely manner.

Regarding the correspondence of the low-variance, low-persistence regime with the "anchoring" of long-run inflation expectations, the variability of those expectations has declined over time. From the late-1970s to early-1980s, movements in the expectations measure tended to be about a percentage point each year; from the mid-1980s to mid-1990s, movements around a half a percentage point in a year were not uncommon; from the late 1990s to 2014, movements more than a couple tenths of a percentage point in a year have been rare.[12] While the variability reduction in the early 1980s occured within the high-variance, high-persistence regime, the subsequent further settling down in variability does appear to be a reaction to the onset of the low-variance, low-persistence regime in the 1990s and its subsequent durability. Certainly, if inflation were to experience an extended period of time in the

[12]The median long-run inflation expectations measure from the University of Michigan survey, not shown, behaves similarly to the SPF measure examined here, although it has more drawbacks. For example, the survey never explicitly defines the inflation measure, and participants must round their responses to the nearest interger.

high-variance, high-persistence regime, it seems likely that long-run inflation expectations would become substantially more variable than they have been in recent years.

Tracking survey-based measures of long-run inflation expectations is a popular way of gauging the likelihood of a break in the inflation process. However, those measures were not available in the late 1960s during the last transition *into* the high-variance, high-persistence regime, so it is unclear how they would behave in such a transition. It bears keeping in mind that the SPF sample is not particularly large and that its composition fluctuates from quarter to quarter, so small movements in the median CPI measure around its 2.2 to 2.3 percent average level since the adoption of inflation targeting might be due to sampling error and thus uninformative. Interestingly, that 2.2 to 2.3 average since 2012 lines up well with the long-run Markov-switching model forecasts.

4.2 Comparison with Alternative Long-Run Tail Probabilities

While the SPF survey measure above does not provide explicit assessments of the probability distribution of inflation outcomes, other sources have become available more recently that do provide such assessments. First, an options market pricing inflation caps and floors using CPI inflation realizations has developed since 2009, and data on these options prices have been provided to researchers at the Federal Reserve Board by the company BGC partners. Kitsul and Wright (2013) discuss how to use this data to estimate cumulative density functions (cdfs) of average inflation over the next five and next ten years. Board staff constructs cdfs using similar techniques daily; this paper reports annual averages of those daily probabilities for comparison with the Markov-switching models. The options-implied cutoffs round to values of 0.5 since they are derived from an estimated discrete distribution where the probability mass around any integer k can be thought of as covering an area extending from k plus or minus 0.5. Second, since 2010, the Primary Dealer Survey conducted by the Federal Reserve Bank of New York has asked respondents to estimate probabilities of CPI inflation averaging above or below certain levels in the five year period starting five years hence. The lowest and highest cutoffs for average inflation reported in this survey are one and three percent.

Table 6 compares the 2014 annual average probabilities from these two sources with analogous probabilities from the Markov-switching model estimated using CPI inflation. In the model, inflation in each period is normally distributed with means and variances dependent on the state, and with k forecast periods to consider, there are 2^{k+2} combinations of state-dependent cdfs.[13] Using probability weights computed as before by iterating forward the 2014 value of $p_{t|t}$, these state-dependent cdfs are then combined in a weighted average to form the overall cdf. Since the state-dependent cdfs are normal, they can be computed using closed form solutions or approximated via simulations using a large number of draws. The paper takes the second approach here, and even with $k = 10$, the simulations are not particularly computationally intensive. The simulated frequencies of average inflation above or below the cutoffs over five or ten year periods are reported in the table.

The options-implied probabilities line up reasonably well with comparable probabilities from the Markov-switching model. The probabilities of 5-year and 10-year average inflation rounding to integers less than one or above four are a bit larger from the options market than from the Markov-switching model simulations. The options-implied probabilities of even more extreme outcomes are somewhat smaller than those from the model, however, and the differences are not enormous, with the Markov-switching probabilities ranging between

[13] The power is $k + 2$ instead of k because the state in the two periods prior to the start of the simulation has an effect on the estimated shocks. These simulations do not employ the techniques in Kim (1994) to collapse the number of combinations of states.

roughly half and twice their options-implied counterparts.

For reference, the third panel of the table reports the average from 2010 to 2013 of options-implied probabilites, which are much larger than the 2014 probabilities. Illiquidity or high risk aversion may have distorted these probabilities in the early days of the inflation options market. Genuinely higher perceived inflation risk may have been a factor as well. Concerns about a Depression-type economic environment might have influenced the lower tail of the options-implied distribution, and concerns about the rapid expansion of government debt and the balance sheet of the Federal Reserve might have influenced the upper tail of the options-implied distribution. If that was the case from 2010-2013, however, those concerns abated considerably in 2014.

The probabilities of five-year, five-year forward inflation averaging below one percent or above three percent are somewhat smaller from the primary dealer survey than from the Markov-switching model simulations. The bottom panel shows that the probabilities from the primary dealer survey have been relatively stable, at least compared to those from the options market. Further, the options-implied probabilities appear to have converged towards those from the dealer survey and the Markov-switching model. Overall, the options-implied probabilities and primary dealer survey probabilities in 2014 appear reasonably well approximated by the Markov-switching model. Of course, there is no guarantee that either the options market or the primary dealer survey provides the best assessment of tail risks; both are relatively new and the probabilities may evolve over time.

4.3 One-Year and Two-Year Ahead Forecast Comparisons

Traditional one-year ahead forecast evaluation is not particularly relevant for the prediction intervals studied in the next section, which substitute a non-parametic distribution derived from actual Board staff forecast errors for the model-estimated distribution in the low-variance, low-persistence state. Nevertheless, comparison of one- and two-year ahead forecast errors from the model with those from alternative sources is informative.

Figure 5 plots one-year ahead CPI inflation forecasts from three sources, the Markov-switching model, the SPF median, and the Federal Reserve Board staff.[14] The forecasts extend from 1981 to 2014 except the Board staff forecasts which stop in 2009 since they typically are released to the public with a five year lag. Over the full sample, the Markov-switching model forecasts track the SPF forecasts reasonably well, except for some volatility in the mid-1980s, when the model may have been too quick to infer a possible regime transition after oil price declines led to a transitorily low inflation reading in 1986. This result reinforces the prior discussion on the importance of recognizing the nature of inflation shocks and tempering estimated probabilities appropriately. Despite that mid-1980s volatility, and the fact that the SPF forecasts reflect more information than the univariate Markov-switching forecasts based on inflation alone, the first inset box shows the RMSE of the SPF forecasts is not much lower than the RMSE of the forecasts from the model over the full sample. And from 1990-2014, the model RMSE is slightly lower than the SPF RMSE, again reflecting the model's more-timely recognition of the 1990s transition to the low-variance, low-persistence regime.

The second inset box in Figure 5 reports forecast error statistics over a period starting in 1999 to align the start date with the PCE price results below, and ending in 2010 to bring the Board staff forecasts into the comparison. In addition to RMSE, the inset box reports key

[14]These Markov-switching model forecasts use the latest parameter estimates, but forecasts from rolling Markov-switching model estimates produced similar results.

percentiles of the distribution of forecast errors.[15] The SPF and Markov-switching forecast statistics are similar, while the Board staff forecasts have a larger RMSE and a more one-sided distribution of forecast errors. The 15th percentile of the distribution of Board staff forecast errors is only minus 0.4, while the 85th percentile is plus 2.3. Like many findings in the literature on forecasting, the result in Romer and Romer (2000) that Board staff inflation forecasts outperform their private sector counterparts appears to have ceased being an accurate description of reality soon after the paper's publication. The sample of Q4/Q4 growth rates examined here is not large, but the relevant question now appears to be why Board staff forecasts haved *underperformed* their private sector counterparts over the past decade and a half.

Figures 6 and 7 report one-year and two-year ahead PCE price inflation forecasts from the Board staff starting in 1997, the year these first began being reported regularly in the main inflation forecast table in the staff's Greenbook. Since these are unavailable from 2010 to 2014, the figures substitute midpoints of central tendencies from fourth-quarter FOMC SEP forecasts.[16] These are plotted along with one- and two-year ahead forecasts from the Markov-switching model estimated on PCE price inflation, along with actual PCE price inflation shifted back either one or two years to align it with the forecasts. From 1997 to

[15]The number of observations comprising the discrete distribution from which these percentiles are computed is small, so these percentiles linearly interpolate between observations when the relevant percentile does not fall exactly on an observation percentile in the discrete distribution. Procedures that do not interpolate, instead bootstrapping draws from the distribution of errors, produce similar percentiles to those reported here.

[16]While the midpoints of these central tendencies are not perfect substitutes for the Board staff forecasts, the FOMC participants do all have access to the staff's Greenbook in real time. Examination of FOMC transcripts from 2007 to 2009 suggests a number of FOMC participants set their forecasts in close alignment with the Greenbook, although others emphasize the differences between their forecasts and the Greenbook. Those differences might define the edges of the central tendency reported in many SEP rounds.

2014, these forecasts are close to a constant equal to μ_{S1}.[17]

The sample in Figures 6 and 7 permits computation of one- and two-year ahead forecast errors from 1999 to 2015, and inset boxes in each figure report forecast error statistics over this period. Faust and Wright (2012) argue that a good benchmark model for forecasting inflation is one where inflation adjusts smoothly from its current level to the SPF long-run expectation, and since μ_{S1} is not far from that expectation, the Markov-switching model forecasts are close to that type of benchmark model. However, the one-year ahead forecasts from the Markov-switching model assume rapid adjustment to the long-run expectation; to examine a forecasting model assuming much slower adjustment, the inset boxes report forecast error statistics for pure random walk forecasts as well.[18]

The results in Figures 6 and 7 mirror those discussed earlier. First, the Markov-switching model forecasts have lower RMSE than the random walk forecasts, not surprising given the lack of inflation persistence over this period. Second, the Markov-switching model forecasts also have lower RMSE than the Board staff forecasts, and the visually-evident negative correlation between the Board staff forecasts and actual inflation in Figure 7 shows why the near-constant Markov-switching model forecasts outperform. Third, the distribution of Board staff forecast errors is again one-sided, more so than the distributions from the models. One out of every six or seven Board staff forecasts was far too low over this sample, by at least 1.5 percentage points, but on the other end of the distribution, the forecasts on the high side missed by only one third to one half as much. However, it is interesting to note that the two-year ahead SEP forecasts have been similar to the Markov-switching model

[17] These Markov-switching model forecasts are based on the latest parameter estimates and data, but rolling estimation using real-time data produces similar results.

