



Empirical Estimates of Changing Inflation Dynamics

Jeff Fuhrer, Giovanni Olivei, and Geoffrey M. B. Tootell

Abstract:

This paper provides an array of empirical evidence bearing on potentially important changes in the dynamics of U.S. inflation. We examine the overall performance of Phillips curves relative to some well-known benchmarks, the efficiency with which the Federal Reserve's Greenbook forecasts of inflation use real activity information, and shifts in the key determinants of the reduced-form "triangle model" of inflation. We develop a structural model-based interpretation of observed reduced-form shifts and conduct a reduced-form assessment of the relationship between core and headline measures of inflation, centering on the persistent "pass-through" of relative price changes into core and headline inflation measures, and a parallel exercise that examines the pass-through of key relative price changes into wage and compensation measures.

JEL Codes: E5, E31, E37

Jeff Fuhrer is an executive vice president and the director of research, Giovanni Olivei is a vice president and economist, and Geoff Tootell is a senior vice president and the deputy director of research, all at the Federal Reserve Bank of Boston. Their email addresses are Jeff.Fuhrer@bos.frb.org, Giovanni.Olivei@bos.frb.org, and Geoff.Tootell@bos.frb.org, respectively.

We would like to thank Todd Clark for helpful comments and suggestions. Denny Lie, Chad Stecher, and Rachel Stern provided excellent research assistance

The views expressed in this paper are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Boston or of the Federal Reserve System.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bos.frb.org/economic/wp/index.htm>.

This version: May 2009

Introduction

This paper provides a battery of empirical tests of Phillips curves. The tests examine the stability of various aspects of the Phillips curve, including: (1) its slope (the coefficient(s) on the real activity measure), (2) the influence of key relative prices that shift the curve, (3) the influence of past inflation dynamics on the formation of inflation expectations, and (4) general equilibrium influences that arise from the conduct of monetary policy and the “I-S” curve linking real activity to real interest rates. The paper uses both single-equation, reduced-form versions of the Phillips curve and multi-equation, constrained rational expectations models that incorporate a Phillips curve.

The paper suggests the following conclusions: (1) There have indeed been shifts in Phillips curve parameters. This conclusion is validated in all forms of the Phillips curve that we examine. (2) The shifts are most pronounced for the effect of the relative price of oil on core inflation. (3) The effect of the real activity variable may have diminished in recent decades. However, it is not zero, and we provide evidence suggesting that the Federal Reserve’s Greenbook forecast may have under-weighted the contribution of the unemployment rate in recent years. (4) The shifts appear to be concentrated in the early 1980s. The Phillips curve has been *relatively* stable over the past two decades. (5) Accounting for a change in parameters in the 1980s dramatically improves the out-of-sample forecasting performance of the Phillips curve. In fact, incorporating this shift in parameters overturns the results of Atkeson and Ohanian (2001). (6) In a structural model of the New Keynesian variety, the source of the shifts in the more reduced-form Phillips curves is a combination of a smaller effect of the relative price of oil in the Phillips curve, a smaller Phillips slope coefficient, and a downward shift in the interest sensitivity of the I-S curve, all of which occurred in the early 1980s. In more recent years, the share of backward-looking or indexing agents in the Phillips curve appears to have declined, although the identification of that effect is difficult in recent years.

The structure of the paper is as follows. Section I discusses the Phillips curve forecasting performance, especially relative to the Atkeson-Ohanian benchmark. Section

II examines the inflation-unemployment tradeoff implicit in recent Greenbook forecasts. Section III looks at reduced-form and structural estimates of shifts in key Phillips curve parameters. Because a key finding from Section III is a clear shift in the influence of the relative price of oil, Section IV looks in more detail at potential changes in the “pass-through” of key relative prices into core and headline inflation, and into wages. Section V offers some concluding remarks.

I . Why use a Phillips curve? Some new results on forecast accuracy

In this section we revisit the forecasting performance of the traditional backward-looking Phillips curve. The main reason to perform this exercise is to provide an interpretation of some of the negative findings in the literature (see Atkeson and Ohanian 2001, and Stock and Watson 2008) on the usefulness of backward-looking Phillips curves to forecast U.S. inflation. While increasingly complemented by more structural models of inflation, backward-looking Phillips curve specifications often play an important role in shaping the inflation outlook and the conduct of monetary policy. We show that the relationship between inflation and unemployment is tighter than some of the previous studies suggest, though there are periods when the Phillips curve relationship failed to materialize. We discuss changes in the relationship that have occurred over time, and in so doing we illustrate specifications that improve on the forecasting performance of well-known benchmarks.

Extant literature on out-of-sample inflation forecasting casts doubt on the usefulness of backward-looking Phillips curves for forecasting U.S. inflation in the post-1984 period. While several studies have reached more nuanced conclusions, the essence of Atkeson and Ohanian’s (2001) results about the inability of Phillips curve specifications to improve upon a simple univariate benchmark (whereby the forecast of inflation in the next four quarters is the value of four-quarter inflation today) has proved difficult to overturn. In the most comprehensive study to date, Stock and Watson (2008) corroborate Atkeson and Ohanian’s findings, with the important qualification that when

the unemployment rate is sufficiently different from the NAIRU, Phillips curve specifications improve substantially upon a univariate benchmark.

To illustrate some important issues concerning the usefulness of traditional backward-looking Phillips curves in forecasting inflation, we start by comparing the in-sample versus the out-of-sample performance of the Phillips curve. We show that there is a deterioration in the performance of the Phillips curve when moving from in-sample to out-of-sample evidence. There are several potential reasons for this finding. For example, it has been noted that out-of-sample tests have low power in small samples.¹ Another potential explanation is model instability. Shifts in the parameters over time may weaken the out-of-sample performance of the Phillips curve. We provide evidence that parameter instability is likely to have played a role in the weakening of the out-of-sample relative to the in-sample evidence over the period 1984 to 2007. Out-of-sample forecasts were generally worse than in-sample forecasts in the first part of the sample, but performed better in the latter part of the sample. This pattern in the relative performance of in-sample versus out-of-sample forecasts suggests a shift in parameters in the early 1980s. We then show how the out-of-sample performance of the Phillips curve can be improved by taking into account potential changes in the Phillips curve's parameters.

The backward-looking Phillips curve specification we consider is standard in the literature. Current inflation is a function of past inflation, the unemployment rate, and other control variables (supply shocks) that act as shifters in the inflation-unemployment rate relationship. Let P_t denote the average price index in quarter t . Then the quarterly rate of inflation at an annual rate, π_t , can be written as $\pi_t = 400 * \ln(P_t / P_{t-1})$. Four-quarter inflation, that is, the percentage growth in prices over four quarters, is given by $\pi_t^4 = 100 * \ln(P_t / P_{t-4})$. We consider inflation forecasts at the four-quarter horizon, which is widely studied in the literature and of interest to policymakers. As a result, our Phillips curve specification takes the form:

¹ See Inoue and Kilian (2004). For evidence on the low power of out-of-sample tests in backward-looking Phillips curve specifications, see Clark and McCracken (2006).

$$(1) \quad \pi_{t+4}^4 - \pi_t = \alpha + \gamma(L)\Delta\pi_t + \beta(L)u_t + \delta(L)\mathbf{x}_t + \varepsilon_t, \quad ,$$

where α is a constant, $\gamma(L)$, $\beta(L)$, and $\delta(L)$ are lag polynomials written in terms of the lag operator L , and $\Delta\pi_t$ is the first difference in inflation. The variable u_t denotes the unemployment rate, \mathbf{x}_t is a vector of supply shock variables, and ε_t is an error term. As written, the specification imposes the condition that the coefficients on quarterly inflation π sum to one. These time t and lagged quarterly inflation terms proxy for expectations of inflation over the next four quarters—a key feature of the backward-looking Phillips curve.

We compare the four-quarter-ahead forecast of inflation using the Phillips curve specification in equation (1) with the Atkeson-Ohanian benchmark, whereby expected inflation over the next four quarters, π_{t+4}^4 , is equal to inflation over the previous four quarters, π_t^4 , plus noise:

$$(2) \quad \pi_{t+4}^4 = \pi_t^4 + \eta_t \quad ,$$

where η_t is an i.i.d. error term. This random-walk benchmark is not as naïve as it first appears. Stock and Watson (2007) document that from 1984 on, the univariate process for inflation, π_t , is well described by an IMA(1,1) process. Over this period, the estimated coefficient for the moving average term in the IMA (1,1) representation is such that at the four-quarter (or longer) horizon, the Atkeson-Ohanian benchmark well approximates a forecast made with the IMA(1,1) process. As such, the Atkeson-Ohanian benchmark can be thought of as an effective shortcut to a non-naïve univariate representation of the inflation process.

We consider two measures of prices, the core CPI and the core PCE indexes. Core measures of consumer prices are of particular interest to policymakers. Also, by

mitigating the role of energy and food price shocks, a Phillips curve relationship specified on core inflation measures should provide a better assessment of the role of economic activity in generating common price movements. The civilian unemployment rate (16 yr. +) is specified either in levels or as a deviation from a time-varying NAIRU. Specifying the unemployment rate in levels is equivalent to assuming a constant NAIRU in equation (1) over the chosen estimation period. We use the CBO measure as our time-varying estimate of the NAIRU.² As supply shock controls, we include in (1) lags of the change in the relative price of oil and lags of the change in the relative price of non-oil imports.

We use the root mean squared forecast error (RMSE) as our metric to evaluate each forecast. The RMSE over the period t_1 to t_2 is

$$RMSE_{t_1, t_2} = \sqrt{\frac{1}{t_2 - t_1 - 1} \sum_{t=t_1}^{t_2} (\pi_{t+4}^4 - \pi_{t+4|t}^4)^2},$$

where $\pi_{t+4|t}^4$ is the forecast of π_{t+4}^4 made at time t .

Table I.1 compares root mean squared errors for four-quarter inflation forecasts for core CPI and core PCE inflation, respectively. The table provides the RMSE for the Atkeson-Ohanian forecast benchmark, and the RMSE from different Phillips curve specifications relative to the Atkeson-Ohanian benchmark. The sample period is 1984 to 2007. We consider Phillips curve specifications with either a time-varying NAIRU or a constant NAIRU, and with or without supply shocks. We also distinguish between in-sample forecasts and (pseudo) out-of-sample forecasts. The in-sample forecasts are obtained from estimating a Phillips curve over the period 1984:Q1 to 2007:Q4. The out-of-sample forecasts, instead, simulate a real-time forecasting exercise.³ This involves

² We experimented with other estimates of the NAIRU, constructed as either one-sided or two-sided filters of the unemployment rate. The results are not sensitive to the chosen measure.

³ To fully simulate a real-time forecasting exercise, pseudo out-of-sample forecasting should use real-time data. Here, as in some previous studies, we use final data. This, among other things, reduces the measurement error bias that a truly real-time exercise entails.

estimating a Phillips curve up to date t , making the four-quarter ahead forecast for date $t + 4$, then moving forward to date $t + 1$ and repeating the estimation to generate the next four-quarter ahead forecast. This process is repeated over the entire sample period, and at each point in time the Phillips curve is estimated over a rolling window of 40 quarters. This means that when we generate the first forecast in the sample, that is, the forecast for inflation over the four quarters 1984:Q1 to 1984:Q4, we estimate a Phillips curve over the period 1973:Q4 to 1983:Q4. Thus, in the earlier part of the sample, the out-of-sample forecasts use data that pre-date 1984 when estimating the Phillips curve. These data are not used when generating the in-sample forecasts, which are based on a Phillips curve estimated over the forecasting sample 1984 to 2007. For simplicity, the Phillips curve specification we estimate in the out-of-sample exercise at each point in time always maintains the same lag lengths. Specifically, we use the same lag lengths as in the Phillips curve specification estimated over the full forecasting sample.

The results in the table show that the in-sample Phillips curve forecasts perform well relative to the Atkeson-Ohanian benchmark. The Phillips curve RMSEs are 13 to 42 percent below the Atkeson-Ohanian benchmark, depending on the particular specification and the inflation measure used. The performance of the Phillips curve, however, deteriorates when we consider the out-of-sample forecasting exercise. The deterioration occurs across all different specifications and measures of inflation, but is much more pronounced for core CPI inflation than for core PCE inflation. For the core CPI inflation measure, the RMSE is higher than the Atkeson-Ohanian benchmark across all specifications. It should also be noted that estimating a Phillips curve with the same lag lengths as the in the in-sample specification might allow the out-of-sample exercise to understate the deterioration in the RMSE relative to the in-sample exercise. This is so because in real time the forecaster also faces model specification uncertainty, which we may have artificially reduced by imposing the (best fitting) in-sample specification.

In sum, while the in-sample performance of the Phillips curve in the post 1984 sample is good relative to the Atkeson-Ohanian benchmark, the results for the out-of-sample forecasting exercise are in line with the existing literature. These latter results

show that the performance of the Phillips curve over the period 1984 to 2007 relative to the Atkeson-Ohanian benchmark is, at best, mixed. Our findings of a deterioration in the out-of-sample vis-à-vis the in-sample performance of the Phillips curve are also similar to the findings in Clark and McCracken (2006).

What accounts for the out-of-sample forecasting performance deterioration of the Phillips curve? A possible answer can be gleaned by comparing the evolution of the in-sample RMSE versus the out-of sample RMSE over time. This is shown in Figures I.1a and I.1b for core CPI and core PCE inflation, respectively. At each point in time, the RMSE is computed over the most recent eight quarters. From the figures, it is evident that the out-of-sample forecasts are worse than the in-sample forecasts in the early part of the sample, that is, from 1984 to the early 1990s. In the most recent years, the out-of sample forecasts have been as good as—when not better than—the in-sample forecasts. Note that data from the early 1980s are used in estimating Phillips curves for the out-of sample forecasts that cover the period 1984 to the early 1990s but that these observations do not enter in the later period, as the window over which Phillips curves are estimated in the out-of-sample forecasting exercise is 40 quarters. This pattern in the RMSE suggests that parameter instability in the early 1980s is likely to have adversely affected the out-of-sample forecasting performance of the Phillips curve. Later in the paper, we provide formal tests on parameter instability, which, overall, confirm this view. The low power of out-of sample forecasting relative to in-sample forecasting could still play some role, but this explanation seems somewhat harder to reconcile with the way in which the out-of-sample forecast has evolved over time relative to the in-sample forecast.

