Inflation Dynamics when Inflation is Near Zero

Jeff Fuhrer, Giovanni Olivei and Geoffrey Tootell

October 14, 2010

1. Introduction

In the wake of the longest post-war recession in US history, US inflation has declined noticeably from its 2008 peak (see the table below for a summary of a variety of inflation measures). At about one percent at the time of writing, inflation is as low as it has been since the early 1960s. With most economists estimating that significant economic slack remains, it is reasonable to wonder about the future trajectory of inflation. Will well-anchored expectations pull inflation up as the recovery proceeds? Will slack resources pull inflation down toward or below zero? Is the US headed for a Japanese-style period of protracted, albeit modest deflation?

Table 1

Recent declines in inflation

Inflation measure	Change since 2008 peak (percentage points)	Current Inflation rate (12-mo. or 4-qtr. chg.)
Core CPI	-1.6	0.95
Headline CPI	-6.1	1.20
Core PCE	-1.6	1.38
Headline PCE	-4.8	1.54
GDP deflator	-2.4	0.84
Cleveland Fed trimmed mean	-3.2	0.50
Cleveland Fed weighted median	-5.9	0.90
ECI private compensation	-1.1	1.82

A key difficulty in answering these questions is that the US has very little recent experience with very low inflation rates. A number of researchers have examined the behavior of inflation over the past twenty years, as it fell from above two percent to an average of about two percent from the late 1990s through the mid-2000s. But the concern of this paper centers on the behavior of inflation as it drops *below* two percent, and the experience in this range is quite limited. One might wish to look for a breakpoint in the timeseries properties of inflation data in recent years, but it is nearly impossible to detect a breakpoint in the last few years of a series.[1]

While many approaches to exploring these questions are possible, we consider a few tacks: First, what do standard models of inflation imply about inflation dynamics? In this regard, we take on board the recent advances, both theoretical and empirical, in modeling inflation, including incorporating the effect of "trend inflation" as in Cogley and Sbordone (2009) and the related notion of the "anchoring" of inflation provided by stable long-run inflation expectations. We also consider the possibility of a declining slope of the Phillips curve and reduced pass-through of energy prices, as documented in Fuhrer and Olivei (2010).

[1] While not reported here, we conduct conventional unknown breakpoint tests for changes in the autoregressive properties of a variety of inflation measures for the US. Simple autoregressive and reduced-form models of the core and headline inflation measures examined in table 1 above typically find either one or no breaks in the series from 1970 to the present, with the breakpoints falling in the period between 1975 and 1980.

Second, we examine the parallels between the long experience with low inflation in Japan and the relatively recent experience in the US. Finally, we consider the buoying effect that downward nominal wage rigidity may have on US price inflation. Here we examine disaggregated data on wages to determine the extent to which firms are able to adjust their wage bill downward; if there is downward rigidity in the firm's wage bill, then this may act to limit the fall in inflation that might otherwise occur.

Section two considers the inflation dynamics implied by a variety of standard models of inflation. The section also examines a potential nonlinearity in the correlation between unemployment and inflation, such that very high rates of unemployment have disproportionately smaller influence on inflation than modest rates. The third section discusses the Japanese experience, with an emphasis on the influence of short- and long-run inflation expectations. The section also examines some parallels between Japanese and US inflation dynamics over the past two decades. Finally, section four examines micro wage data for evidence of two kinds of wage rigidity. The first, which is well-documented, is the resistance of individual wages to nominal wage cuts. The second, potentially of more relevance to price-setting behavior and inflation, examines preliminary evidence on the downward rigidity (or lack thereof) of the firm's wage bill. The final section concludes.

2. Inflation dynamics implied by standard models

While there remains considerable debate surrounding the appropriate way to model inflation, it may still be instructive to consider the implications that standard models of inflation bear for the future trajectory of US inflation. This paper is not the place to summarize the exhaustive (and exhausting) literature on inflation models. Instead, we will use two "workhorse" models of inflation to examine the experience in the US following recessions.

First, consider the old-style backward-looking Phillips curve, which we can represent in simplified form as

$$\pi_t = \pi_{t-1} - a(U_t - \bar{U}) + cX_t$$

where π_t is the annualized percentage change in a price index, U is the civilian unemployment rate, X represents important relative price shifts, and $\bar{U}$ is an estimate of the non-accelerating-inflation rate of unemployment, or NAIRU. Actual empirical implementations of this model will often include additional lags of inflation and unemployment, but the gist of the model is simple: Absent the temporary effects of large relative price changes, the change in the inflation rate depends on the level of unemployment relative to the NAIRU. The implication of this model for post-recession inflation is equally simple: As long as the unemployment rate remains elevated above the NAIRU, inflation will continue to fall.