[18] These random walk inflation forecasts for each of the next two years are equal to the current-year inflation estimate from each December Greenbook or fourth-quarter SEP.

forecasts since 2011, as shown in Figure 7.

5 Prediction Intervals

5.1 Parametric versus Non-parametric: Considerations

While the Markov-switching model forecasts perform better than the Board staff forecasts over the past decade and a half, the Board staff and forecasts employ much more information. One example is the last one-year ahead SEP forecast from December 2014 that predicted relatively low inflation in 2015, shown in Figure 6. Oil prices had plunged in late 2014, and it was predictable that this would translate into a plunge in energy prices early in 2015 that likely would hold down PCE price inflation for the year as a whole. Such predictability can arise from a variety of factors that are difficult to capture in a parsimonious time series model. Forecasters in the real world also may respond to incentives that lead them to set their forecasts not to maximize accuracy and minimize RMSE, but to serve some other goal. These considerations, if expected to continue going forward, should be reflected in properly constructed prediction intervals, leading to the strong presumption that as little model structure should be imposed on them as possible, since no model can incorporate all these factors. Less structure favors prediction intervals constructed from the non-parametric distribution of errors the forecasters have actually made, which reflect asymmetries and other non-standard features. Distributions derived from a model or that impose symmetry or normality should be used only as a last resort.

Still, although the bar should be high for imposing unecessary structure that rules out past patterns continuing into the future, it is important to weigh carefully whether past patterns are likely to continue going forward. In the case at hand, the most pronounced pattern in the distribution of Board staff forecast errors from 1999-2014 is its asymmetry, the fact that the Board staff forecast misses at one end of the distribution, when the forecasts were too low, were much larger than the misses at the other end of the distribution, when the forecasts were too high. Evaluation of the possible causes of this pattern is required to shed light on whether it is likely to persist.

At least three classes of potential explanations might account for the asymmetric pattern in the Board staff forecast errors. First, the "true" asymptotic distribution of forecast errors may be symmetric, with the asymmetry over the 16 years from 1999-2014 just due to chance. The sample here is small, and the increase in energy and other commodity prices over this period was plausibly unanticipated and likely accounts for some of the asymmetry. However, it probably does not explain all of it. A sizable portion of that increase in energy and other commodity prices reversed by the middle of 2015, and adding a low 2015 inflation value to the sample does not change the asymmetry in the distribution of Board staff forecast errors very much. Furthermore, the Board staff forecast errors for "core" PCE price inflation—PCE price inflation excluding food and energy from the consumption bundle—retains some positive asymmetry from 1999-2014, even though Board staff unemployment rate forecast errors over this period show considerable positive skewness which should have led to *negative* skewness in the inflation forecast errors according to a Phillips curve relation.[19] And, as discussed further below, Board staff forecast errors have been more asymmetric than those

[19]As before, calculations substitute SEP forecasts for Board staff forecasts over the last five years of the sample. Over that 1999-2014 sample, the 15th percentile of the distribution of one-year-ahead Board staff forecast errors for "core" PCE price inflation is minus 0.3, while the 85th percentile is plus 0.6. For two-year-ahead forecasts, the 15th percentile is minus 0.6, while the 85th percentile is plus 0.8. From 1999-2008, the latest "core" PCE price inflation measure used to compute forecast errors strips out "food services", sometimes called "food away from home", because that category was not included in "core" until 2009. Price inflation in that category has tended to be higher than the rest of core, so removing that category from core reduces the forecast error asymmetry somewhat.

from the private sector and from impartial models, with all the forecasts evaluated against the same set of shocks, including those from energy prices.

The second class of explanations for the asymmetry in the Board staff forecast errors are explanations that imply the "true" asymptotic distribution of forecast errors is asymmetric. Such an asymmetry could arise for a number of reasons, including economic mechanisms like downward nominal wage rigidity feeding through to downward rigidities in consumer prices—see Akerlof, Dickens, and Perry (1996). If this class of explanations has some validity, then inflation forecasts set to minimize RMSE likely would produce a forecast error distribution that is asymmetric.

The third class of explanations is that the Board staff sets its forecasts in an idiosyncratic manner that generates some of the forecast error asymmetry. The simplest way to examine this class of explanations is to compare the distribution of Board staff forecast errors with the distribution of forecast errors either from outside forecasters who are more likely to minimize RMSE or from impartial models like a constant or a random walk, as in the inset boxes of Figures 5, 6, and 7. All these forecasts are evaluated against the same set of possibly asymmetric inflation realizations, and while all of the distributions exhibit some asymmetry, that asymmetry is most pronounced in the distribution of Board staff forecast errors. For one-year ahead CPI inflation, the largest discrepancies between the SPF and Board staff forecasts, in 2003, 2004 and 2008, drive this differential asymmetry.

These results suggest the Board staff may have either systematically misjudged the likelihood of different inflation outcomes, or set its inflation forecasts in a way that does not maximize accuracy and minimize RMSE. The first possibility might stem from a mistaken belief that inflation expectations would fall persistently and hold down inflation after recessions. The second possibility might stem from the Board staff setting its forecast according to an asymmetric loss function, in which the cost of a forecast that is too high is perceived to be higher than the cost of a forecast that is too low. Since the Board staff forecasts are prepared for FOMC meetings, with lower inflation forecasts possibly leading to looser monetary policy than would otherwise be the case, one possibility is that the Board staff either consciously or unconsciously made a cost-benefit analysis that favored running a looser monetary policy than would obtain under an RMSE-minimizing forecast, and then altered its forecasts to produce the outcome it favored. In other words, the Board staff may have set its forecasts based on stategic considerations, with the goal of moving monetary policy in an easing direction that the Board staff viewed as desirable based on a perceived asymmetric loss function.

The relatively-low Board staff inflation forecasts in 2003, 2004, and 2008 suggest that these idiosyncratic factors influencing the Board staff inflation forecasts—either the misjudgments about the likelihood that low inflation would persist or a preference for policy easing based on an asymmetric loss function—were strongest in those years. However, the Board staff forecasts have been lower than the forecasts from the SPF or the models in most years since 1999 (see Figures 5, 6 and 7), suggesting these idiosyncratic factors may be relatively stable, and thus likely to continue to influence Board staff forecasts going forward.

Of these three classes of explanations for the asymmetry in the Board staff inflation forecasts, only the first would justify imposing symmetry on the distribution of forecast errors in the low-variance, low-persistences state going forward. The evidence favorable to the other two hypotheses suggests that would be unwarranted.

5.2 Interval Construction

Based on the above considerations, this section computes prediction intervals substituting the non-parametic distribution derived from actual Tealbook forecast errors from 1999-2014 for the model-estimated distribution in the low-variance, low-persistence state. Using the appropriate probability weights, the overall cdf used to compute forecast error percentiles is a weighted average of (1) the discrete distribution of forecast errors from 1999-2014, linearly interpolating between observations, and (2) a normal cdf computed using model parameters in the high-variance, high-persistence state. These cdfs are computed using closed form solutions rather than simulations.

Since the model shows very high probabilities that inflation was in the low-variance, low-persistence state throughout 1999-2014, the discrete distribution is a non-parametric estimate of the forecast error distribution conditional on that state, $g\left(\pi_t - \widehat{\pi_t^{bs}}|S_t = 1\right)$. Here $\widehat{\pi_t^{bs}}$ are the actual Board staff forecasts, allowing asymmetries and any other features of those forecast errors to show through to the distribution and intervals. Of course, the cost to this approach is the lumpiness it introduces through use of a discrete distribution, mitigated somewhat by the interpolation between observations.

Since no forecast observations are available from the high-variance, high-persistence state, those errors must be approximated using the parametric model:

$$g\left(\pi_t - \widehat{\pi_t^{bs}}|S_t = 2\right) \approx N\left(\pi_t - \left(\pi_{t-1} + \theta_1\varepsilon_{t-1} + \theta_2\varepsilon_{t-2}\right), \sigma_{s2}^2\right).$$

While this is an approximation, the results from the last section suggest the variance of the error term and the random walk properties of the distribution are approximately correct, and those are the key features pinning down the forecast uncertainty beyond one period ahead.

The probabilities from the parametric model are then used to combine the non-parametric $g\left(\pi_t - \widehat{\pi_t^{bs}}|S_t = 1\right)$ with the parametric $g\left(\pi_t - \widehat{\pi_t^{bs}}|S_t = 2\right)$ to produce the overall cdf and forecast error percentiles.

The top panel of table 7 reports forecast error percentiles around the December 2014 SEP forecasts for inflation in 2015 and 2016. With the estimated probability of the high-variance, high-persistence state so low in 2014, these percentiles are very close to those computed from the non-parametric distribution $g\left(\pi_t - \widehat{\pi_t^{bs}}|S_t = 1\right)$ alone. The median forecast error is positive, so the 50th percentiles of the distributions are above the SEP midpoint forecasts of 1.3 percent in 2015 and 1.85 percent in 2016. The lower end of each distribution is more compact than the upper end, reflecting the asymmetry discussed above. For reference, the bottom half of the table reports prediction intervals from the purely parametric Markov switching model, which makes no use of the distribution of actual forecast errors. The one-year ahead prediction intervals are similar to the preferred non-parametric intervals at the top of the table, although those are centered around a forecast that is a half percentage point lower. Comparing the two-year ahead intervals based around similar forecasts shows the asymmetries in the non-parametric distribution more clearly.

The remaining panels of the table use alternative simulations to illustrate how the prediction intervals change with a hypothetical increase to 33 percent in the probability of the high-variance, high-persistence state. Several points are noteworthy. First, while the 15th, 50th, and 85th percentiles do move with this probability increase, the changes in these percentiles are not large. A high-variance, high-persistence probability close to 100 percent clearly would matter much more for even these middle percentiles, but it is the tail percentiles, the 5th and 95th, that respond most strongly to the probability increase examined here. And those tail percentiles are largely driven by the model itself, not the empirical distribution of forecast errors: the 5th and 95th percentiles from the recommended hybrid version of the model in the top panels are quite close to those from the purely parametric model in the bottom panels. Second, the width of the 90 percent prediction interval defined by those tail percentiles now grows much more markedly with the forecast horizon, due to the random walk in the high-variance, high-persistence state. Finally, the second simulation changes the jumping off value for inflation from 1.1 percent, the estimate of 2014 PCE price inflation at the time of this writing, to 3.0 percent.[20] The jumping off value has a large effect on the tail percentiles, again due to the random walk in the high-variance, high-persistence state.