An important question is what features of the Phillips curve could be responsible for parameter instability. Is the instability in the specification coming from the lags of inflation, the unemployment rate, or the supply shocks? Again, we will provide formal tests later, but some insight can be gained by looking at the evolution of the RMSE in the out-of-sample forecasts relative to the in-sample forecast when only some coefficients in the Phillips curve specification are estimated in the out-of-sample exercise, while the

other coefficients are kept fixed at the values estimated in the in-sample exercise. This is shown in Figures I.2a and I.2b for core CPI and core PCE inflation, respectively. The figures depict the RMSE of the in-sample forecasts and the RMSE of out-of-sample forecasts when only the subset of coefficients pertaining in turn to the lags of inflation, the unemployment rate, and the supply shocks are estimated, with the remaining coefficients set equal to the full-sample (1984 to 2007) estimates. The figures show that instability in the supply shocks coefficients at the very beginning of the sample is likely to have played a role. There is also some evidence that a shift in the parameters on lagged inflation contributed to the deterioration of out-of-sample forecasts. The evidence on shifts in the unemployment rate coefficients is somewhat weaker, though earlier in the sample the RMSE for core CPI inflation when only the unemployment rate coefficients are estimated is generally above the RMSE for the in-sample exercise. We will show that these findings align well with more formal tests presented later in the paper.

So far we have shown, as in previous literature, that out-of-sample forecasts of inflation using Phillips curves are not always better than simple univariate benchmarks. The changing dynamics of inflation make reliance on the Phillips curve difficult, but before discarding the framework entirely one should first investigate whether Phillips curve specifications that account for shifts in dynamics are useful for forecasting inflation. This avenue of research is still in its infancy. Stock and Watson (2007, 2008) illustrate a simple univariate model for inflation where parameters adapt to account for breaks. The model has appeal from an economic standpoint in that it decomposes inflation into a permanent and a transitory component. The model's forecasting performance is similar to the Atkeson-Ohanian benchmark over the sample we consider, and we are not aware of any work that checks whether the performance of the model improves once activity variables are included.

In principle, there are several ways to account for shifts in inflation dynamics. A simple way that preserves the standard features of the Phillips curve framework is a time-varying coefficients version of equation (1), which we write as follows:

$$(1') \quad \pi_{t+4}^4 - \pi_t = \gamma_t(L)\Delta\pi_t + \beta_t(L)(u_t - \text{nairu}_t) + \varepsilon_t, \quad ,$$

where we take the unemployment rate as a deviation from the NAIRU, so that the constant term can be dropped from the specification. Here, the coefficients on the lag polynomials $\gamma_t(L)$ and $\beta_t(L)$ are time-varying and follow random walk processes. In order to keep the specification relatively simple, we omit supply shocks.

The time-varying coefficients Phillips curve in (1'), estimated via standard Kalman-filter techniques, seems to capture inflation dynamics fairly well in out-of-sample forecasts. Table I.2 compares the out-of-sample performance of the Phillips curve (1') relative to the Atkeson-Ohanian benchmark for core CPI and core PCE inflation, respectively. As before, equation (1') is estimated over a rolling window of 40 quarters. We allow for only small time variation in the coefficients, but even so the RMSEs for both core CPI and core PCE inflation are more than 50 percent lower than in the Atkeson-Ohanian benchmark. Figures I.3a and I.3b depict the evolution of the RMSE for the Atkeson-Ohanian benchmark and for the Phillips curve specification (1'). The RMSE for the Phillips curve forecast is generally below the RMSE for the Atkeson-Ohanian random walk forecast.

The above findings illustrate that it is possible to specify processes for inflation that, by taking into account changing inflation dynamics, improve noticeably upon widely used univariate benchmarks. In terms of a Phillips curve framework, however, we still need to assess the importance of including the unemployment rate gap in (1') vis-à-vis specifying a relationship, as shown below, in which inflation follows the time-varying univariate autoregressive process:

$$(1'') \quad \pi_{t+4}^4 - \pi_t = \gamma_t(L)\Delta\pi_t + \varepsilon_t, \quad ,$$

where, as before, the time-varying coefficients on quarterly inflation are constrained to

sum to unity. The out-of-sample RMSEs of the Phillips curve specification (1') relative to the RMSEs of the univariate process (1'') when considering core CPI and core PCE inflation are 0.76 and 0.83, respectively. The evolution of the RMSE of the univariate process (1'') versus the RMSE of the Phillips curve specification (1') is also depicted in Figures I.4a and I.4b. The figures show that the RMSEs of the Phillips curve specification tend to be below the RMSEs of the simple autoregressive specification. Still, over some periods the forecasting performance of the Phillips curve does not differ from the forecasting performance of the time-varying univariate autoregressive process. This is true for the late 1990s, and for the most recent period (in this last instance, more so for core PCE than for core CPI inflation). These episodes highlight that not all is well with the simple backward-looking Phillips curve. For example, from mid-2003 until mid-2005, the Phillips curve specifications we have considered in this section would have predicted a fall in inflation, as the unemployment rate was relatively far away from the NAIRU. Contrary to the Phillips curve prediction, inflation picked up over this period. While some omitted control variables could in principle account for the rise in inflation, it is not clear what these variables are.

To summarize, in this section we have shown that the forecasting performance of the backward-looking Phillips curve is not as poor as some studies indicate. We have highlighted, however, that to improve noticeably over well-known univariate benchmarks one has to explicitly take into account changes in inflation dynamics for out-of-sample Phillips curve forecasts of inflation. The most important changes in inflation dynamics are likely to have occurred in the early 1980s, as will be shown more formally later in the paper. The best way to accommodate changing inflation dynamics in the Phillips curve specification, however, is still an open question that deserves more study. It is likely, though, that even with improved specification some episodes in the post-1984 sample will remain hard to explain in the light of an inflation-unemployment tradeoff.

II. The inflation-unemployment tradeoff: A view from the Greenbook

Even when one adopts a Phillips framework to model inflation dynamics, considerable debate remains about the size of the unemployment-inflation tradeoff. This link from the nominal to the real side of the economy is obviously crucial, in that it informs the conduct of monetary policy in the pursuit of its dual mandate. Is there evidence that this tradeoff has changed over time? Figure II.1 shows the relationship between the change in core PCE inflation and the unemployment rate gap over the period 1994:Q1 to 2007:Q4 in the context of a baseline Phillips curve specification such as (1). Note that the two series are partialling out the effect of other right-hand-side variables in the Phillips curve specification, so that regressing one series on the other gives an estimate of the slope of the Phillips curve. This relationship, estimated over the period 1984 to 2007, produces a statistically significant and economically meaningful slope for the unemployment rate gap. However, over the more recent period depicted in the figure, the estimated slope of the Phillips curve would be close to zero. The figure shows that the correlation between inflation and unemployment is generally negative but for the period 2003:H2 to 2005:H1, when inflation was rising despite a relatively high unemployment rate. In essence, this episode turns out to be influential when estimating a small slope for the Phillips curve over the period 1994:Q1 to 2007:Q4.

This point can be seen, too, in the context of Figures II.2a and II.2b, which show estimates of the slope of the Phillips curve obtained from the time-varying coefficients specification (1'') for core CPI and core PCE inflation, respectively. At each point in time, the figures depict the estimated slope from a rolling regression using 40 quarters of data, where each estimate is associated with the last observation in the rolling window. It is evident that there is a drop in the estimated slope when the period 2003:H2 to 2005:H1 starts entering into the rolling window. The conclusion to be drawn from this episode in terms of the economic relevance of the tradeoff between inflation and unemployment over the most recent 10 to 15 years is certainly debatable. But it is important to note that the failure of the inflation-unemployment tradeoff to materialize during this period

seems to be concentrated in time. The figures also show that, pre-2003:H2, there is some evidence of a decline in the slope when using core CPI as the inflation measure, but not when using core PCE. As of 2003:Q1, for both core CPI and core PCE inflation, the estimated slope of the Phillips curve was about -0.3 , which is economically meaningful.

While this reduced-form evidence is not clear-cut in pointing to a flattening of the Phillips curve, more structural evidence suggests that the slope of the Phillips curve may have lessened in the most recent decade. Later in the paper we provide some evidence to that effect. Even more important, Tetlow and Ironside (2007) document that the Federal Reserve Board's workhorse macroeconomic model, FRB/US, has been modified over the course of the most recent decade to encompass a flatter Phillips curve. From 1997 to 2003, the FRB/US estimate of the sacrifice ratio doubled, essentially implying that over this period the slope of the Phillips curve lessened by about one half.⁴

The belief that the Phillips curve has flattened appears to be informing the Federal Reserve's Greenbook inflation forecast as well. Figure II.3a shows the estimated slope of the inflation-unemployment tradeoff when we estimate a Phillips curve on Greenbook forecasts of core CPI inflation over time. Specifically, we estimate a relationship of the form:

$$(2) \quad \pi_{t+4,t}^{4,GB} - \pi_t^{RT} = \alpha + \gamma(L)\Delta\pi_t^{RT} + \beta(L)u_t^{RT} + \delta(L)\mathbf{x}_t^{RT} + \varepsilon_t,$$

where $\pi_{t+4,t}^{4,GB}$ is the Greenbook forecast of core CPI inflation over the next four quarters, and the variables indexed by "RT" denote real-time information available at the time the Greenbook forecast was made. Equation (2) is essentially the same as the Phillips curve specification (1), with the difference that (2) uses Greenbook forecasts of inflation instead of actual inflation. In equation (2), the unemployment rate is expressed in levels, and we use the relative change in crude oil prices as a control variable. The sample unit when

⁴ The slope of the Phillips curve determines the effect that deviations of unemployment (or output) from its equilibrium will have on inflation. The smaller is the coefficient, the more unemployment needs to move in order to effect a change in inflation. Thus a smaller coefficient entails a higher "sacrifice ratio" — the number of unemployment point-years required to reduce the inflation rate by one percentage point.

estimating (2) is given by one Greenbook forecast. In performing this exercise, we use a rolling window that includes 10 years of Greenbook forecasts.⁵ The results indicate a decline over time in the estimated slope of the Phillips curve. When forecasting inflation, the Greenbook was obviously using more information than the information included in equation (2). Omitting relevant information would bias the estimated slope whenever the omitted information is correlated with the unemployment rate. In this respect, a relevant omission in (2) could be a time-varying NAIRU. We lack information about the Greenbook's estimate of the NAIRU at the time of each Greenbook forecast. As a check on the plausibility of our specification, we back out from (2) an estimate of the NAIRU over the 10-year rolling window. This estimated NAIRU is depicted in Figure II.3b, and its evolution over time appears to be roughly in line with the sparse evidence we have on the Greenbook's evolving estimate of the NAIRU.

The change in the perceived size of the inflation-unemployment tradeoff that one can infer from Greenbook forecasts—in addition to the FRB/US model, which forms the basis for the Greenbook's alternative simulations—has important policy implications. With a flatter Phillips curve, inflation responds less to the unemployment rate, and thus the costs of bringing down inflation to the desired target in terms of lost employment are much higher.

It is not straightforward to gauge whether, *ex-post*, the Greenbook's perception that the inflation-unemployment tradeoff has flattened considerably was right. But it is possible to check whether the Greenbook has incorporated information about the unemployment rate efficiently into its inflation forecast. In this regard, it is important to note that over the period 1996 to 2003, Greenbook forecasts of the unemployment rate have been fairly accurate, exhibiting no glaring bias. As a result, Greenbook's inflation forecasts errors were, on average, not the result of errors in forecasting the unemployment rate. Having established this point, which is essential when examining the inflation forecasts error within the context of a Phillips curve framework, we then ask whether the Greenbook was incorporating information about the labor market—

⁵ We stop the analysis in 2003 because Greenbook data are made available to the public with a five-year lag.

available at the time the forecast was made—efficiently into the inflation outlook. For this purpose, we define the Greenbook inflation forecast error, e_t^4 to be:

$$(3) \quad e_t^4 = \pi_{t+4}^4 - \pi_{t+4,t}^{4,GB},$$

where, as before, $\pi_{t+4,t}^{4,GB}$ is the Greenbook forecast of core CPI inflation over the next four quarters, and π_{t+4}^4 is the corresponding actual value. We then regress the forecast error, e_t^4 , on information available at the time the forecast was made. This information is given by the most recent value of the unemployment rate, u_t^{RT} , and by core CPI inflation over the most recent quarters, $\pi_t^{4,RT}$. These values are in real time, and are taken from the same Greenbook from which the forecast $\pi_{t+4,t}^{4,GB}$ is taken.⁶ We consider past inflation in addition to the unemployment rate because these two variables are correlated. To the extent that past inflation helps to explain the Greenbook's inflation forecast errors, omission of past inflation from the regression would lead to bias in the estimated coefficient for the unemployment rate.

The regression results over the period 1996 to 2003 are as follows:

$$e_t^4 = -.233 * u_t^{RT} - .351 * \pi_t^{4,RT} + 1.80, \quad R^2 = .39, \quad N = 64,$$

(0.066) (0.152) (0.407)

where standard errors are in parentheses. Over this sample, there is evidence that both the current level of the unemployment rate and four-quarter inflation explain a considerable portion of the Greenbook forecast error. The regression indicates that when the current level of the unemployment rate was high, predicted inflation tended to be higher than actual inflation. Similarly, when the most recent four quarters of inflation were high, predicted inflation tended to be higher than actual inflation. This last result

⁶ Because we consider core CPI inflation and the civilian unemployment rate, real-time values are the same as current vintage values, aside from minor seasonal adjustments in the unemployment rate.

suggests that the Greenbook's perception of inflation persistence was likely too high. The negative coefficient on the unemployment rate is consistent with the Greenbook, maintaining a flatter Phillips curve than the one operating on actual data. Still, this is not the only possible explanation. A tendency to overstate the NAIRU during this period could also have produced such a result. While we cannot distinguish between the two hypotheses, we note that the regression estimates are not driven by the experience of the late 1990s, when estimates of the NAIRU were being revised downward. Figure II.4 shows actual core CPI inflation, the Greenbook forecast, and the Greenbook forecast adjusted ex-post to correct for the inefficiency with which lagged inflation and the lagged unemployment rate were factored into the forecast. The adjustment matters most in the early part of the sample and over the period 2002–2003. But this inefficiency in the Greenbook's inflation forecast, if still present, could become relevant again in the current context of a rising unemployment rate.