6 Conclusion

This paper constructs Markov-switching models that estimate the probability of inflation returning to a high-variance, high-persistence regime similar to the 1970s and 1980s, and uses those models to generate prediction intervals around Federal Reserve Board staff forecasts of PCE price inflation. Not surprisingly, the model probabilities show the low-variance regime with a stable mean has governed the behavior of inflation since the 1990s, and as long as the probability of the high-variance, high-persistence regime remains very low, the model's prediction intervals around one- and two-year ahead forecasts will be well-approximated by those computed from the distribution of forecast errors in the low-variance regime. The paper computes that distribution in the low-variance regime non-parametrically, and shows how to combine that non-parametric distribution with the parametric information from the Markov-switching model. If the model-based probability of the high-variance, high-persistence regime

[20] For simplicity, estimated shock values are left unrevised.

were to increase substantially, the width of the prediction intervals, particularly those using the far tails of the mixture distribution, would widen substantially as well.

Additional results in the paper that may be innovations to some information sets include:

1. Over the full 1948-2014 sample, time series evidence strongly suggests the four-quarter inflation measures examined here follow MA(2) processes in differences. On average, about half of the typical innovation to the inflation process reverses after two years, with reversals following large innovations to oil and other commodity prices contributing importantly to this result.

2. Over the last 20 years of the sample, a wide range of evidence suggests there is no random walk component to inflation of any quantitative significance for forecasting purposes. Estimating the probability that a random walk reappears in the inflation process, is, of course, one of the main goals of the paper.

3. Survey measures of long-run inflation expectations, often used to make inferences about whether the current low-variance, low-persistence inflation regime is likely to continue, were unavailable in the late 1960s during the last transition out of such a regime, and were comparatively slow to recognize the transition into the current low-variance regime in the 1990s. The Markov-switching models were much faster in recognizing that regime transition in the 1990s, so they may provide more accurate and timely inferences on regime transitions going forward as well.

4. The tail risks to inflation over the next 5 or 10 years, as measured by financial market derivatives, collapsed in 2014. After that collapse, those measures of tail risks are now much closer to those estimated using the Markov-switching model.

5. The Romer and Romer (2000) result that Federal Reserve Board staff forecasts have outperformed their private sector counterparts has not held out of sample, at least for the four-quarter growth rates targeted by the forecasts of FOMC members. Since 1999, Board staff forecasts have been less accurate (as measured by RMSE) than forecasts from both the private sector and the Markov-switching models. The distribution of Board staff forecast errors has taken on a noticeable asymmetry as well: the forecasts at one end of the error distribution, those that were too low, missed by much more than the forecasts at the other end of the distribution that were too high. This asymmetry is largely idiosyncratic to the Board staff—private-sector forecast error distributions or those from impartial models like a random walk or a constant exhibit much less asymmetry—and may stem either from consistent misjudgments about the likelihood that low inflation would persist or from a desire by Board staff to justify loose monetary policy, particularly in 2003, 2004, and in the aftermath of the Great Recession.

Bibliography

Atkeson, Andrew, and Ohanian, Lee. "Are Phillips curves useful for forecasting inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review* 25, pp. 2-11.

Akerlof, George; Dickens, William; and Perry, George. "The Macroeconomics of Low Inflation" *Brookings Papers on Economic Activity*, 1996:1, pp. 1-76.

Bachmann, Rudiger; Berg, Tim; and Sims, Eric. "Inflation Expectations and Readiness to Spend: Cross-Sectional Evidence" *American Economic Journal: Economic Policy*, 2015, pp. 1-35.

Benati, Luca. "Investigating Inflation Persistence Across Monetary Regimes" *Quarterly Journal of Economics*, 2008, pp. 1005-1060.

Bernanke, Ben. "The Economic Outlook and Monetary Policy," speech delivered at the Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 27th, 2010.

Chan, Joshua; Koop, Gary; and Potter, Simon. "A New Model of Trend Inflation," working paper, 2012.

Cogley, Timothy; Primiceri, Giorgio; and Sargent, Thomas. "Inflation-Gap Persistence in the US" *American Economic Journal: Macroeconomics*, 2010, pp. 43-69.

Davig, Troy and Doh, Taeyoung. "Monetary Policy Regime Shifts and Inflation Persistence," *The Review of Economics and Statistics*, 2014, pp. 862-875.

Evans, Martin, and Wachtel, Paul. "Inflation Regimes and the Sources of Inflation Uncertainty" *Journal of Money, Credit and Banking*, 1993, pp. 475-511.

Faust, Jon, and Wright, Jonathan. "Forecasting Inflation," *Handbook of Economic Forecasting,* 2012.

Hamilton, James. "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, 1989, pp. 357-384.

Juster, F. Thomas and Wachtel, Paul. "Inflation and the Consumer" *Brookings Papers on Economic Activity*, 1972:1, pp. 71-121.

Kilian, Lutz, and Vigfusson, Rob. "Do Oil Prices Help Forecast U.S. Real GDP? The Role of Nonlinearities and Asymmetries," *Journal of Business and Economic Statistics*, 2013, pp. 78-93.

Kim, Chang-Jin. "Unobserved-Component Time Series Models with Markov-Switching Heteroscedasticity: Changes in Regime and the Link Between Inflation Rates and Inflation Uncertainty," *Journal of Business and Economic Statistics*, (1993), pp. 341-349.

Kim, Chang-Jin. "Dynamic linear models with Markov-Switching," *Journal of Econo-*

metrics, 1994, pp. 1-22.

Kitsul, Yuriy, and Wright, Jonathan. "The Economics of Options-Implied Inflation Probability Density Functions" *Journal of Financial Economics*, 2013, pp. 696-711.

Lanne, Markku. "Nonlinear Dynamics of Interest Rate and Inflation," *Journal of Applied Econometrics*, 2006, pp. 1157-1168.

Reifschneider, David and Tulip, Peter. "Gauging the Uncertainty of the Economic Outlook from Historical Forecasting Errors," Finance and Economics Discussion Series 2007-60 (Washington: Board of Governors of the Federal Reserve System, November 2007).

Romer, Christina and Romer, David. "Federal Reserve Information and the Behavior of Interest Rates" *American Economic Review* 90, 2000, pp. 429-457.

Shafir, Eldar; Diamond, Peter; and Tversky, Amos. "Money Illusion," *Quarterly Journal of Economics*, 1997, pp. 341-374.

Stock, James; and Watson, Mark. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking*, 2007, pp. 3-34.

Table 1: MA(2) Process Estimates for the Change in Inflation

$$\pi_t = \pi_{t-1} + \varepsilon_t + \theta_1 \varepsilon_{t-1} + \theta_2 \varepsilon_{t-2}$$

	θ_1	θ_2	σ_ε
1948-2014			
CPI inflation	**-0.05**	**-0.51**	**1.98**
	(0.11)	(0.11)	(0.17)
PCE price inflation	**-0.09**	**-0.47**	**1.68**
	(0.11)	(0.11)	(0.14)
1995-2014			
CPI inflation	**-0.84**	**-0.03**	**0.91**
	(0.26)	(0.26)	(0.15)
PCE price inflation	**-0.82**	**-0.18**	**0.71**
	(0.30)	(0.23)	(0.13)

Table 2: AR(1) Process Estimates for Inflation

$$\pi_t = \alpha + \gamma \pi_{t-1} + u_t$$

	α	γ
1995-2014		
CPI inflation	**2.26**	**0.03**
	(0.62)	(0.25)
PCE price inflation	**1.63**	**0.12**
	(0.49)	(0.24)

Table 3: Out-of-Sample RMSEs from 1995 to 2014 October Blue Chip Consensus Forecasts of Annual CPI Inflation from 0 to 4 years ahead

	1999-2014				
Years Ahead (k)	0	1	2	3	4
$RMSE(k)$	0.06	0.72	1.03	0.96	1.01

Note: Root Mean Square Forecast Errors (RMSEs) from 1995 to 2014 October Blue Chip long range consensus forecasts made in year T of annual average CPI inflation in years $T+k$, $k = 0, 1, 2, 3, 4$.

Table 4: Markov Switching Models for Inflation, 1948-2014

	μ_{s1}	σ_{s1}	θ_1	θ_2	σ_{s2}	p_{11}	p_{22}	$p_{1,init}$
CPI inflation	**2.24**	**0.89**	**0.14**	**-0.65**	**2.40**	**0.963**	**0.938**	**0.00**
	(0.18)	(0.12)	(0.18)	(0.18)	(0.32)	(0.038)	(0.047)	(0.99)
PCE price inflation	**1.84**	**0.73**	**0.03**	**-0.57**	**2.16**	**0.967**	**0.932**	**0.00**
	(0.13)	(0.09)	(0.16)	(0.15)	(0.28)	(0.032)	(0.047)	(1.01)

Table 5: Probability of High-Variance Random-Walk PCE Price Inflation Regime in Year T+k, Conditional on Starting in the Low-Variance Inflation Regime in Year T

	Probabilities, in Percentage Points (%)	
k	Markov-Switching Model	UCSV model
0	0	0
1	3	3
2	6	8
3	9	13
4	11	16
5	13	19
6	15	21
7	17	23
8	19	24
9	20	26
10	21	27
100	32	42

Note: The probability in year k from the Stock-Watson (2007) UCSV model is the probability that the permanent innovation variance ($\sigma^2_{\varepsilon,t}$) from the model exceeds 0.76, starting $\sigma^2_{\varepsilon,t}$ in year 0 at its estimated 2014 value of 0.13 and assuming the process evolves as in the model: $\ln\sigma^2_{\varepsilon,t} = \ln\sigma^2_{\varepsilon,t-1} + \nu_{\varepsilon,t}$, where $\text{var}(\nu_{\varepsilon,t}) = 0.2$. The threshold value of 0.76 is the average permanent innovation variance in the high-variance, high-persistence state, or the weighted average value of the time series of estimated median values for $\sigma^2_{\varepsilon,t}$ from the UCSV model estimated on total PCE price inflation from 1960-2014 with weights equal to the smoothed probabilities from the Markov switching model estimated in table 4. The average permanent innovation variance in the low-variance state is much lower, at 0.19.