In sum, simple backward-looking Phillips curves provide mixed evidence of a change in the inflation-unemployment tradeoff. But more structural models often indicate a reduction in the slope of the Phillips curve over the past 10 to 15 years. Moreover, the Greenbook forecasts of inflation in the most recent years appear to hinge on a relatively flat Phillips curve. While the actual size of the tradeoff is still debatable, we have shown that, over the period 1996 to 2003, the Greenbook has not incorporated information about the unemployment rate efficiently into the inflation forecast. While there is more than one potential explanation for this finding, the finding is consistent with the Greenbook's positing a slope in the Phillips curve that is too flat.

III. Estimates of changing inflation dynamics

A. Reduced-form Phillips curve

In considering how inflation may behave over the monetary policy horizon, it is of interest to know the future effects on inflation of economic slack, of expectations, and

of important relative price (or other) shocks.⁷ One cannot presume that the influence of these determinants on inflation has remained constant across time; for example, there are good reasons to expect that the effect of oil shocks on inflation should have changed in recent decades. There has also been considerable debate about whether the slope of the Phillips curve has changed in recent decades, with the leading argument centering on a decline in the slope of the Phillips curve that would imply a significantly higher sacrifice ratio. Finally, most macroeconomic theories of inflation suggest that changes in the systematic response of the central bank to output and inflation will change the behavior of inflation. To the extent that monetary policy behaves differently now than it did in the 1960s and 1970s, the reduced-form dynamics of inflation now would be expected to differ from those in the earlier period.

In this section we examine more formally whether the influence of inflation’s key determinants, as viewed through the Phillips curve lens, has changed in recent decades. We consider two variants of the Phillips curve in this section, one a “traditional,” backward-looking and somewhat reduced-form representation already used in the previous sections, the other a “structural” rational expectations representation, in the tradition of “New Keynesian Phillips Curve” models (see, for example, Galí and Gertler 1999 and Christiano, Eichenbaum, and Evans 2005).

The traditional Phillips curve takes the form below. Inflation, π_t , depends on its own lags (proxying expectations and a variety of frictions in the price-setting process), an economic slack measure (here proxied by the unemployment rate gap), and identified relative-price shifters. As before, we write the equation in first-differenced form, which implicitly imposes the constraint that the sum of the lagged inflation coefficients equals unity:

$$(4) \quad \Delta\pi_t = \gamma(L)\Delta\pi_{t-1} + \beta(L)(u_{t-1} - nairu_{t-1}) + \delta^o(L)\Delta rp_{t-1}^o + \delta^{no}(L)\Delta rp_{t-1}^{no} + \varepsilon_t.$$

⁷ This taxonomy of effects loosely corresponds to Robert Gordon’s “triangle” model of inflation, see Gordon (1985)

As is conventional in such specifications, we include several lags of the unemployment rate, as well as lagged changes in the relative price of oil and non-oil imported goods.⁸

We examine two measures of core inflation, the CPI and PCE, and perform an array of tests of stability of all the key coefficients in the Phillips curve.⁹ The top panel of Table III.1a shows the results for the core PCE. The first five lines of results report the p -value for the test that the coefficients are stable across the break at the indicated time. We take a reported p -value of 0.05 or lower as statistical evidence of a change in the relationship. The “Great Moderation” break point is set to 1984:Q1, in accord with an array of empirical literature on the reduction in the variability of many macroeconomic time series at least up until the current (2008 and 2009) recession.¹⁰ The Greenspan era breakpoint is set to the beginning of Alan Greenspan’s tenure as Chairman of the Federal Reserve, in 1987:Q3, and the post-1994 break point is chosen as representing a period of accelerating productivity.^{11, 12}

For the core PCE, the “known” breakpoint tests suggest that the effect of the relative price of oil changed significantly, sometime in the mid-1980s. The tests develop no compelling evidence of a shift in the effect of inflation’s other key determinants. The unknown breakpoint tests uncover a breakpoint in 1982, consistent with the known

⁸ Alternative specifications that employ “trend inflation measures”—here the FRB/US measure of long-run inflation expectations—yield somewhat different results, depending on the specification. One specification allows both lagged inflation and trend inflation to enter, with coefficients summing to one. The weight on trend inflation was estimated at 0.2 for PCE inflation, and 0.06 for CPI.

⁹ See the complete data definitions for Sections III and IV in the data appendix at the end of the paper.

¹⁰ In the present exercise, the estimation sample ends in 2007:Q4, right at the onset of the recession (the NBER-dated peak is December 2007).

¹¹ In each case, we compute a Wald test that incorporates a heteroskedasticity and autocorrelation-consistent estimate of the parameter covariance matrix. The test regression includes a dummy times the variable or variables for which a shift is entertained. The test statistic is constructed as $W = (R\hat{\beta})'(R\hat{\Omega}R')^{-1}(R\hat{\beta})$, where R is constructed so as to impose the constraint $R\beta = 0$ for the dummied regressors of interest, and $\hat{\Omega}$ is the HAC estimate of the covariance matrix. The statistic is distributed as a chi-square with degrees of freedom equal to the number of coefficients constrained.

¹² Results for a Phillips curve that defines the unemployment gap as the difference between the civilian unemployment rate and the CBO’s estimate of the NAIRU produce the same qualitative results.

breakpoint test findings.¹³ Wald tests for coefficient stability on the individual coefficient groups develop evidence that the effect of changes in the relative price of oil shifted significantly in the 1980s. The sum of the coefficients on the relative price of oil (not reported) dropped from about 0.1 in the 1970s to approximately zero after the mid-1980s. There is also some evidence that the pattern of lags in the Phillips curve shifted in the 1980s. Because this term in the Phillips curve proxies for a number of possible inflation determinants—expectations, contracting lags, indexation—it is more difficult to interpret this shift. In the discussion of the structural results below, we will attempt to uncover more of the structure underlying these lags, by making explicit expectations, price-setting, and monetary policy actions.

For the CPI, Table III.1b, the results are similar. For the known breakpoints, the core CPI develops rejections of stability only for the relative price of oil, in this case at all three breakpoints. As with the core PCE, the coefficient sums that reflect the effect of oil prices on core inflation diminish until they are near zero as the sample progresses toward the present. The results for the unknown breakpoint are quite similar to those for the PCE, with a shift in 1980. The significance attached to the test for a shift in the lag coefficients is higher than for the PCE, while stability of the oil price effect is rejected with very high significance. As with the PCE results, the estimated sum of coefficients on oil prices drops from a bit below 0.1 for the earlier samples to zero for the past two decades.

In sum, the effects of relative oil prices have been a key source of instability in inflation dynamics, from the perspective of this traditional Phillips curve. The estimated impact of a change in oil prices is currently insignificantly different from zero, a statistically significant shift from the sizable impact in the 1970s and early 1980s. To be sure, there have likely been other shifts as well. In the CPI Phillips curve, we develop

¹³ The breakpoint test uses bootstrapped critical values constructed under the null of no break for the full-sample Phillips curve estimates. The “wild bootstrap” method we employ follows O’Reilly and Whelan (2005). Multiple breakpoint tests using asymptotic critical values suggest a breakpoint in the mid-1970s in addition to the break detected in the early- to mid-1980s.

some evidence that the unemployment effect has shifted down in the 1980s. Such a shift is less evident in the core PCE Phillips curve.

B. Structural Phillips curve

As suggested above, interpreting shifts in the backward-looking Phillips curve can be difficult, as the framework leaves implicit many structural features of price-setting. This is particularly true for the lagged inflation terms in the traditional Phillips curve. In addition, the framework is mute regarding the behavior of other aspects of the economy that bear on inflation—notably the systematic behavior of monetary policy and the transmission channel from monetary instruments to output to inflation. A simple example demonstrates how the combination of price-setting behavior, monetary policy, and the real economy jointly determine the behavior of inflation.

A stylized version of the backward-looking Phillips curve above makes inflation a function of lagged inflation and the output gap:

$$\pi_t = \pi_{t-1} + a\tilde{y}_t,$$

where \tilde{y}_t represents the deviation of output from potential. The output gap in turn is a function of the policy rate:

$$\tilde{y}_t = -bi_t,$$

and the policy rate i_t follows a stylized Taylor rule:

$$i_t = c\pi_t.$$

With this much-simplified economic structure, it is easy to see that inflation will follow the simple first-order autoregressive process:

$$\pi_t = A\pi_{t-1}$$

$$A = \frac{1}{1 + abc}$$

Inflation will more persistently deviate from its desired level (here assumed for simplicity to be zero)— A will be larger—when monetary policy targets inflation less aggressively (c is smaller), when the policy rate’s effects on output (b) are smaller, and when the effect of output on inflation (a) is smaller. This trivial model demonstrates simply and intuitively the way that the key features of the macroeconomy together determine the reduced-form behavior of inflation.

This section takes a somewhat more structural view on the potentially shifting dynamics of inflation by making many of these aspects explicit. Doing so entails some risk, as the resulting model is more restricted and in a sense places higher demands on the aggregate data in attempting to identify these additional economic behaviors.

The “hybrid” model—so-called because it includes both forward- and backward-looking elements—comprises three components, along the lines of the skeletal model sketched above:

- I. A Phillips curve with explicit expectations, in the “New Keynesian” Phillips curve tradition. It allows for a fraction of backward-looking or indexing price setters, motivating a lagged inflation term;
- II. An “I-S” curve, which one can derive from a canonical consumer optimization problem, which also allows for some backward-looking behavior, motivated in many cases by habit formation in consumer spending; and
- III. A monetary policy (Taylor 1993) rule, constructed along conventional lines.¹⁴

The model is summarized in the following set of equations:

¹⁴ We also explored an alternative version of the model, in which the output gap depends on a long-term real interest rate, defined according to a simple term structure relationship linking the long real rate to expected short-term real rates. The qualitative results presented in the tables below are unchanged; identification of the I-S curve was superior in the version presented here.

$$\begin{aligned}
(5) \quad \pi_t &= \mu\pi_{t-1} + (\beta - \mu)\pi_{t+1,t} + \tilde{\gamma}_t + \varepsilon_t \\
\tilde{y}_t &= \mu_y \tilde{y}_{t-1} + (\beta - \mu_y)\tilde{y}_{t+1,t} - a(i_t - \pi_{t+1,t} - \bar{\rho}) + v_t \\
i_t &= bi_{t-1} + (1-b)[a_\pi(\pi_t^4 - \bar{\pi}) + a_y \tilde{y}_t + \bar{\rho} + \bar{\pi}] + \eta_t
\end{aligned}$$

Because the overall behavior of inflation can be influenced by price-setting behavior, summarized in the top equation, or by the monetary transmission mechanism, summarized in the second equation, or by the behavior of monetary policy, summarized in the third, we look for breaks in the coefficients of all of these equations.¹⁵

The results in Table III.2 take a form similar to those in the preceding table for the backward-looking Phillips curve.¹⁶ For the core PCE, shown in the top panel, the results suggest clearly that the model overall has not been stable, either across known breakpoints, or looking at breakpoints detected using unknown breakpoint methodology.¹⁷ Somewhere in the 1980s, a very significant parameter shift is detected.

Interestingly, the tests suggest that the shift is *not* primarily due to significant shifts in the parameters of the forward-looking Phillips curve, but rather to a shift in the transmission mechanism (the I-S curve)—and perhaps to some change in monetary policy.

For the core CPI, the results are similarly clear. The data strongly reject stable coefficients for the entire model across the full sample. Both known and unknown breakpoint tests indicate a break somewhere in the 1980s. The source of the break appears to be the same as that for the PCE. The test for stability of the I-S curve parameters strongly rejects stability for both known breakpoint dates.

¹⁵ The tests reported are likelihood-ratio tests for the hypothesis that the estimated coefficients in the first sample are equal to those in the second. The likelihood ratio is computed by imposing the first-sample coefficients on the second sample, and taking the ratio of the unconstrained likelihood to the constrained likelihood for the second sample.

¹⁶ Results using an output gap based on the CBO estimate of potential output differ little from the results presented.

¹⁷ This method simply searches for the breakpoint associated with the largest likelihood ratio. We do not know of research that allows one to compute the critical values of this likelihood ratio test for systems of equations with rational expectations. Bai (1999) discusses critical values for a likelihood-ratio test for single-equation models.

Typically, such structural models do not account for the effects of key relative price shifts, as in the backward-looking models. In some cases, this is theoretically justified—particularly when inflation is linked to real marginal cost. In this case, marginal cost should capture the effect of any input costs on inflation, so the addition of energy or other input prices would be superfluous. In our case, inflation is linked to the output gap, as is conventional in many NKPC specifications. Here, there may be a theoretical and empirical role for relative prices. Thus, we consider a model that allows for the explicit effect of the relative price of oil on inflation, augmenting the forward-looking Phillips curve as follows:

$$\pi_t = \mu\pi_{t-1} + (\beta - \mu)\pi_{t+1,t} + \tilde{\gamma}_t + p_o\Delta r p_t^o + \varepsilon_t.$$

The results of the tests in this case are interesting, as they now uncover some shift in Phillips curve parameters.

Again, as shown in table III.3, for both core PCE and core CPI, the hypothesis of stability of all the coefficients is strongly rejected for both known and unknown breakpoints. But now, stability of the Phillips curve coefficients is rejected for both measures and both breakpoints. Because the only change to the Phillips curve is the addition of an oil price term, this suggests that the effect of the oil price on inflation, omitted in the first specification, is a significant source of instability in these estimates, echoing the results of the reduced-form Phillips curve models. Both the known and the unknown breakpoint tests point to a shift in structure in the early 1980s.

Figure III.1 provides a reasonably clear picture of the changes that appear to have occurred over the past 30 years. The figure displays the estimated value of a particular structural model parameter at the breakpoint date indicated on the horizontal axis. The solid line shows the parameter's estimated value over the first part of the (split) sample, and the dashed line shows the estimated value over the second part of the sample. For example, the top-left panel shows the estimates for μ , the backward-looking weight (or indexation parameter) in the Phillips curve. The value of the solid line for

1990:Q1 is the estimate of μ for the sample 1966:Q1–1989:Q4; the value of the dashed line for the same period is the estimate of μ for the sample 1990:Q1–2007:Q4.

The results that emerge from this figure should be of some interest to policymakers. First of all, most of the later sample estimates are lower than those of the early sample: Inflation lately appears somewhat less responsive to output gaps (γ), the effect of oil prices is smaller (p_o), and the interest-sensitivity of output is smaller. The transmission channel of monetary policy—the effect of interest rates on real activity, and correspondingly the effect of real activity on inflation—appears somewhat muted, with the change occurring somewhere in the early 1980s. Second, many of these key parameters are reasonably stable after that breakpoint. That is, the dashed lines, while they vary with normal sampling variation, are relatively stable over the past 25 years. Two notable exceptions are the policy response to inflation (a_π), to which we will return shortly, and the backward-looking component of inflation (μ).

Beginning with the top left panel, the backward-looking weight in inflation, sometimes referred to as the “intrinsic persistence” of inflation, is insignificantly higher in the earlier sample than in most of the later period. However, beginning in the early 1990s, its value drops from about 0.5 to 0.2 or lower. This decline could reflect a change in the price-setting behavior of economic agents. In conjunction with better-run monetary policy, this could explain the observations by some of a decline in reduced-form inflation persistence in recent years. It may also reflect the difficulty in assessing the inherent persistence in inflation during a period when inflation has ranged between one and three percent. Finally, it is important to note that the decline in this parameter is not robust across alternative model specifications.¹⁸

¹⁸ The inference here is somewhat complicated. An alternative version of the model in which the inflation equation is $\pi_t = \mu\pi_{t-1} + (1 - \mu)\pi_{t+1,t} + \tilde{y}_t + \varepsilon_t$ does not display a similar decline in μ . This seemingly small change in the term pre-multiplying expected inflation from $\beta \approx 1$ to 1 appears to have important implications for this aspect of inflation dynamics. Note also that the estimated inflation response falls below 1 in the early 1990s, an apparent violation of the so-called “Taylor principle.” However, in this model, with $\beta < 1$, inflation stabilizes independent of the response of monetary policy, a feature of this widely used class of models that is not well articulated in most treatments.

The response of inflation to output gaps is everywhere lower in the later samples, as shown in the top right-hand panel of the figure. This corroborates the widespread belief that the slope of the Phillips curve has become quite shallow, implying a large sacrifice ratio. The effect of the relative price of oil on inflation, shown in the middle left panel, is clearly lower in the later samples than the earlier, and is generally difficult to distinguish from zero for most of the samples that include data from 1985 and forward.

The responsiveness of output to the real interest rate, shown in the middle right panel, is clearly smaller in the recent data than in the earlier samples. In fact, it is difficult to estimate a non-zero interest elasticity, hinting at the general difficulties in identifying some of these structural equations in recent periods.

Apart from some early sample noise, the target rate of inflation, shown in the lower left panel, has behaved consistently with common wisdom about the Fed's inflation goal: It is clearly lower in the later samples, and appears to have declined somewhat in recent decades, with the current estimate hovering around 2 percent.¹⁹ The rise in the monetary policy response to inflation, shown in the bottom right panel, is striking. It arises from the brief period in the data from 1992 to 1995, during which inflation was declining, but the funds rate was rising from 3 to 6 percent. In this relatively simple model, that episode is interpreted as a modest decline in the inflation goal, coupled with a significant increase in the response of policy to inflation.^{20,21}

In sum, we find reasonably consistent evidence across both reduced-form and structural models that the influence of key relative prices (particularly oil) on inflation has significantly diminished in recent years. We find somewhat more mixed evidence

¹⁹ The overall sample begins in 1966. For breakpoints dated 1975 through 1980, we use 30 to 40 observations to identify 10 parameters. The inflation target is difficult to identify in samples that include the early 1980s, the non-borrowed reserves operating procedure period.

²⁰ Note that the vertical axis of the panel is truncated at five. The maximum response estimated in this period is about 10.

²¹ Not shown are estimates for the forward- and backward-looking weights in the I-S curve; these do not differ significantly between estimation periods, with the backward-looking component centering on about 0.6. In addition, the lagged interest rate term in the policy rule remains reasonably stable at around 0.8 throughout the entire sample period.

that the influence of real activity has diminished, although more generally that influence appears to have varied over time and is not as robust as one would like. In terms of structural explanations of reduced-form results, the shifts uncovered in the more reduced-form model appear to arise from shifts in several underlying structural parameters. The structural model displays instability in the parameters of both the Phillips curve and the I-S curve. Much of the parameter instability centers on the early- to mid-1980s, and for many parameters, estimates have remained reasonably stable since that time. The exceptions are the backward-looking component of the Phillips curve and the responsiveness of the policy rate to inflation, both of which appear to have shifted (downward and upward, respectively) in the early 1990s. We emphasize, however, that the identification of these structural features of the model is somewhat delicate and specification-dependent, particularly over the most recent period of relative macroeconomic stability.

To sum up, these results suggest that:

(1) The persistent influence of large oil price increases on both core and headline measures is likely quite small, but probably not zero;

(2) Real activity, while exhibiting a smaller influence than in the 1970s and early 1980s, remains an important determinant of inflation. When linked to the results of the Greenbook's inflation forecast, this suggests that a key risk at present may be the possibility of *under*-estimating the influence of weak real activity on inflation in coming quarters.

IV. The relationship between “core” and “headline” inflation: A reduced-form test for changing pass-through of important relative prices

One of the most consistent findings in the preceding section is that the influence on inflation of key relative prices—especially the relative price of oil—has declined markedly in recent decades. The magnitude and persistence of the effect of oil and food prices on trend inflation have received renewed attention because of the recent large

swings in these prices. In this section we take a closer look at the pass-through of food and energy prices into core inflation measures and into wages. A natural way to examine pass-through of food and energy prices into core inflation is to look at the relationship between core and headline inflation measures. Accordingly, this section uses a simple reduced-form regression linking core and headline inflation measures to assess the pass-through of these key relative prices into inflation. The section concludes by reporting results from an analogous regression test that assesses the pass-through of relative price fluctuations into wages.

A. Pass-through into prices

An important element of the monetary policy discussion over the past several years has been the extent to which significant relative price changes “pass through” into core measures of inflation. Put differently, an important empirical question is the degree to which relative price changes have had or will have persistent effects on inflation. For example, if surges in food and energy prices have only transient effects on core and headline inflation measures, then they may be of relatively little concern to policymakers. On the other hand, if such surges tend to become embedded in the inflation process, affecting wages and expectations and leading to persistent increases in key inflation measures, then policymakers should be concerned.²²

In the two previous sections, we explored the extent to which relative prices affect core inflation measures in constrained structural models. In both reduced-form Phillips curves, and in more structured New-Keynesian Phillips curves, we developed evidence of a shift in the relationship between oil prices and core inflation.

In this section, we present results from a relatively straightforward empirical test that may shed some light on the effects of relative prices on key inflation measures and

²² In this section, we put aside the issue of the structural mechanisms by which such relative price changes might have persistent effects on non-food and energy prices. To be sure, the causes would likely center on the extent to which monetary policy accommodates or is expected to accommodate such relative price increases, and whether price-setters’ expectations are well anchored, perhaps on their view of the policymaker’s implicit inflation goal. But a more complete description of how relative price changes might lead to persistent changes in inflation lies beyond the scope of this paper.

on whether such pass-through effects have changed in recent years. The overall or “headline” measures of inflation must, by definition, reflect the relative price changes for all the goods in the consumer basket, at least for the period in which the price changes occur. Whether such changes have more lasting effects on either headline or core inflation measures is an empirical question. Thus, a regression such as the following,

$$\pi_t^J = \sum_{i=1}^n a_i \pi_{t-i}^H + \sum_{i=1}^n b_i \pi_{t-i}^C + \text{controls} , \quad J = H, C ,$$

in which a core or headline measure of inflation is regressed on lags of both core and headline measures, can provide evidence on the extent to which relative price changes as reflected in the headline measure have persisted in headline measures or “passed through” into core measures of inflation.²³ To the extent that relative price effects persist, lags of the headline measure, which incorporate weighted effects of relative price changes, will influence subsequent core or headline inflation measures, and thus will receive greater weight in the regression.²⁴ In the results that follow, we regress both headline and core (excluding food and energy) inflation measures for the CPI and the PCE on lags of both core and headline inflation, controlling for a variety of other presumed influences on inflation, including the unemployment rate and changes in the relative prices of energy and non-oil imported goods. The rationale for including the latter controls is that they might pick up the predictable component of persistent shocks to these relative prices.

Figure IV.1 motivates an investigation into the extent of pass-through, and the extent to which it may have changed in recent years. In the years prior to 1985, both

²³ Rich and Steindel (2007) address the issue of which measures of core inflation best capture the trend or persistent component of inflation, a closely related topic. Their test differs from ours, but their results emphasize the difficulty in defining a reliable core measure of inflation.

²⁴ As more controls and constraints are added to the regression—such as measures of resource utilization, and constraints on the sum of the lagged inflation coefficients—the regression takes on the form of a conventional backward-looking Phillips curve.

headline and core inflation moved together in the presence of large shocks to the relative prices of energy and food. Following 1985, the core rate appears to respond less to such shocks and is thus noticeably less volatile. It also appears that core inflation serves as the anchor to which the headline measure will move in subsequent quarters.

Table IV.1 provides representative regression results for the test regression.²⁵ In the earlier part of the sample, it is clear that relative price changes as captured by the short-term movements in the headline inflation rate feed substantively and significantly into subsequent movements in both the core and the headline measures. From 1970 to 1985, the weight on the movements in headline inflation for subsequent movements in core inflation was 0.9 for the CPI and 0.7 for the PCE. Correspondingly, the influence of the lagged core inflation movements was much smaller and not statistically significant. In a sense, headline led core inflation around.

The picture changes quite substantially in the later periods. The weight on the headline inflation measure in the core inflation regression falls to 0.1 to 0.2, and is statistically indistinguishable from zero. In fact, perhaps most surprisingly, over the past 15 to 20 years, the weight on lagged headline inflation in the *headline* inflation regression has fallen dramatically in magnitude, and is no longer significantly different from zero. Not only have relative price effects stopped passing through into core, they apparently have had very little persistent effect on headline inflation.

One possible interpretation of the significance of headline inflation in the regressions in which headline inflation is the dependent variable is that the lagged headline inflation terms proxy for changes in key relative prices that were quite persistent, and thus persist in headline inflation measures. To control for this possibility, variants of the regressions in Table IV.1 are run that explicitly include lags of relative prices (oil, non-oil imported goods, and food). The qualitative conclusions from Table IV.2 are quite similar in these regressions, although the headline measure is no longer a significant predictor of the core inflation rates, even in the earlier samples.

²⁵ An array of other results, not presented here, confirm the qualitative conclusions throughout. Regardless of the specification, the constraints, and the controls, the influence of past headline inflation on core *and* on headline inflation has diminished to a level that is insignificantly different from zero in recent years.

Representative results are presented below. The weight on headline inflation remains small and statistically indistinguishable from zero in the later period. As in the previous results, the core has a weight of 0.8 or better, and is always statistically significant.

A key difference for these results is that both headline inflation measures now consistently respond to the lagged trajectory of the core measure in the more recent data, as shown in the rightmost columns of the second panel of the table. If this were not the case, as in Table IV.1, it would suggest a “de-coupling” of core and headline inflation, an interpretation we would be reluctant to accept at face value.

Overall, the evidence from this section suggests that currently it is headline inflation that converges to core inflation, and that the relative prices of food and energy leave little if any lasting imprint on the core inflation measures.²⁶ Of course, this is a reduced-form implication of the recent behavior of inflation, which must depend in part on the behavior of monetary policy. The implications of a model that explicitly articulates monetary policy were discussed in the previous section.

B. Pass-through into wages/compensation

Another channel for the potential pass-through of relative price changes to inflation is through wages. In principle, workers could demand compensation for the real wage losses suffered as the prices of key goods rise. The following regressions look for evidence of such a channel historically, again examining significant subsamples to see if any such relationship has changed over time.

The dependent variable in the regressions reported in Table IV.3 is a measure of the change in nominal compensation (the ECI or the non-farm compensation series), denoted Δw_t in the table. The form of the regression is based on the first-order condition for the use of labor input in a production process: Under certain assumptions, the real wage paid to labor should equal the marginal product of labor. The regression is run in

²⁶ An extensive specification search did not uncover any specifications for the PCE in which the headline had a sizable and significant effect on core inflation over the past two decades. We were able to find one specification for the CPI in which the headline was imputed a 0.2 weight with a significant coefficient.

first-difference form, with growth in the average productivity of labor ($\Delta prod_t$) proxying for the marginal product.²⁷ The inflation rate that feeds into the change in real wages is allowed to be either the headline (π^H) or core (π^C) inflation rate. A variety of controls, including the unemployment rate and the key relative prices from the regressions above, are allowed to enter as well:

$$\Delta w_t^J = \sum_{i=1}^n a_i \pi_{t-i}^H + \sum_{i=1}^n b_i \pi_{t-i}^C + \sum_{i=1}^2 c_i \Delta prod_{t-i} + \sum_{i=1}^2 d_i (u_{t-i} - nairu_{t-i}) + \sum_{i=1}^k e_i rp_{t-i}^o + \varepsilon_t$$

$J = ECI, NFC$

These regressions generally fit less well than the inflation-based regressions above, and the results are not as clear-cut.²⁸ However, one cannot reject the hypothesis that the relative price changes embodied in headline inflation have had no effect on either wage measure over the past two decades.

V. Conclusions

We document significant changes in the dynamics of U.S. inflation as described by an array of various Phillips curve specifications. These specifications cover the conventional backward-looking Phillips curve and a hybrid “New-Keynesian” Phillips curve, this last embedded into a simple general equilibrium representation of the macroeconomy. The different specifications, overall, yield similar results concerning the nature of the changes in inflation dynamics. They point to the impact of energy and food prices on core inflation measures having changed over time, with a much diminished pass-through of food and energy prices into core inflation. They also produce evidence of changes in the effect of past inflation on current inflation, and of potential changes in the sacrifice ratio.

²⁷ The cointegrating discrepancy from the long-run relationship linking wages, prices, and productivity, which should enter such a regression on theoretical grounds, entered insignificantly in almost every case. In fact, the presence of a cointegrating relationship among wages, prices, and productivity is weak, except in the case in which wages, prices, and productivity are all taken from the non-farm accounts. We do not study the dynamics of the implicit deflator for non-farm output in this paper.