Table 6: Estimated probabilities of CPI inflation averaging above or below certain values over various horizons

	Probabilities, in Percentage Points (%)					
	cdf(k), or prob($\overline{\pi_{CPI}} < k$)			1-cdf(k), or prob($\overline{\pi_{CPI}} > k$)		
	$k = -1.5$	$k = -0.5$	$k = 0.5$	$k = 3.5$	$k = 4.5$	$k = 5.5$
Horizon:						
	Implied by Financial Derivatives, 2014					
next 5 years	0.5	1.4	5.1	6.1	1.6	0.6
next 10 years	0.6	1.7	5.1	12.8	3.3	1.0
	Markov-Switching Model, 2104					
next 5 years	1.0	1.9	3.6	4.6	2.3	1.1
next 10 years	1.2	2.3	4.6	6.2	3.1	1.5
	Memo: Implied by Financial Derivatives, 2010-2013					
next 5 years	2.2	4.5	10.7	14.1	5.1	2.1
next 10 years	2.1	4.3	9.9	21.1	8.4	3.6

	cdf(k), or prob($\overline{\pi_{CPI}} < k$)	1-cdf(k), or prob($\overline{\pi_{CPI}} > k$)
	$k = 1.0$	k=3.0
Horizon:		
	NY Fed Primary Dealer Survey, 2014	
5 yr, 5 yr forward	3.3	8.7
	Markov-Switching Model, 2104	
5 yr, 5 yr forward	8.7	12.7
	Memo: NY Fed Primary Dealer Survey, 2010-213	
5 yr, 5 yr forward	4.2	9.4

Table 7: Prediction Intervals

Using Non-Parametric Distribution of Actual Board Staff Forecast Errors in Low-Variance State

Forecast:	Years Ahead:	Forecast Error Percentiles: 5th	15th	50th	85th	95th
		2014 estimates: $\widehat{p_2} = 0.03$; $\pi_T = 1.1$				
$\hat{\pi}_{T+1} = 1.3$	1	0.7	0.8	1.6	2.9	3.3
$\hat{\pi}_{T+2} = 1.85$	2	1.0	1.2	2.1	3.5	3.9
		Alternative Simulation: $\widehat{p_2} = 0.33$; $\pi_T = 1.1$				
$\hat{\pi}_{T+1} = 1.3$	1	-0.9	0.7	1.6	3.1	3.6
$\hat{\pi}_{T+2} = 1.85$	2	-1.6	1.0	2.0	3.6	4.7
		Alternative Simulation: $\widehat{p_2} = 0.33$; $\pi_T = 3.0$				
$\hat{\pi}_{T+1} = 1.3$	1	0.7	0.8	1.8	3.5	5.5
$\hat{\pi}_{T+2} = 1.85$	2	0.1	1.1	2.2	3.9	6.5

Memo: Purely Parametric Markov-Switching Model

Forecast:	Years Ahead:	Forecast Error Percentiles: 5th	15th	50th	85th	95th
		2014 estimates: $\widehat{p_2} = 0.03$; $\pi_T = 1.1$				
$\hat{\pi}_{T+1} = 1.8$	1	0.5	1.0	1.8	2.6	3.1
$\hat{\pi}_{T+2} = 1.8$	2	0.4	1.0	1.8	2.7	3.3
		Alternative Simulation: $\widehat{p_2} = 0.33$; $\pi_T = 1.1$				
$\hat{\pi}_{T+1} = 1.8$	1	-0.9	0.6	1.8	2.8	3.7
$\hat{\pi}_{T+2} = 1.8$	2	-1.6	0.6	1.8	2.9	4.7
		Alternative Simulation: $\widehat{p_2} = 0.33$; $\pi_T = 3.0$				
$\hat{\pi}_{T+1} = 2.0$	1	0.4	1.1	2.0	3.6	5.5
$\hat{\pi}_{T+2} = 2.0$	2	0.0	1.0	2.0	3.7	6.5

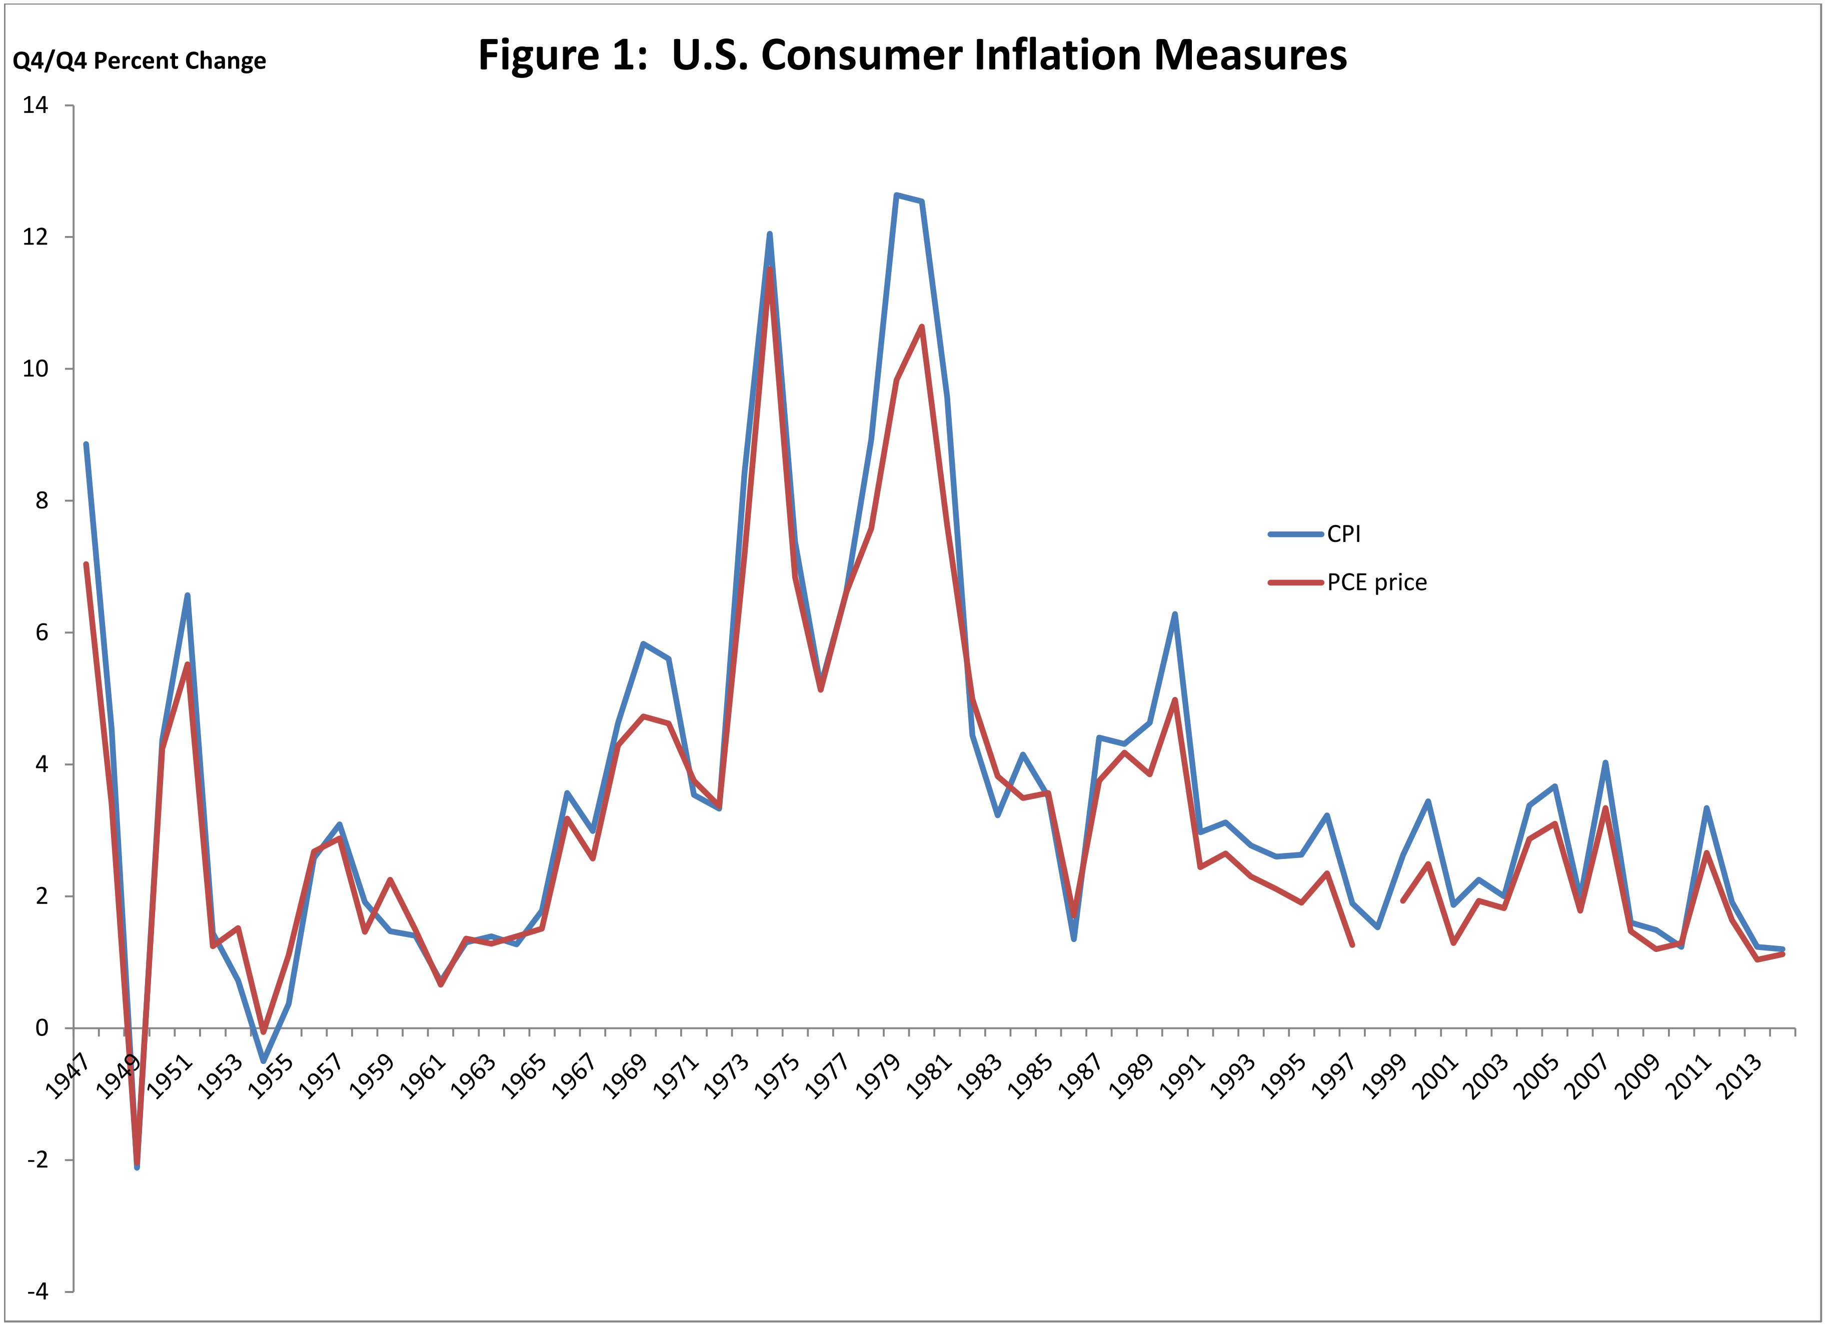

Figure 1: U.S. Consumer Inflation Measures
Q4/Q4 Percent Change
14
12
10
8
6
4
2
0
-2
-4
1947
1949
1951
1953
1955
1957
1959
1961
1963
1965
1967
1969
1971
1973
1975
1977
1979
1981
1983
1985
1987
1989
1991
1993
1995
1997
1999
2001
2003
2005
2007
2009
2011
2013
CPI
PCE price

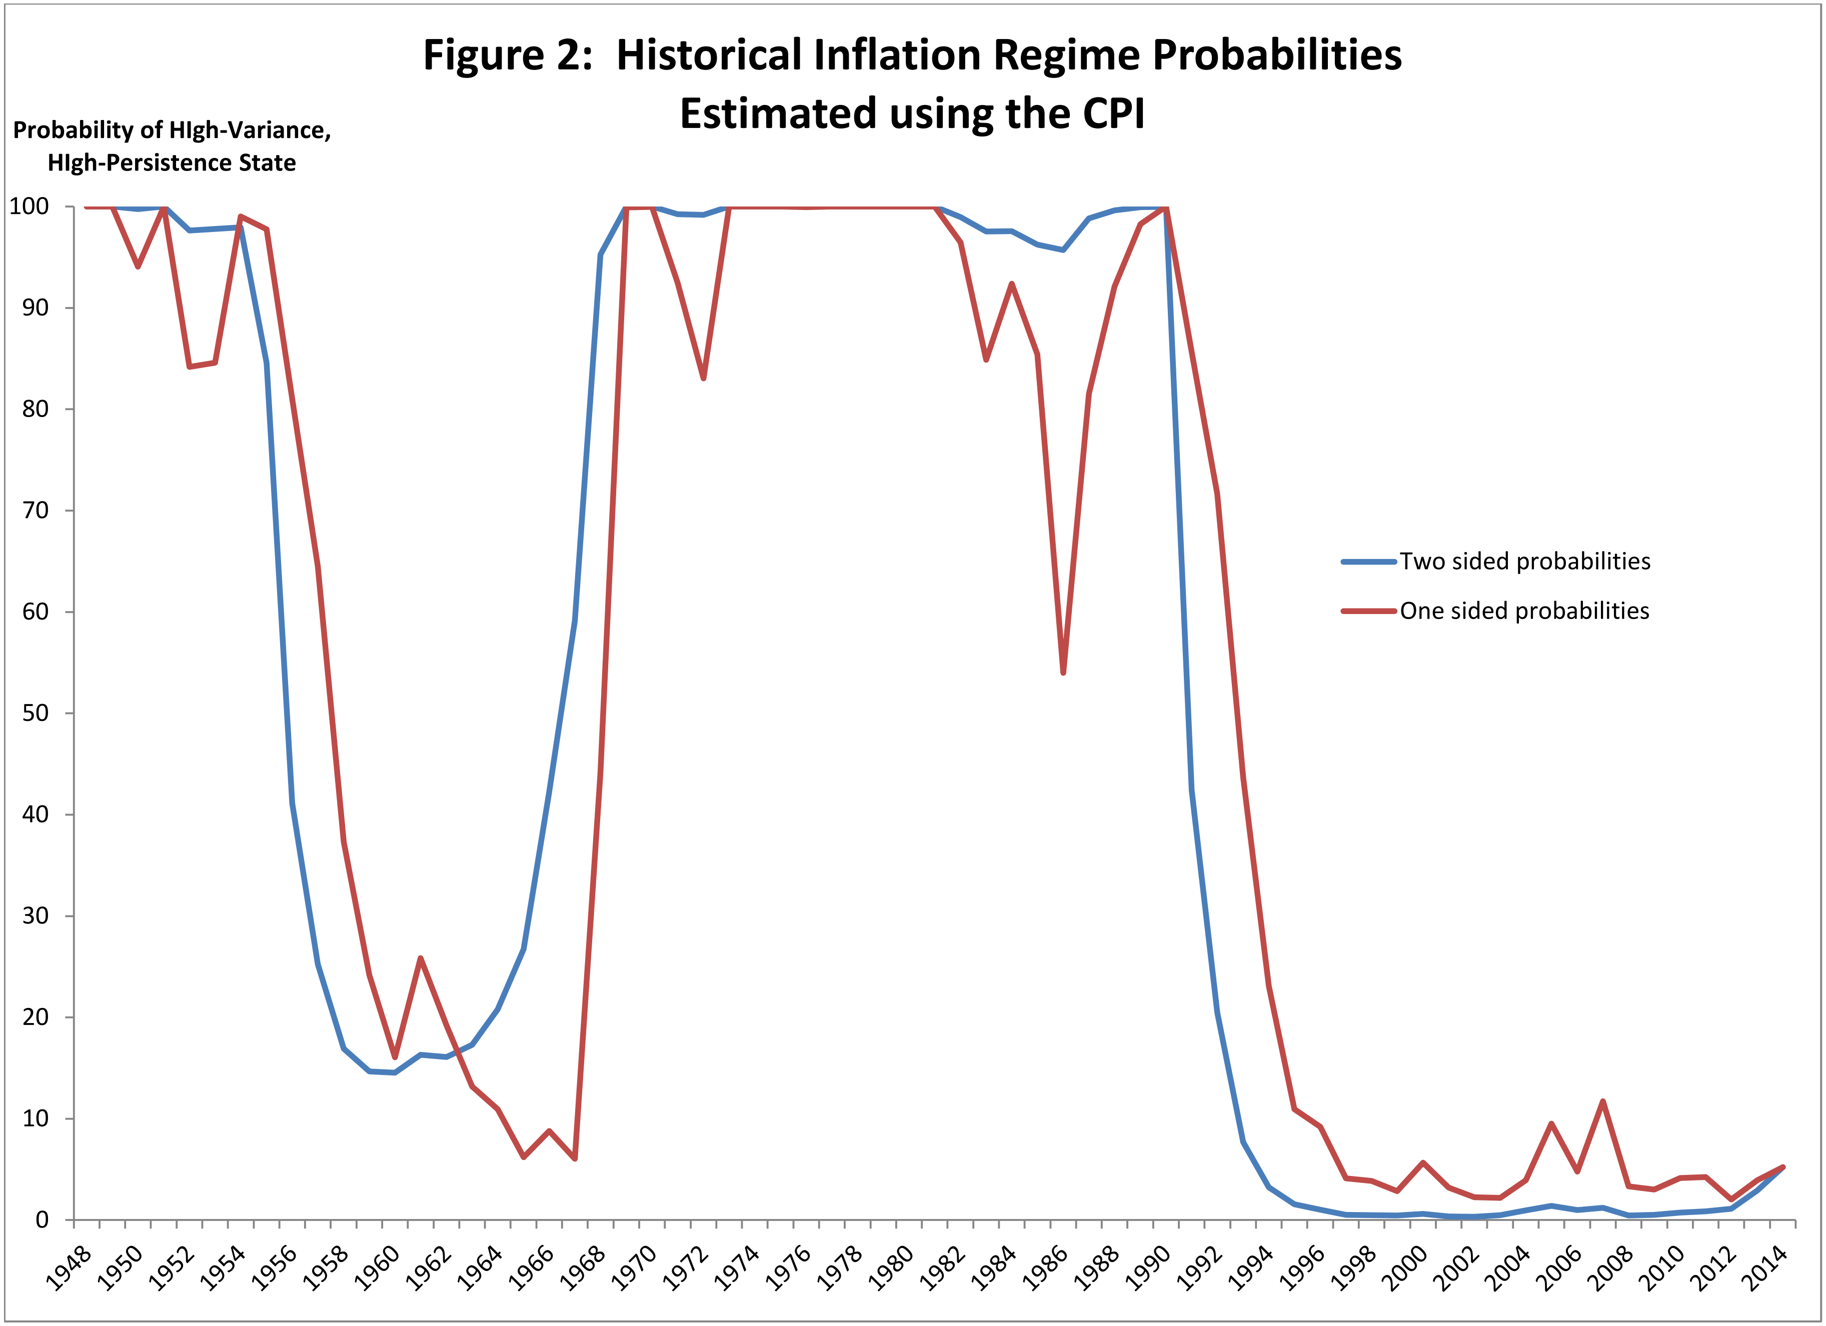
Figure 2: Historical Inflation Regime Probabilities
Estimated using the CPI
Probability of HIgh-Variance,
HIgh-Persistence State
100
90
80
70
60
50
40
30
20
10
0
1948
1950
1952
1954
1956
1958
1960
1962
1964
1966
1968
1970
1972
1974
1976
1978
1980
1982
1984
1986
1988
1990
1992
1994
1996
1998
2000
2002
2004
2006
2008
2010
2012
2014
Two sided probabilities
One sided probabilities