²⁸ In particular, neither inflation measure enters significantly in these regressions, whether the sum of the inflation terms is constrained to unity or not. In general, these so-called wage-price Phillips curves meet with rather meager empirical success.

Taking into account changing inflation dynamics produces more accurate out-of-sample forecasts. Indeed, we show that the Atkeson-Ohanian results concerning the inability of Phillips curves to produce better forecasts in the post-1984 sample than a naïve random walk-type of forecast are overturned once time-varying coefficients are allowed into an otherwise standard backward-looking Phillips curve. Provided changing inflation dynamics are taken into account, the Phillips curve framework continues to have, with a few notable exceptions, relevant economic content for explaining the dynamics of inflation and for inflation forecasting.

The estimated extent of changing inflation dynamics can differ according to the Phillips curve specification used—an issue that is even more pressing when the Phillips curve is estimated in the context of a general equilibrium framework. The uncertainty raised by this variation complicates the forecasting exercise noticeably in real time. Here, the Federal Reserve Board’s Greenbook inflation forecasts are instructive. Recent Greenbook inflation forecasts are predicated on a very high sacrifice ratio. Some Phillips curve specifications do indeed point to an increase of the sacrifice ratio in recent years, but we show that the Greenbook inflation forecasts seem to have relied on a Phillips curve slope that is potentially too flat.

Overall, while the evidence of changes in inflation dynamics is overwhelming, how to best model such changes is still an open issue. Traditional backward-looking Phillips curves provide flexibility in this regard, but their quasi-reduced form makes it difficult to attribute changes in inflation dynamics to particular structural features of the economy. More-structural versions of the Phillips curve, especially when embedded into a general equilibrium framework, can provide more economic content to the changes. In this case, however, the risk is that the instability could be the figment of a mis-specified model.

The task of how to model changing inflation dynamics is somewhat simplified by the fact that the most notable changes seem to have occurred in the early 1980s. But stability of the Phillips curve is still relative, with potentially important changes concerning inflation expectations formation and the sacrifice ratio having occurred in

more recent years. At this point, we still have too few observations to be confident of these changes. Needless to say, the current recession should provide fertile ground for continued assessment of the economic content and forecasting power of the Phillips curve, and for better ascertainment of more recent changes in inflation dynamics.

References

- Atkeson, Andrew and Lee Ohanian (2001). "Are Phillips curves useful for forecasting inflation?" *Federal Reserve Bank of Minneapolis Quarterly Review*, vol. 25, no. 1, Winter, pp. 2–11.
- Bai, Jushan (1999). "Likelihood ratio tests for multiple structural changes." *Journal of Econometrics*, vol. 91, pp. 299–323.
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (2005). "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, vol. 133, no. 1, pp. 1–45.
- Clark, Todd and Michael McCracken (2006). "The predictive content of the output gap for inflation: Resolving in-sample and out-of-sample evidence." *Journal of Money, Credit, and Banking*, vol. 38, no. 5, pp. 1127–1148.
- Galí, Jordi and Mark Gertler (1999). "Inflation dynamics: A structural econometric analysis." *Journal of Monetary Economics*, vol. 44, pp. 195–222.
- Gordon, Robert J. (1985). "Understanding inflation in the 1980s." *Brookings Papers on Economic Activity*, vol. 1985, no. 1, pp. 263–302.
- Inoue, Atsushi and Lutz Kilian (2004). "In-sample or out-of-sample tests of predictability: Which one should we use?" *Econometric Reviews*, vol. 23, pp. 371–402.
- O'Reilly, Gerard and Karl Whelan (2005). "Testing parameter stability: A wild bootstrap approach." Research Technical Paper 8/RT/05, Central Bank Financial Services Authority of Ireland.
- Rich, Robert and Charles Steindel (2007). "A comparison of measures of core inflation." Federal Reserve Bank of New York *Economic Policy Review*, December, pp. 19–38.
- Stock, James and Mark Watson (2007). "Why has U.S. inflation become harder to forecast?" *Journal of Money, Credit and Banking*, vol. 39, Issue s1, pp. 3–33.
- Stock, James and Mark Watson (2008). "Phillips curve inflation forecasts." Presented at the 52nd annual economic conference of the Federal Reserve Bank of Boston, June.
- Taylor, John B. (1993). "Discretion versus policy rules in practice." *Carnegie-Rochester Conference Series on Public Policy*, vol. 39, pp. 195–214.
- Tetlow, Robert and Brian Ironside (2007). "Real-time model uncertainty in the United States: The Fed, 1996–2003." *Journal of Money, Credit and Banking*, vol. 39, no. 7, pp. 1533–1561.

Data definitions for Sections III–IV

Inflation: One-quarter percentage change in the quarterly averages of the total and core measures of the consumer price index (CPI) and the personal consumption expenditures chain-type price index (PCE).

Unemployment: Quarterly average of the civilian unemployment rate.

Relative price of oil: One hundred times the log difference between the West Texas intermediate price of oil and the consumer price index.

Relative price of non-oil imports: One hundred times the log difference between the NIPA non-petroleum import chain-type price index and the consumer price index.

Output gap: One hundred times the log difference between real GDP and HP-filtered real GDP, smoothing parameter 1600.

Policy rate: Quarterly average of the effective federal funds rate.

Table I.1

Root Mean Squared Errors for Phillips Curve Specifications Relative to Atkeson-Ohanian Bench

	<u>Core CPI Inflation</u>	<u>Core PCE Inflation</u>
<i>RMSE of AO Forecast</i>	0.53	0.44
Relative RMSE of Phillips Curve Specifications		
<i>In-Sample Forecasts :</i>		
Time-Varying NAIRU	0.683	0.816
Constant NAIRU	0.707	0.827
Time-Varying NAIRU and No Supply Controls	0.699	0.881
Constant NAIRU and No Supply Controls	0.718	0.869
<i>Pseudo Out-of-Sample Forecasts :</i>		
Time-Varying NAIRU	1.055	0.848
Constant NAIRU	1.230	0.850
Time-Varying NAIRU and No Supply Controls	1.029	0.936
Constant NAIRU and No Supply Controls	1.147	0.900

Table I.2

Root Mean Squared Errors for Out-of-Sample Phillips Curve Specifications with Time-Varying Coefficients Relative to Atkeson-Ohanian Benchmark

	<u>Core CPI Inflation</u>	<u>Core PCE Inflation</u>
<i>RMSE of AO Forecast</i>	0.53	0.44
Relative RMSE of Phillips Curve Specifications	0.385	0.425

Table III.1a				
Coefficient stability in conventional Phillips curve				
<i>p</i> -value for F-test				
Null hypothesis: Stable coefficients				
Core PCE				
	"Known" breakpoint			
	Great Moderation	Greenspan era	post-1994	
All coefficients	0.97	0.96	0.58	
Lags	0.40	0.67	0.55	
Unemployment (U)	0.99	0.76	0.52	
RP of oil	0.00	0.056	0.21	
RP of non-oil imports	0.43	0.40	0.92	
Unknown breakpoint detected at: 1982:Q2				
<i>p</i> -value for F-test for coefficient shifts at breakpoints				
	Lags	Unemp.	Oil	Non-oil
1982:Q2	0.048	0.98	0.00	0.13

Table III.1b				
Coefficient stability in conventional Phillips curve				
<i>p</i> -value for F-test				
Null hypothesis: Stable coefficients				
Core CPI				
	“Known” breakpoint			
	Great Moderation	Greenspan era	post-1994	
All coefficients	0.99	0.45	0.97	
Lags	0.41	0.80	0.82	
Unemployment (U)	0.53	0.61	0.19	
RP of oil	0.00	0.036	0.045	
RP of non-oil imports	0.47	0.38	0.98	
Unknown breakpoint: 1980:Q2				
<i>p</i> -value for F-test for coefficient shifts at breakpoints				
	Lags	Unemp.	Oil	Non-oil
1980:Q2	0.01	0.89	0.00	0.87

Table III.2		
Coefficient stability in structural model		
Baseline Hybrid model		
<i>p</i> -value for likelihood ratio test		
Null hypothesis: Stable coefficients, 1966–2007:Q4		
Core PCE		
	“Known” breakpoint	
	Great Moderation	Greenspan era
All coefficients	0.000	0.000
Phillips coefficients:	0.20	0.29
Taylor rule coefficients:	0.27	0.077
I-S curve coefficients	0.000	0.0049
Unknown breakpoint: <i>sup</i> of likelihood ratio at 1981:Q1		
Core CPI		
	“Known” breakpoint	
	Great Moderation	Greenspan era
All coefficients	0.000	0.000
Phillips coefficients:	0.29	0.45
Taylor rule coefficients:	0.71	0.33
I-S curve coefficients	0.000	0.0052
Unknown breakpoint: <i>sup</i> of likelihood ratio at 1980:Q3		

Table III.3		
Coefficient stability in structural model		
Augmented hybrid model		
<i>p</i> -value for likelihood ratio test		
Null hypothesis: Stable coefficients, 1966–2007:Q4		
Core PCE		
	“Known” breakpoint	
	Great Moderation	Greenspan era
All coefficients	0.000	0.000
Phillips coefficients:	0.015	0.037
Taylor rule coefficients:	0.27	0.085
I-S curve coefficients	0.000	0.0031
Unknown breakpoint: <i>sup</i> of likelihood ratio at 1981:Q1		
Core CPI		
	“Known” breakpoint	
	Great Moderation	Greenspan era
All coefficients	0.000	0.0023
Phillips coefficients:	0.0024	0.0040
Taylor rule coefficients:	0.88	0.72
I-S curve coefficients	0.000	0.019
Unknown breakpoint: <i>sup</i> of likelihood ratio at 1983:Q1		

Table IV.1
Pass-through regressions

$$\pi_t^J = \sum_{i=1}^n a_i \pi_{t-i}^H + \sum_{i=1}^n b_i \pi_{t-i}^C + \sum_{i=1}^2 c_i U_{t-i} + \varepsilon_t$$

$$J = [H, C]$$

4 lags

$$\sum a_i + \sum b_i = 1$$

No explicit relative price effects ($d_i=0$)

Dependent variable	Weight on headline	Signif.?^a	Weight on core	Signif.?^a
1970:Q1–1985:Q4				
Headline CPI	1.34	Yes	−0.34	No
Headline PCE	1.28	Yes	−0.28	No
Core CPI	0.90	Yes	0.10	No
Core PCE	0.68	Yes	0.32	No
1985:Q1–2008:Q1				
Headline CPI	0.32	Yes	0.68	Yes
Headline PCE	0.20	No	0.80	Yes
Core CPI	0.11	No	0.89	Yes
Core PCE	0.14	No	0.86	Yes
1995:Q1–2008:Q1				
Headline CPI	0.54	No	0.46	No
Headline PCE	0.30	No	0.70	No
Core CPI	0.19	No	0.81	Yes
Core PCE	0.18	No	0.82	Yes
2000:Q1–2008:Q1				
Headline CPI	0.10	No	0.90	No
Headline PCE	−0.28	No	1.28	No
Core CPI	0.20	No	0.80	Yes
Core PCE	0.15	No	0.85	Yes

^a Throughout, “Yes” denotes significance at the 5% level or better

Table IV.2				
Pass-through regressions including key relative prices as regressors				
Dependent variable	Weight on headline	Signif.?	Weight on core	Signif.?
1970:Q1–1985:Q4				
Headline CPI	1.22	Yes	–0.22	No
Headline PCE	1.05	Yes	–0.05	No
Core CPI	0.57	No	0.43	No
Core PCE	0.38	No	0.62	No
1990:Q1–2008:Q1				
Headline CPI	0.12	No	0.88	Yes
Headline PCE	–0.01	No	1.01	Yes
Core CPI	0.19	No	0.81	Yes
Core PCE	0.22	No	0.78	Yes

Table IV.3 Pass-through of price changes into wages 4 lags $\sum a_i + \sum b_i = 1$ No explicit relative price effects ($e_i=0$)				
	ECI		NFC	
Sample	Headline	Core	Headline	Core
	$\sum a_i$	$\sum b_i$	$\sum a_i$	$\sum b_i$
1970:Q1–2008:Q1			0.54	0.46
Significance?			Yes	No
1980:Q2–2008:Q1	0.43	0.57	0.48	0.52
Significance?	Yes	Yes	No	Yes
1990:Q1–2008:Q1	0.5	0.5	0.48	0.52
Significance?	No	Yes	No	No
1995:Q1–2008:Q1	0.51	0.49	0.69	0.31
Significance?	No	No	No	No
2000:Q1–2008:Q1	0.52	0.48	1.3	–0.33
Significance?	No	No	No	No
1970:Q1–1989:Q4			0.53	0.47
Significance?			No	No
Include explicit relative price effects				
1970:Q1–2008:Q1			0.55	0.45
Significance?			No	No
1980:Q2–2008:Q1	0.39	0.61	0.42	0.58
Significance?	No	Yes	No	Yes
1990:Q1–2008:Q1	0.4	0.6	0.38	0.62
Significance?	No	Yes	No	No
1995:Q1–2008:Q1	0.36	0.64	0.91	0.093
Significance?	No	No	No	No
2000:Q1–2008:Q1	0.44	0.56	1.8	–0.78
Significance?	No	No	No	No
1970:Q1–1989:Q4			0.54	0.46
Significance?			No	No
“Yes” denotes significant at the 5% level or better				