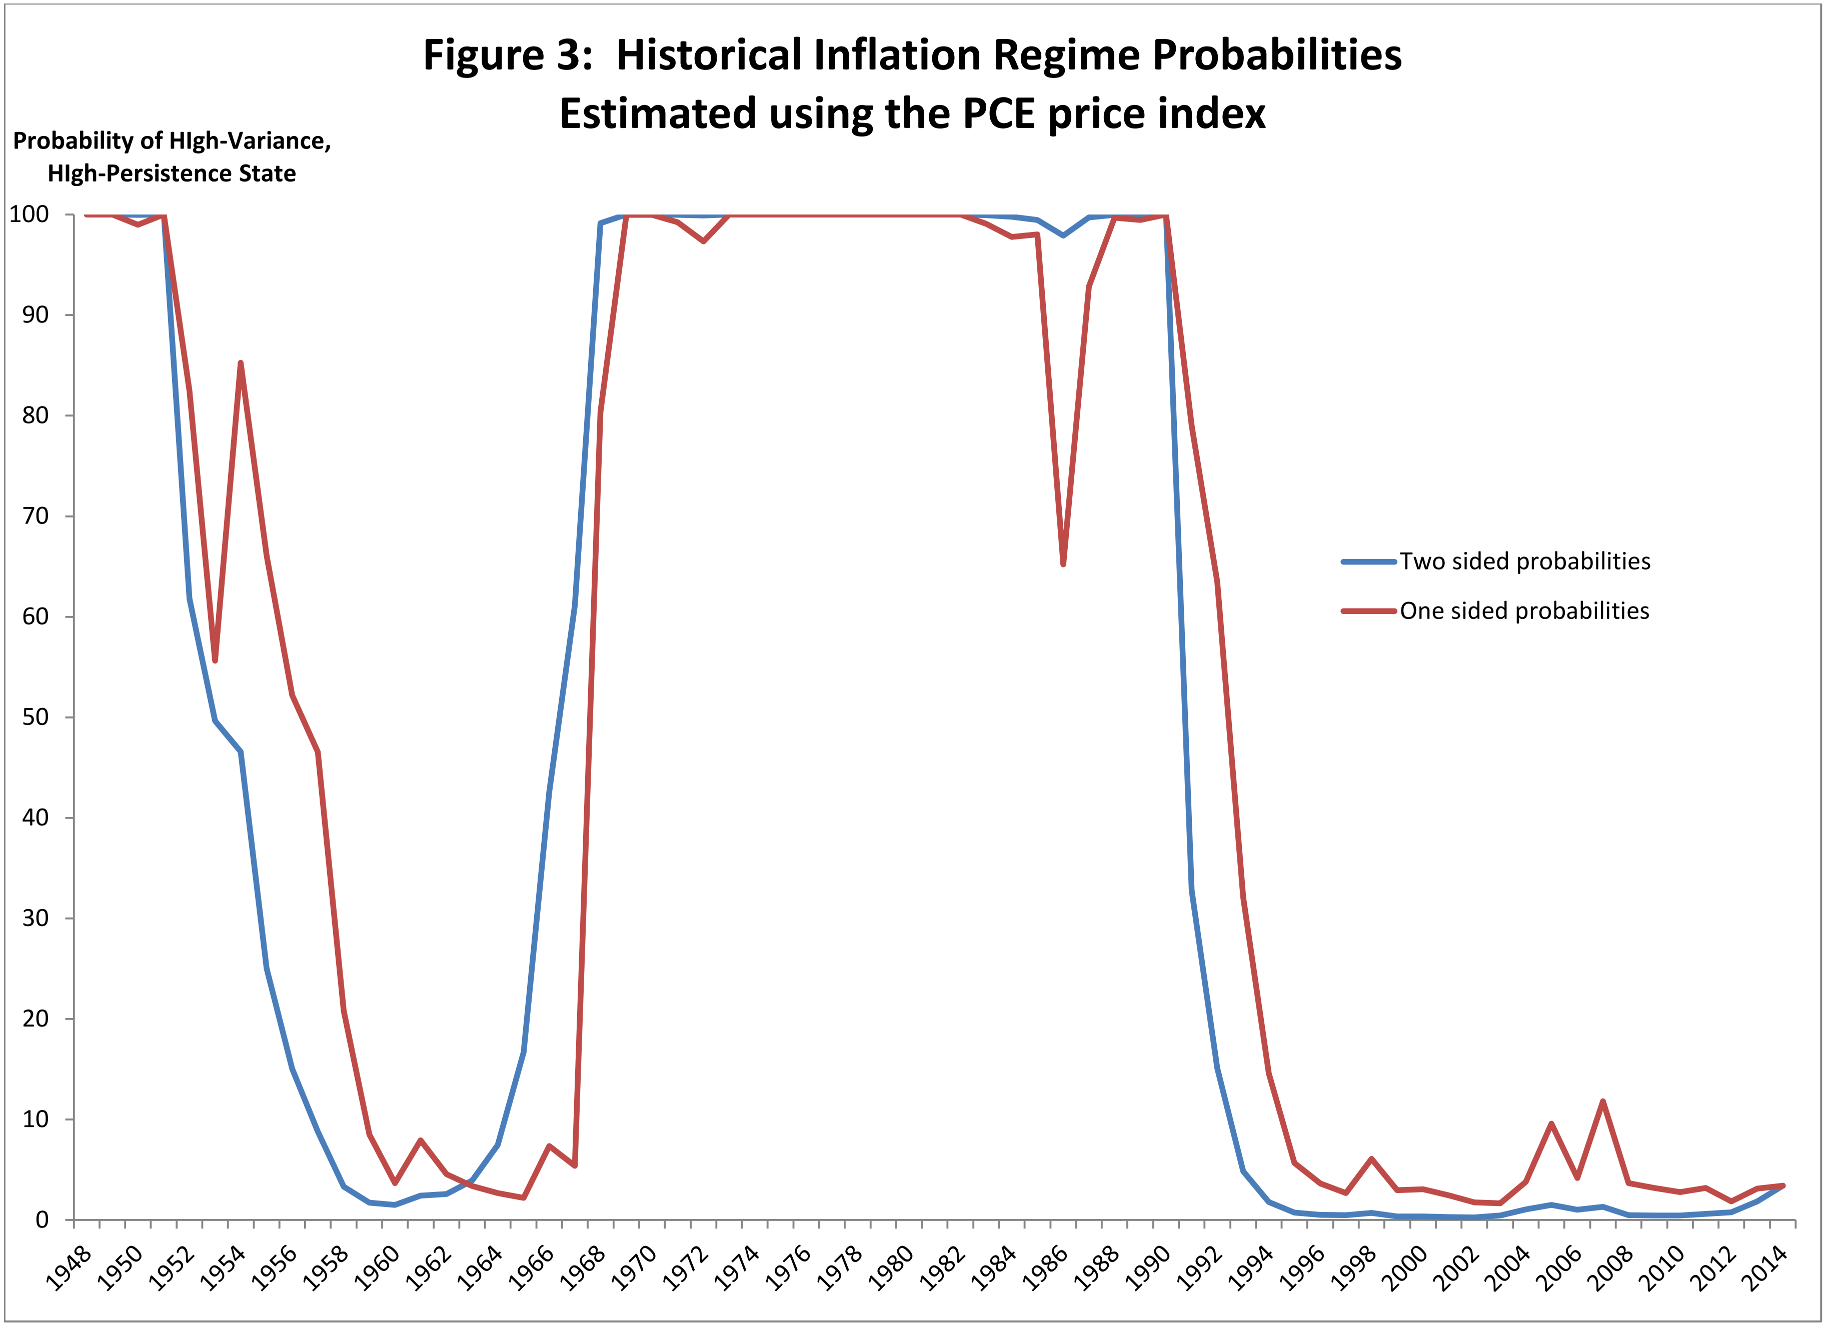

Figure 3: Historical Inflation Regime Probabilities
Estimated using the PCE price index
Probability of HIgh-Variance,
HIgh-Persistence State
100
90
80
70
60
50
40
30
20
10
0
1948
1950
1952
1954
1956
1958
1960
1962
1964
1966
1968
1970
1972
1974
1976
1978
1980
1982
1984
1986
1988
1990
1992
1994
1996
1998
2000
2002
2004
2006
2008
2010
2012
2014
Two sided probabilities
One sided probabilities

Figure 4: Long Run (Next 10-Year Average) CPI Inflation Forecasts

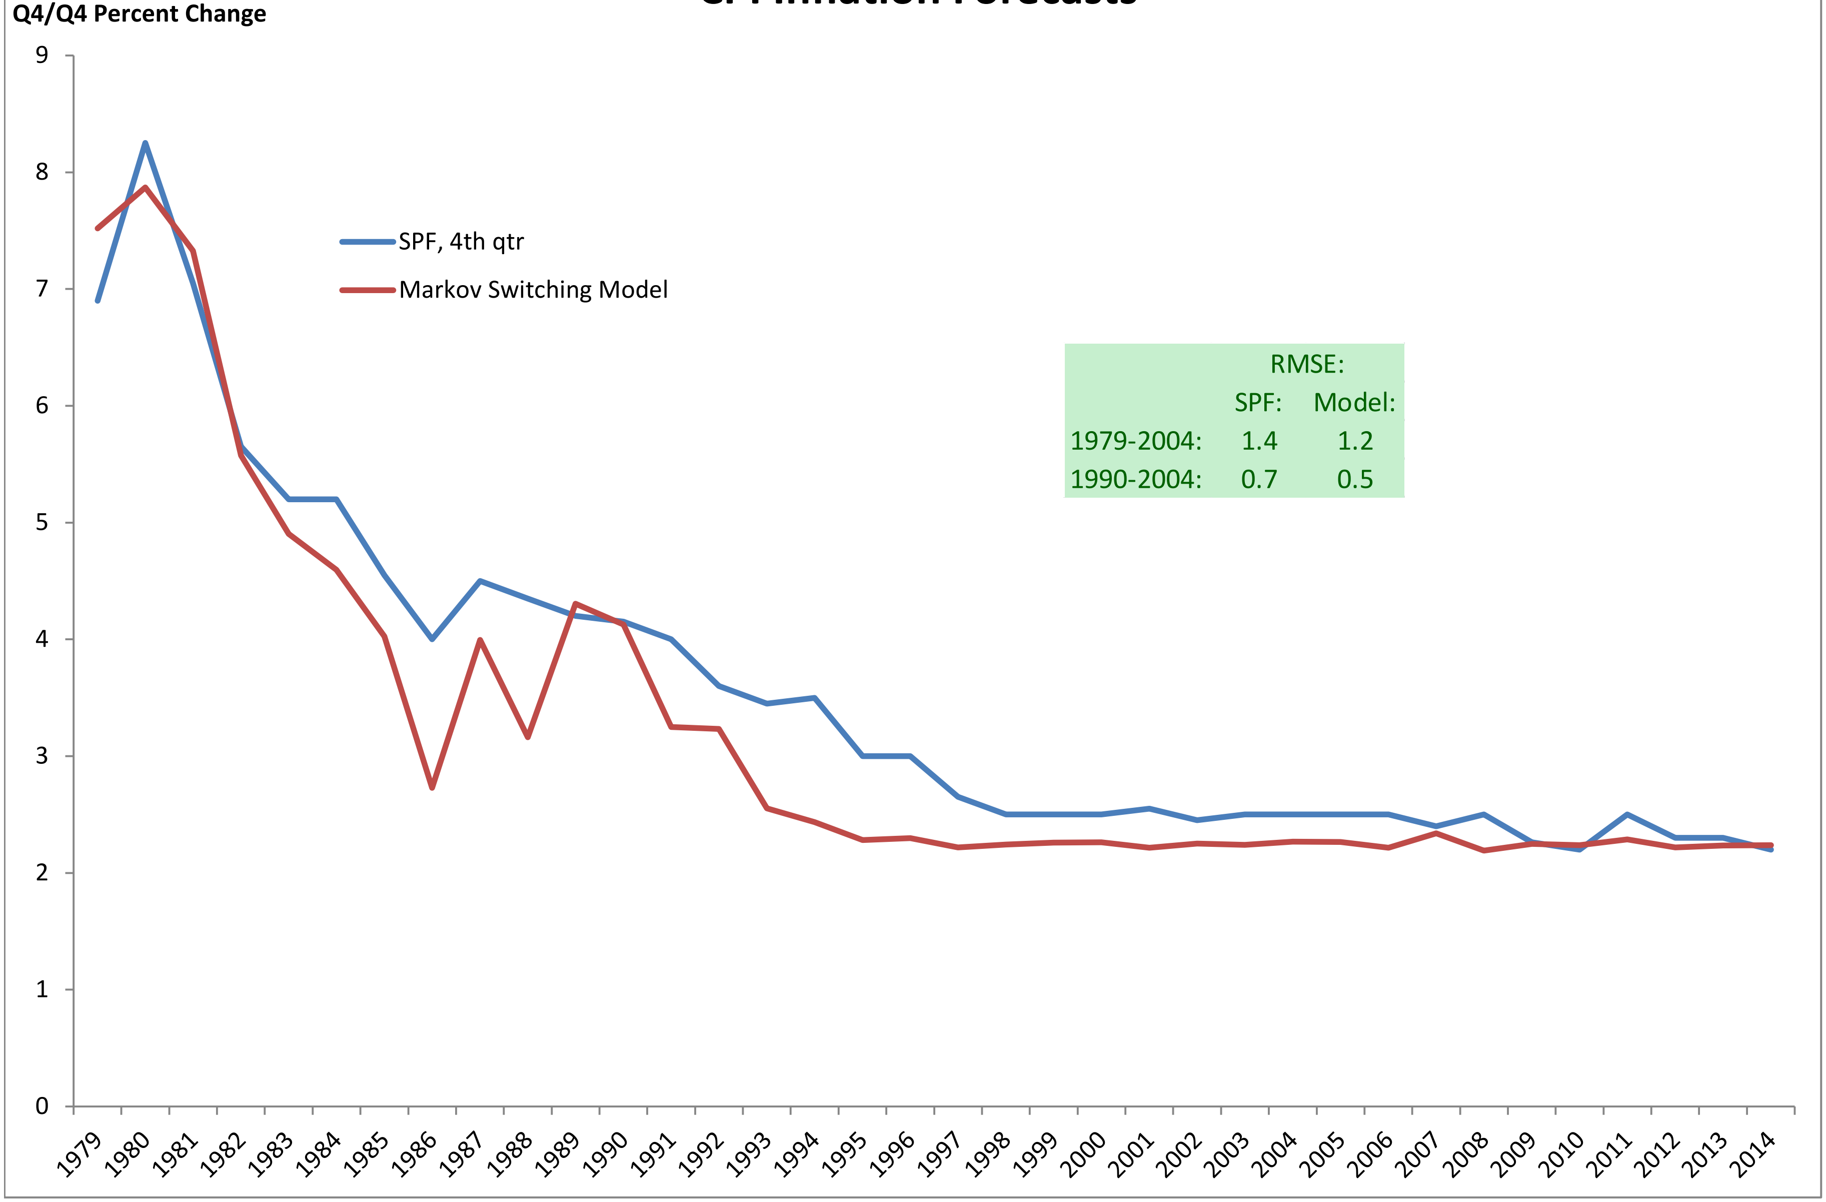

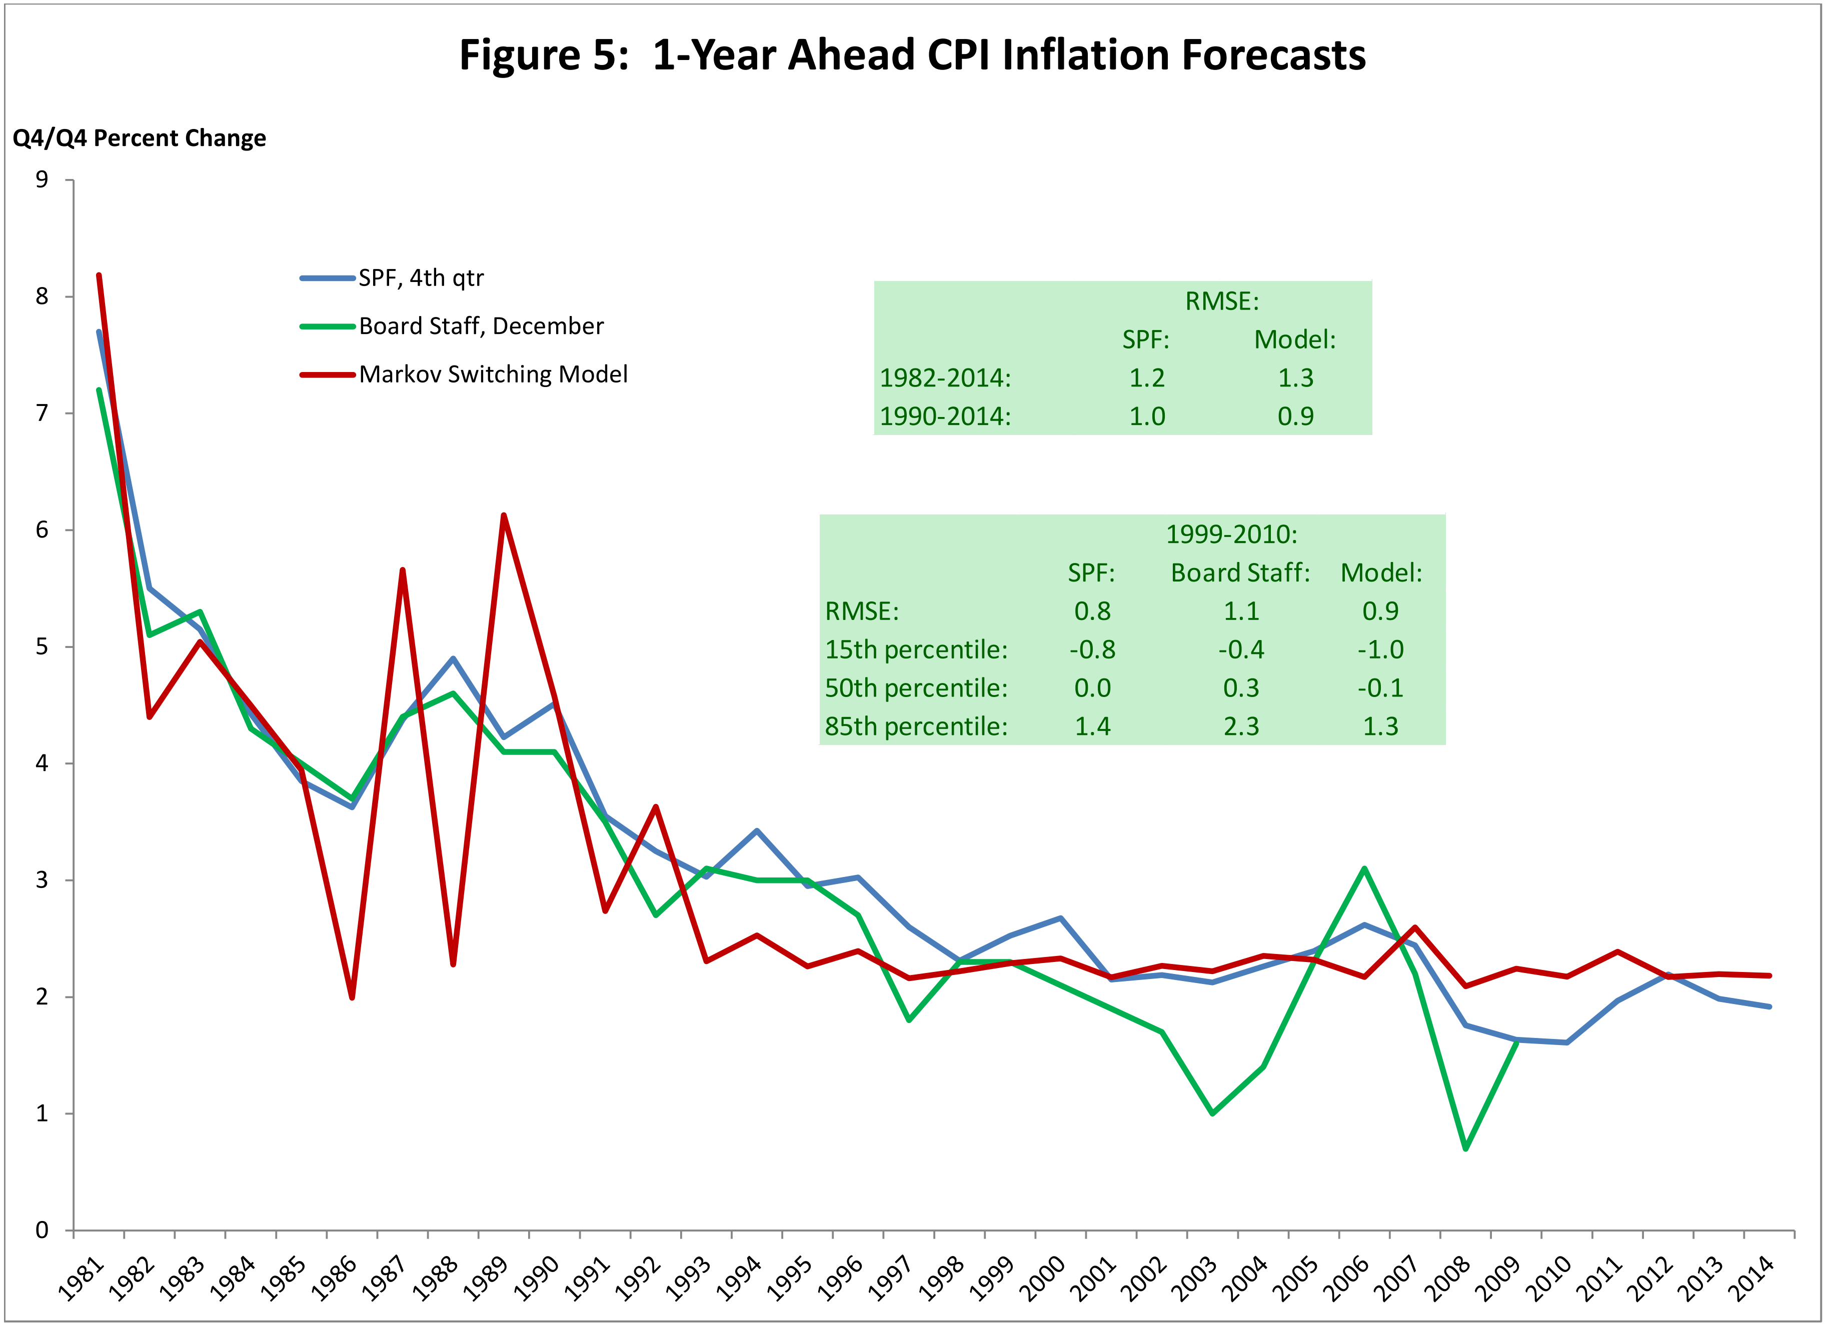
Figure 5: 1-Year Ahead CPI Inflation Forecasts
Q4/Q4 Percent Change
9
8
7
6
5
4
3
2
1
0
SPF, 4th qtr
Board Staff, December
Markov Switching Model
RMSE:
SPF:
Model:
1982-2014:
1.2
1.3
1990-2014:
1.0
0.9
1999-2010:
SPF:
Board Staff:
Model:
RMSE:
0.8
1.1
0.9
15th percentile:
-0.8
-0.4
-1.0
50th percentile:
0.0
0.3
-0.1
85th percentile:
1.4
2.3
1.3
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014