Figure I.1a Core CPI Inflation - RMSE

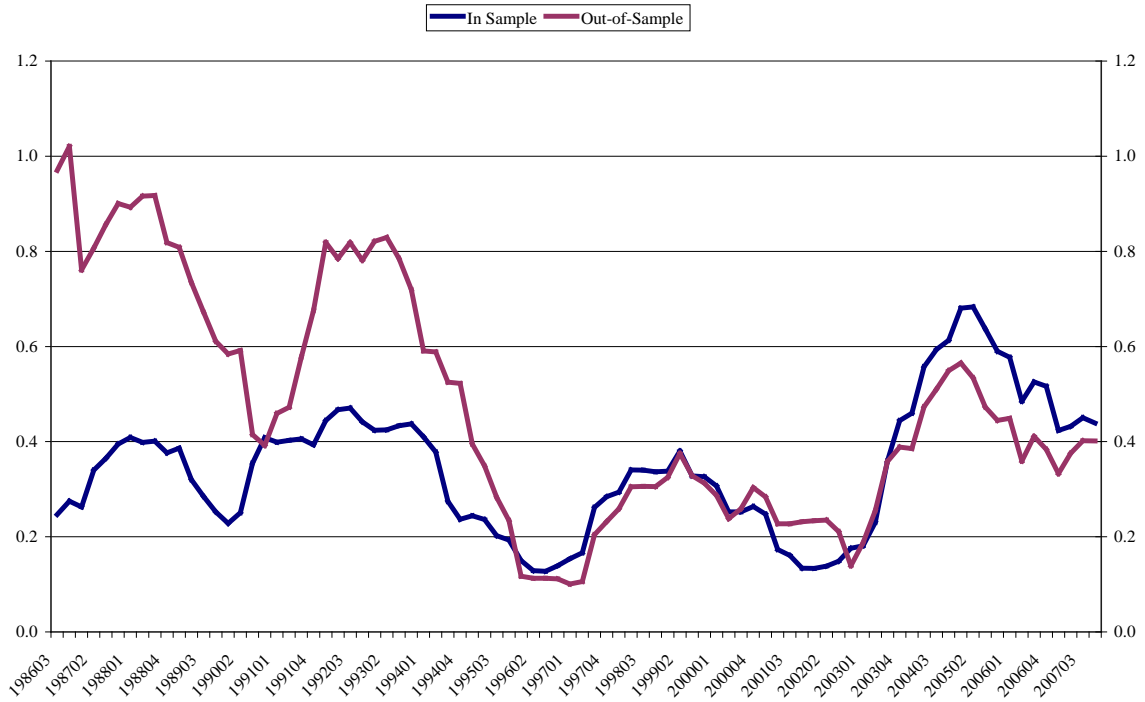
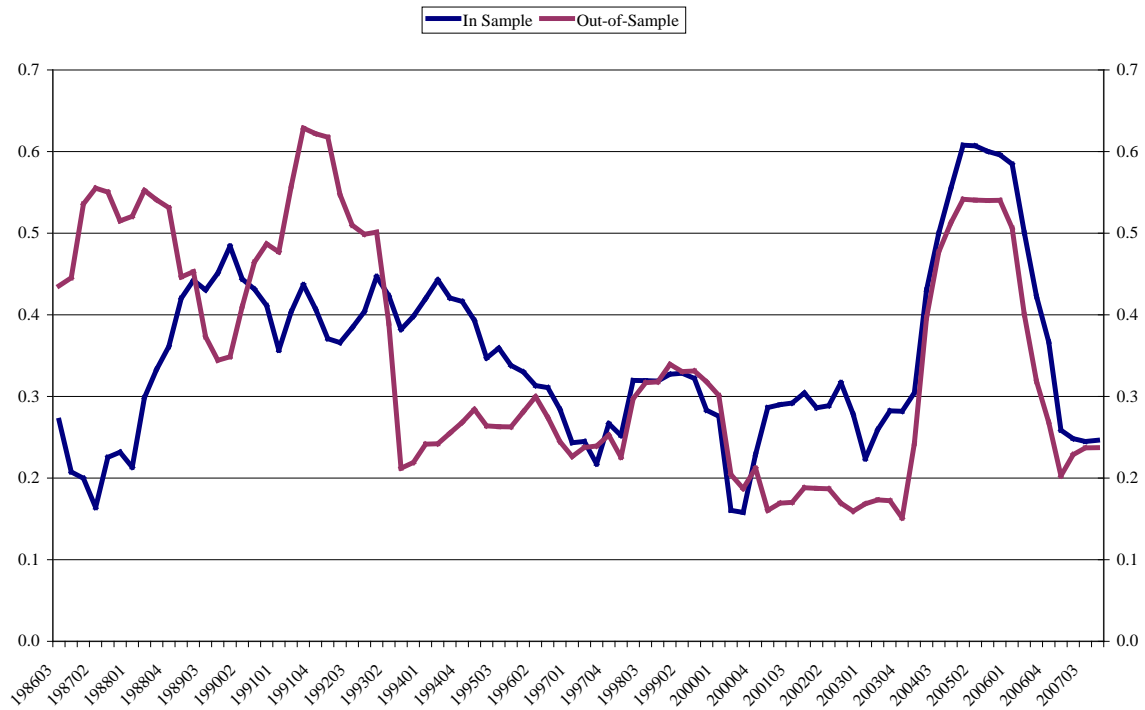


Figure I.1b Core PCE Inflation - RMSE



Note: The figures compare the root mean squared errors (RMSE) for inflation forecasts over a rolling window of eight quarters.

Figure I.2a Core CPI Inflation - RMSE

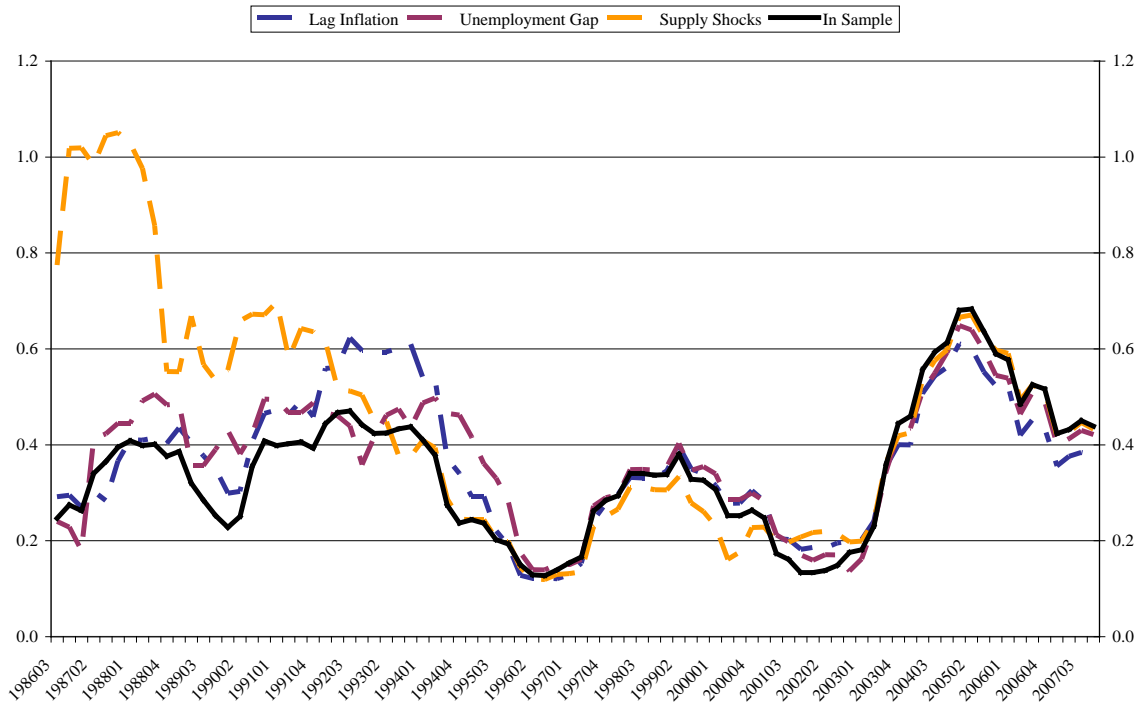
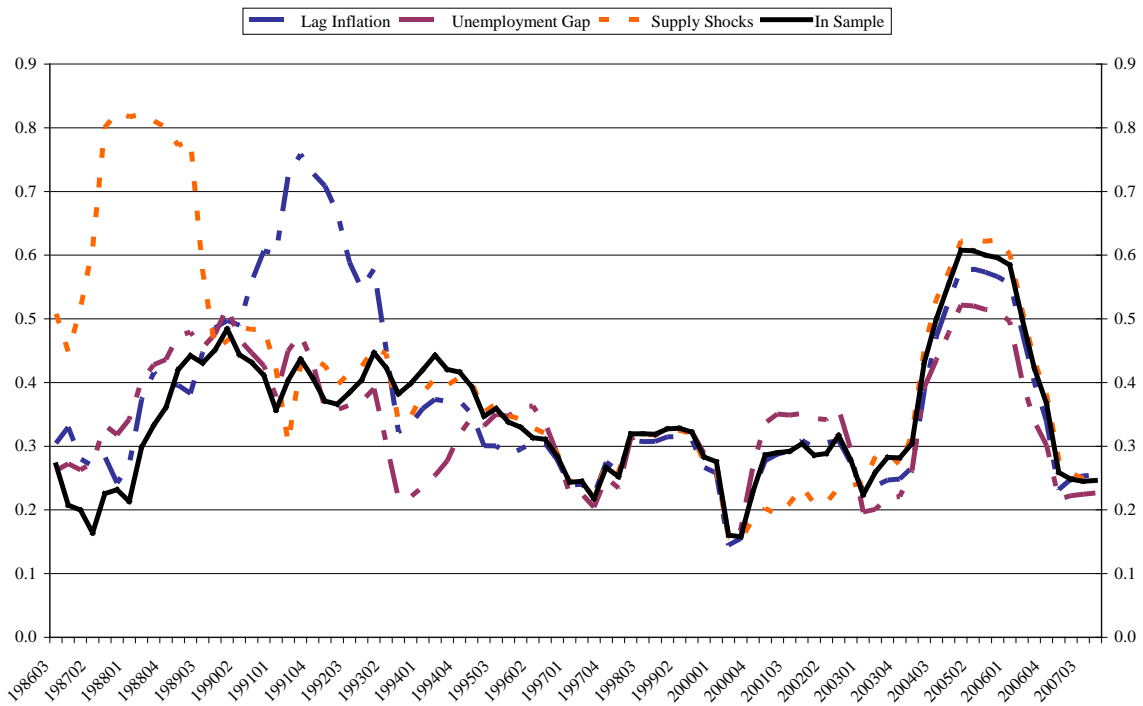


Figure I.2b Core PCE Inflation - RMSE



Note: The figures compare the root mean squared errors (RMSE) for inflation forecasts over a rolling window of eight quarters.

Figure I.3a Core CPI Inflation - RMSE of Philips Curve vs. AO Benchmark

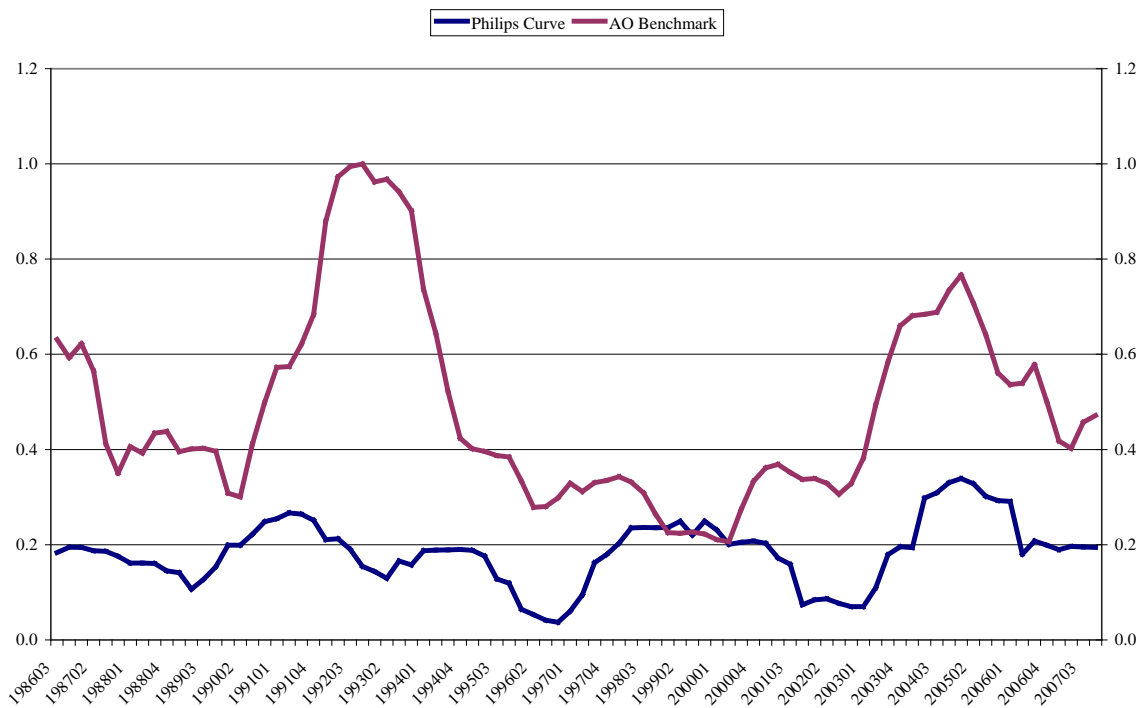
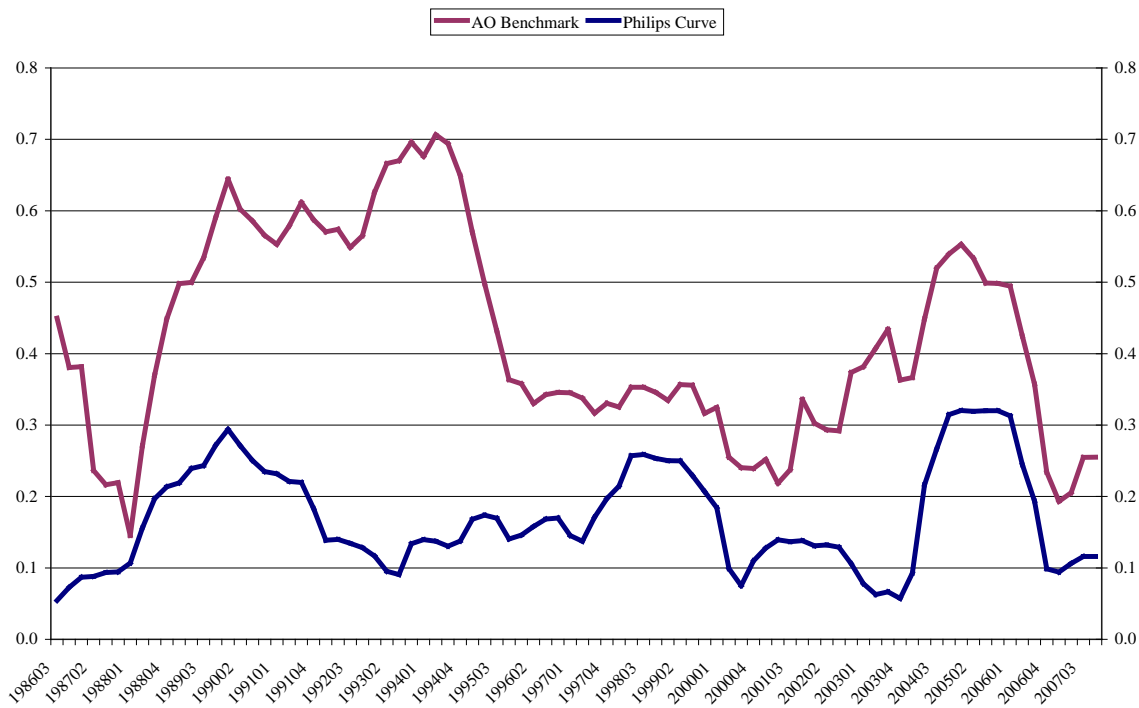


Figure I.3b Core PCE Inflation - RMSE of Philips Curve vs. AO Benchmark



Note: The figures compare the root mean squared errors (RMSE) for inflation forecasts over a rolling window of eight quarters.