Figure 6: 1-Year Ahead PCE Price Inflation Forecasts

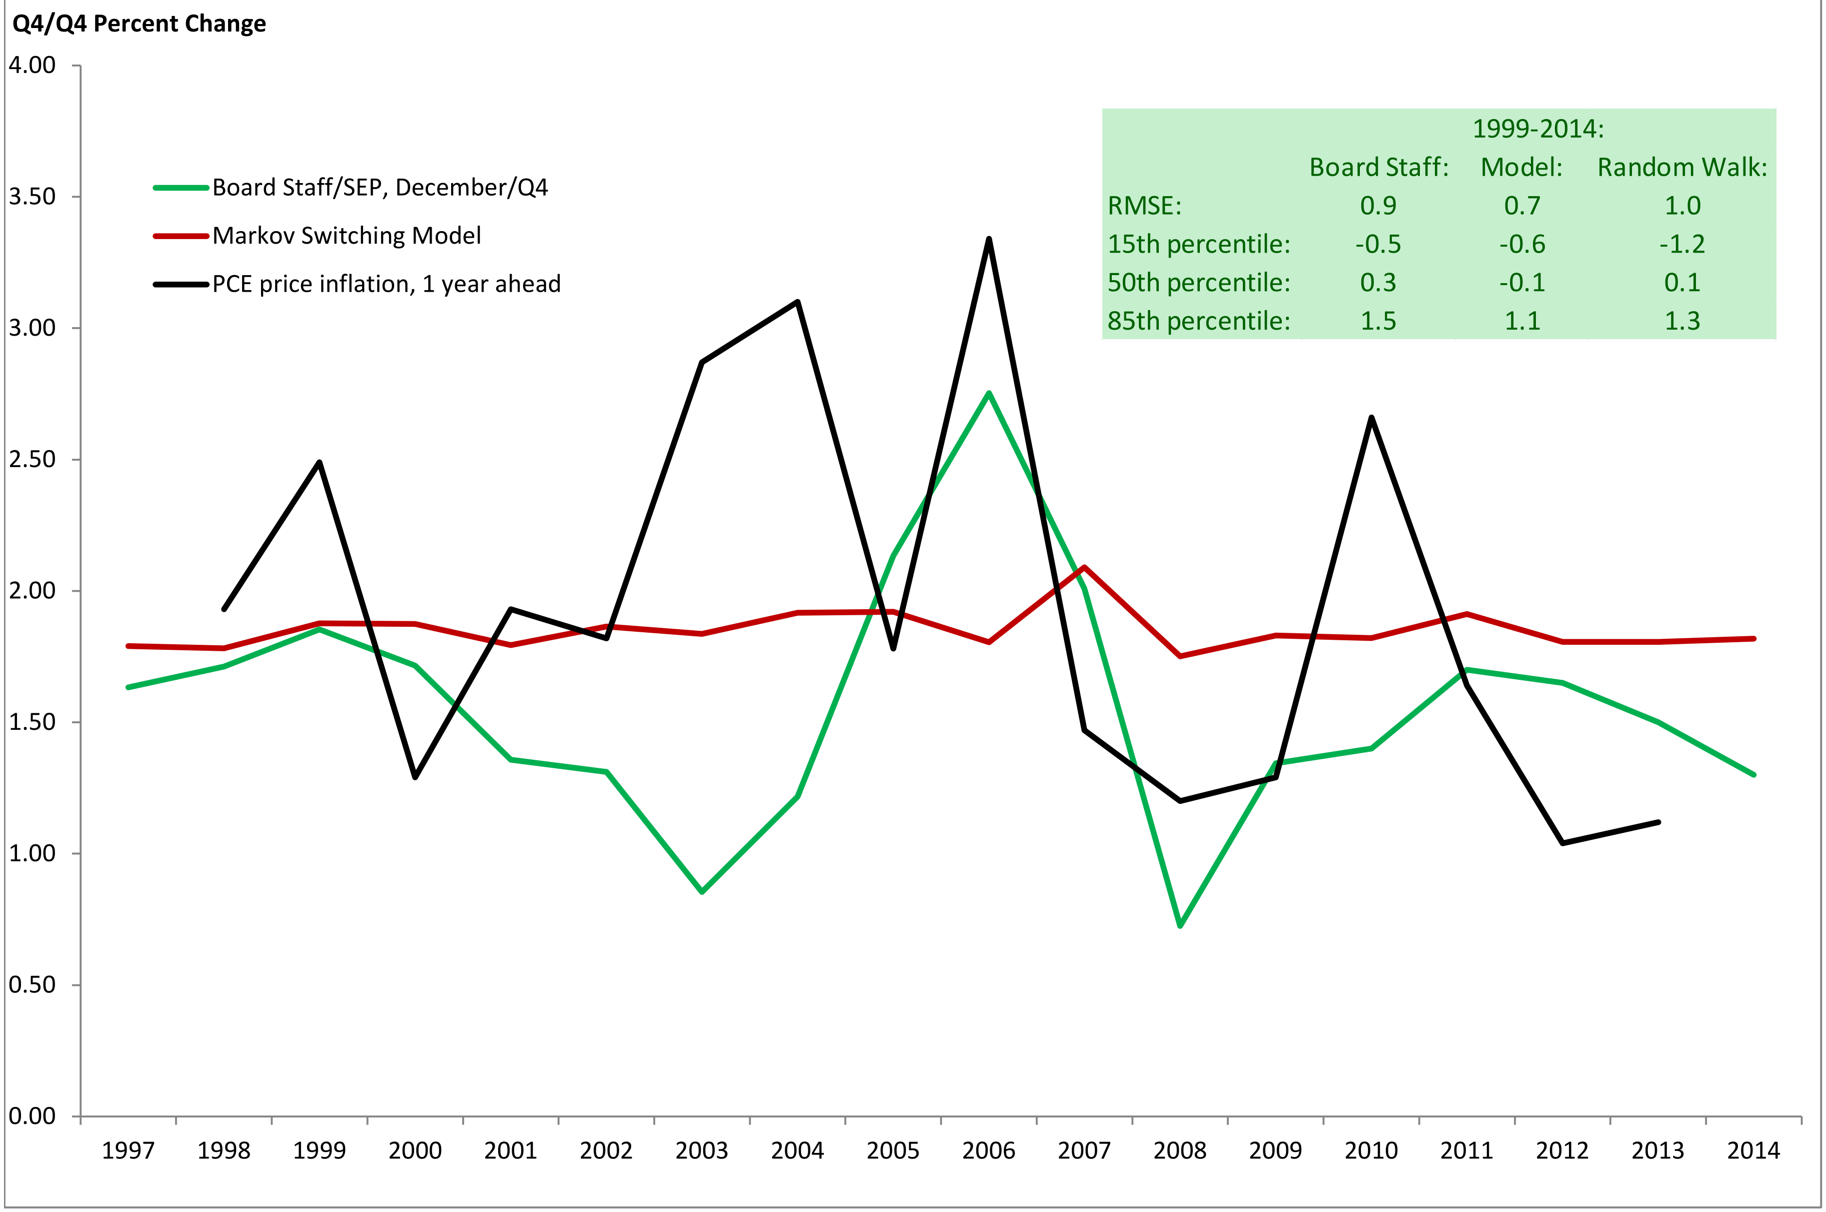

Figure 7: 2-Year Ahead PCE Price Inflation Forecasts