Figure I.4a Core CPI Inflation - RMSE Univariate AR Process

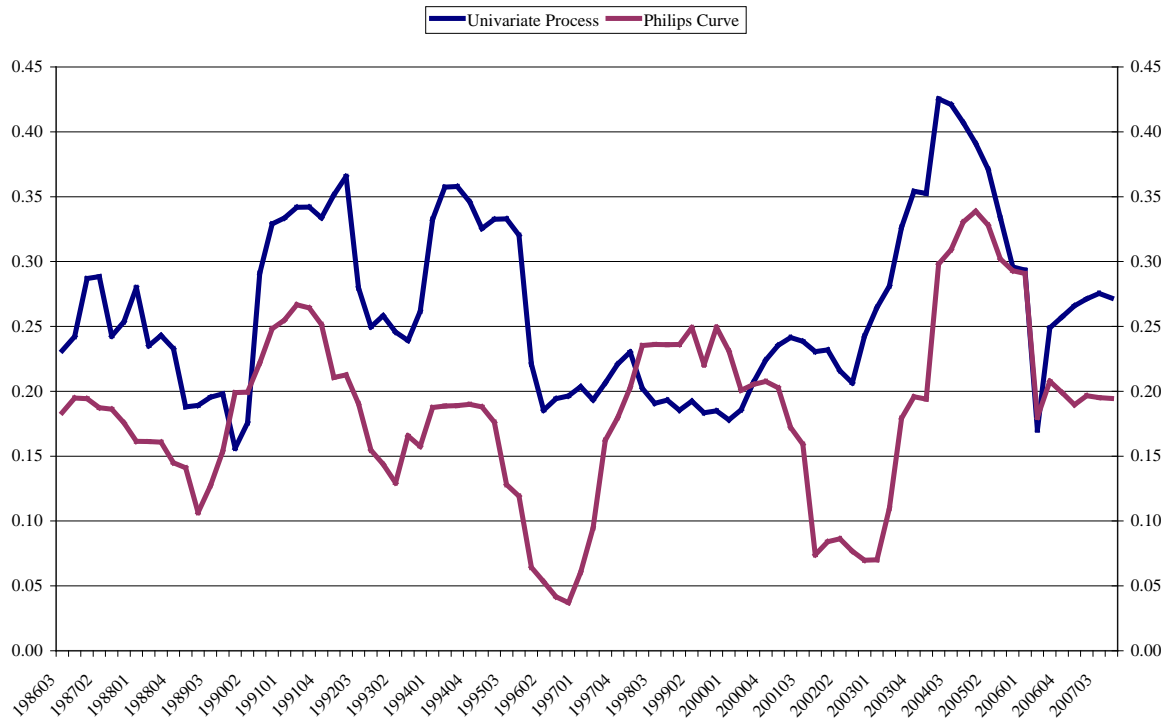
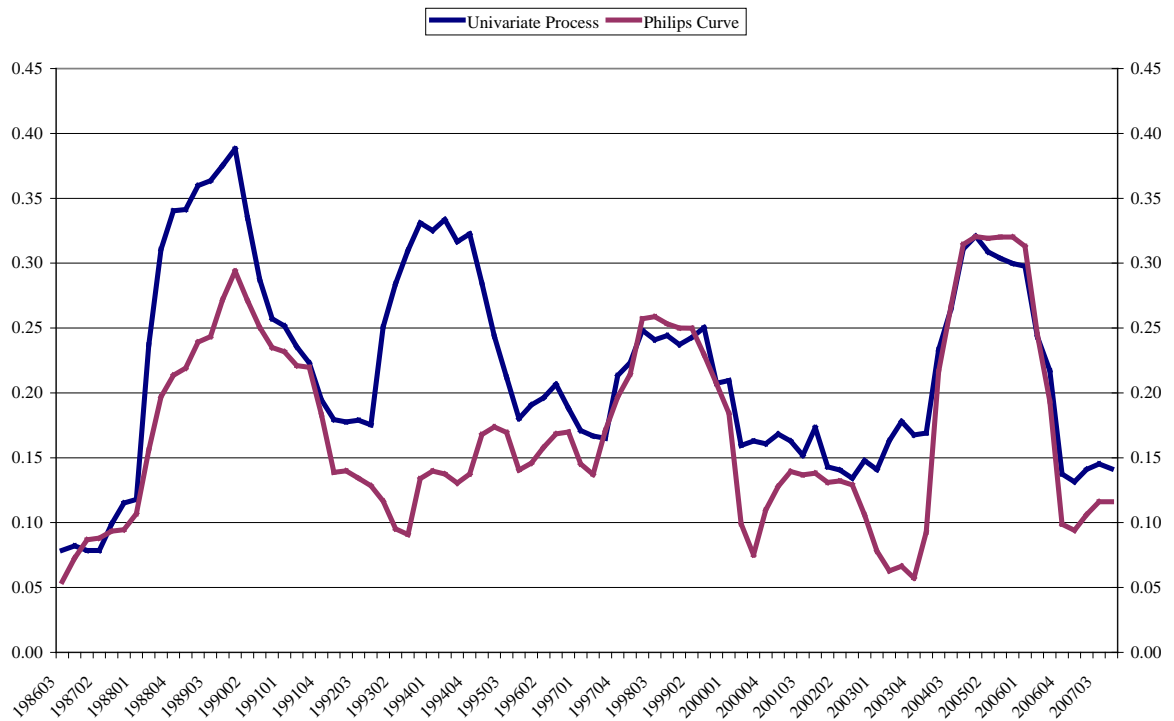
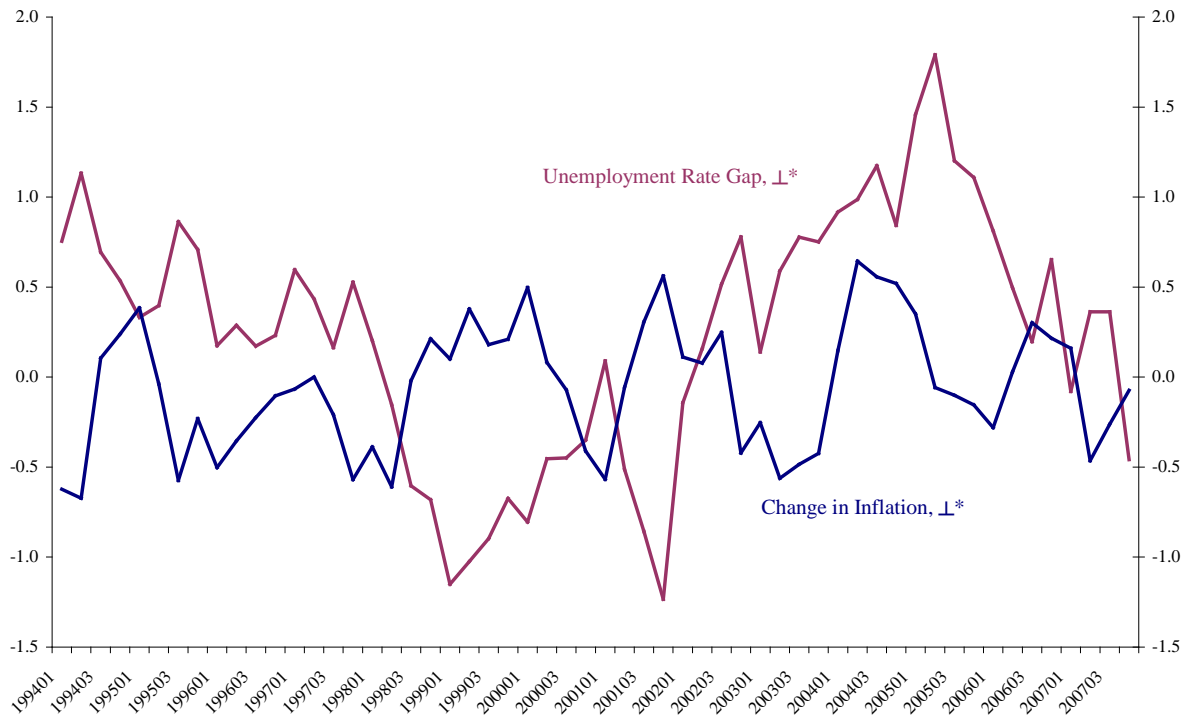


Figure I.4b Core PCE Inflation - RMSE Univariate AR Process



Note: The figures compare the root mean squared errors (RMSE) for inflation forecasts over a rolling window of eight quarters.

Figure II.1 The Inflation-Unemployment Tradeoff, 1994 to 2007



- * Portion of change in inflation not explained by lagged inflation and supply shocks.
- * Portion of unemployment rate gap not explained by lagged inflation and supply shocks.

Figure II.2a Core CPI Inflation - Estimated Slope of Philips Curve

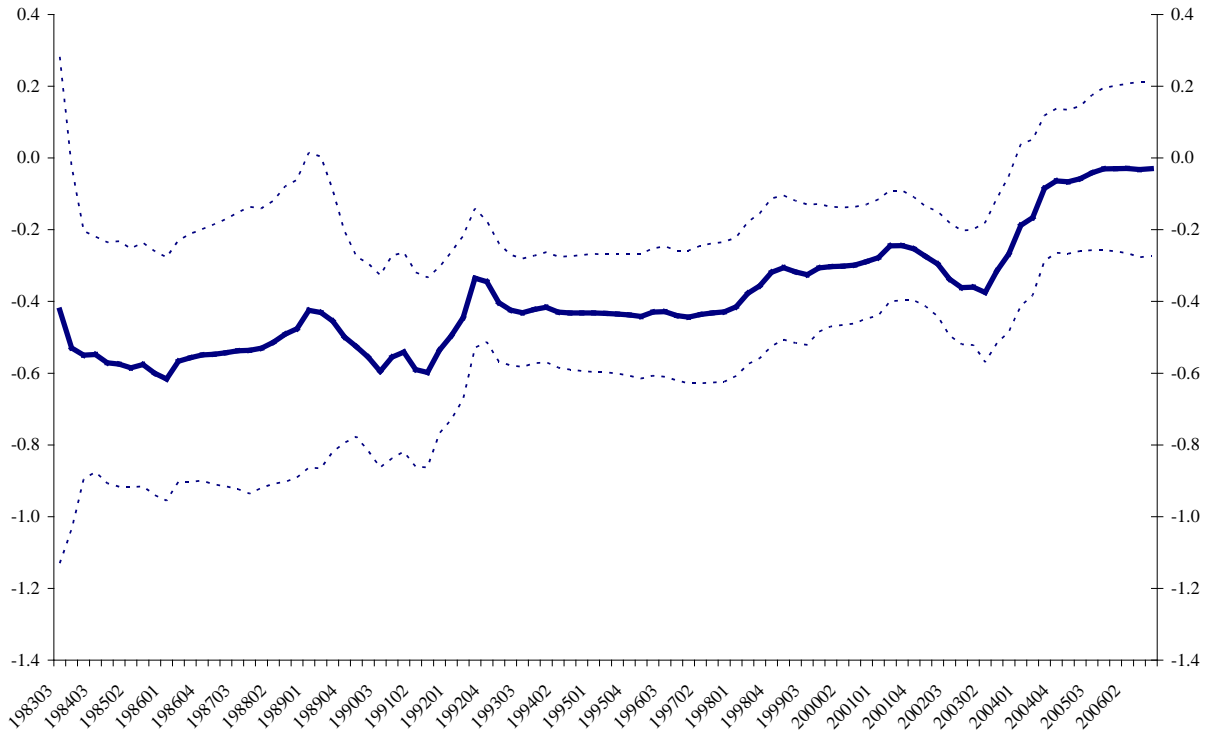


Figure II.2b Core PCE Inflation - Estimated Slope of Philips Curve

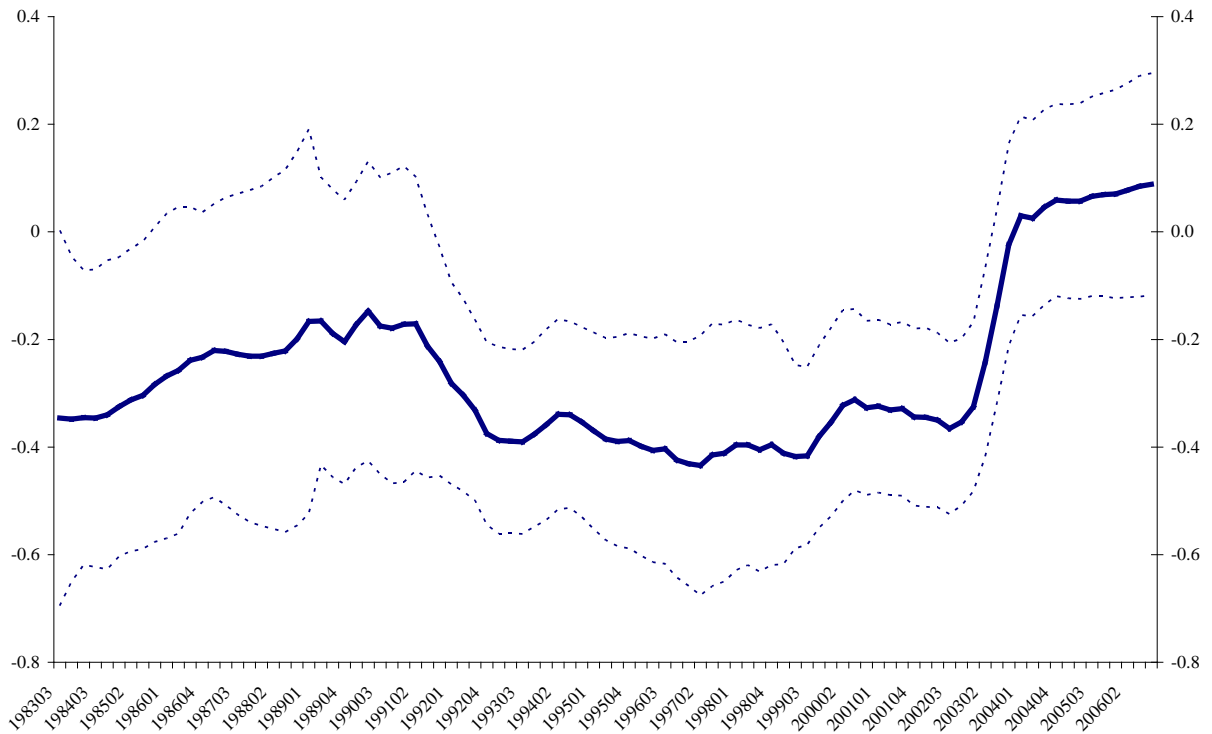


Figure II.3a Estimated Slope of the Phillips Curve Derived from Greenbook Forecasts

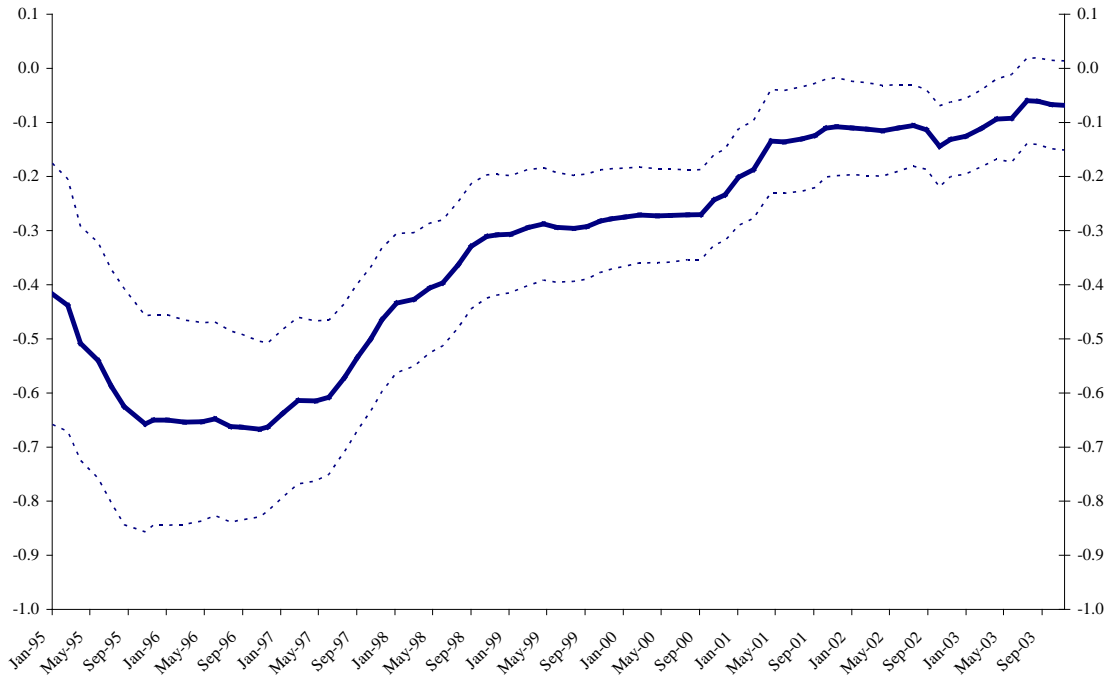


Figure II.3b Estimated NAIRU Derived from Greenbook Forecasts

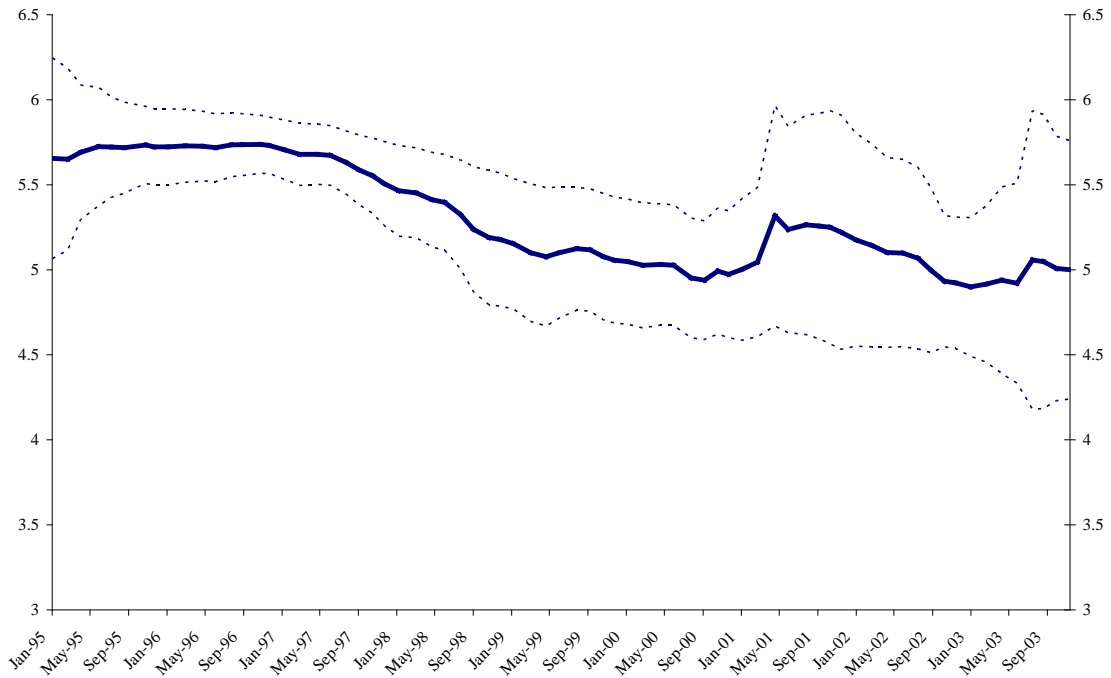


Figure II.4 Actual and Projected Core CPI Inflation

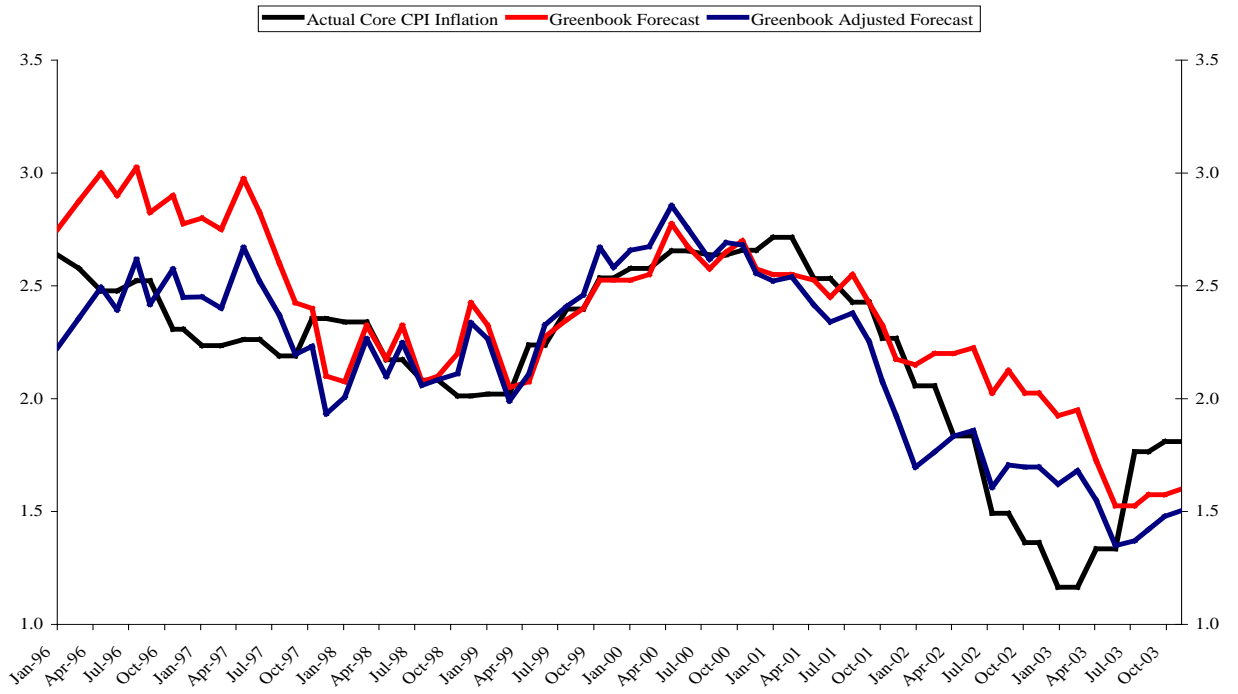


Figure III.1
Estimated model parameters at
various break dates
Augmented hybrid model, PCE

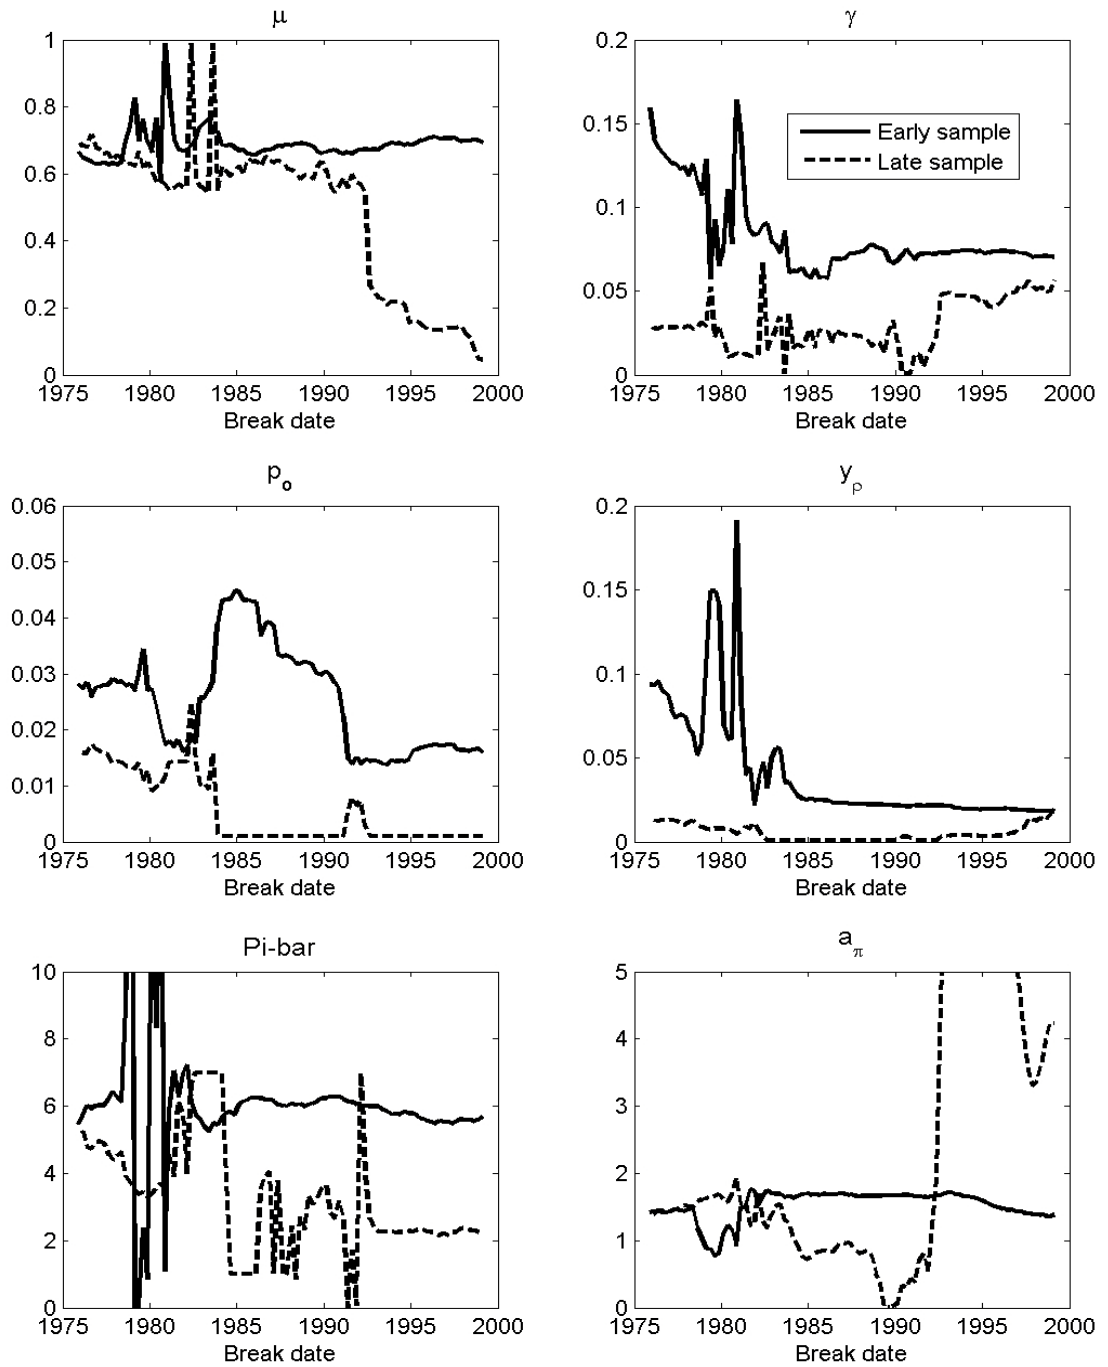


Figure IV.1

CPI-U: All Items Less Food and Energy
% Change - Year to Year SA, 1982-84=100

CPI-U: All Items
% Change - Year to Year SA, 1982-84=100