The first figure displays inflation data over the past sixty years, with recession periods shaded in gray. As is clear from the figure, in every recession since the 1950s, inflation has fallen in the years following recessions. That observation is roughly consistent with the logic of the simple Phillips curve sketched above. Recent studies have suggested that the slope of the Phillips curve, summarized in the parameter *a* above, has likely declined in recent years. This implies that the magnitude of the decline in inflation for a given gap between unemployment and the NAIRU will be smaller, but the directional implications are the same. If this model still retains some validity, the implications for inflation today are quite clear: With unemployment quite elevated above the NAIRU by virtually all accounts, inflation will continue to fall.

As is now well known, it may not be appropriate to think of the old-style Phillips curve as a structural economic relationship. A vast literature, summarized in Fuhrer, Kodrzycki, Little and Olivei (2010), has worked to develop models of inflation that reflect the underlying behavior of price-setting firms, and that provide explicit expectations mechanisms, usually assuming rational expectations. The role of expectations in these models can be quite powerful and bears important implications in current circumstances.

Consider the now-canonical model of inflation that combines the Calvo (1983) rational expectations price-setting model with a notion of trend inflation introduced by Cogley and Sbordone (2009). As Cogley and Sbordone (CS) show, the Calvo model that takes into account the possible presence of a time-varying trend inflation rate takes the form

$$\pi_t - \bar{\pi}_t = \tilde{\rho}(\pi_{t-1} - \bar{\pi}_{t-1} - g_t^{\bar{\pi}}) + b_1 E_t(\pi_{t+1} - \bar{\pi}_{t+1}) + \gamma s_t$$

Where much of the notation is as above, and the $\bar{\pi}_t$ terms reflect an estimate of "trend inflation," *s* is an estimate of real marginal cost, and g_t^{π} is the growth rate of the trend inflation term. Because the time-varying inflation trend is important in this model, we include a monetary policy rule in which the time-varying inflation trend is the (potentially) time-varying inflation target. That target in turn is assumed to follow a random walk:

$$r_t = \max\{0, \omega r_{t-1} + (1-\omega)[a_\pi(\pi_t - \bar{\pi}_t) + a_y \tilde{y}_t + \bar{r}_t]\}$$

$$\bar{r}_t = \bar{\pi}_t + \bar{\rho}$$

$$\bar{\pi}_t = \bar{\pi}_{t-1} + \varepsilon_t$$

where $\bar{\rho}$ is the equilibrium real rate of interest. Output is determined by a hybrid optimizing I-S curve, with weights of 0.6 and 0.4 on the backward- and forward-looking terms. The model is closed with a reduced-form equation that links marginal cost and the output gap.

As demonstrated elsewhere, the behavior of this model varies considerably depending on the weights on the backward-looking and forward-looking terms. A model that is primarily backward-looking will display the properties similar to those described above for the old-style Phillips curve: As long as the unemployment gap remains positive, inflation will continue to decline. However, a model with significant forward-looking weight will display a different dynamic: At the beginning of a recovery, because output and marginal cost are expected to improve, inflation will begin to rise. Even though a substantial output gap may linger for several years, this model implies that inflation will rise because the gap is expected to shrink going forward. In addition, the level of inflation relative to the long-run target (which plays the role of trend inflation in this model) will influence the direction in which inflation moves in the early years of a recovery. If inflation is well below the target, then it will tend to revert towards the target, where the speed of reversion will depend importantly on the size of the coefficient b_1.

In this model, expectations are perfectly anchored—the public understands the Fed's inflation goal and incorporates it into its price-setting. Figure 2 displays the result of simulating this model from realistic initial conditions that are roughly consistent with those observed today: marginal cost is more than five percent below its pre-recession level, inflation is low, and the funds rate can be constrained by the zero lower bound. As the figure suggests, when the weight on the backward-looking model is significant, inflation continues to fall significantly below zero in the presence of very depressed marginal cost, and the funds rate is pinned at zero for quite a while.

If one believes alternatively that inflation is well-characterized by a purely forward-looking version of the Cogley-Sbordone model, then the intuition developed above holds, as can be seen in Figure 3. Now the expected improvement in the trajectory of marginal cost gradually lifts the inflation rate from the beginning of the simulation. Inflation never drops below zero, the funds rate need not drop to zero, and the economy and inflation move towards equilibrium much more quickly.

The conclusions from this exercise are that (a) in conventional models, well-anchored (indeed, perfectly-anchored) inflation expectations need not constrain inflation from dropping significantly in the presence of depressed marginal cost and/or a significant output gap; (b) the extent to which the models imply a large drop in inflation depends critically on the amount of forward-looking behavior in the specification.

Which version of this conventional model—the purely forward-looking or the hybrid—has provided a better description of the historical experience at the beginning of economic recoveries? To answer this question, we slightly modify the model so that it attains better empirical fit over the past several recessions. Instead of marginal cost, we allow an estimate of the unemployment gap to enter the inflation equation, so that the model becomes

$$\pi_t - \bar{\pi}_t = \tilde{\rho}(\pi_{t-1} - \bar{\pi}_{t-1} - g_t^{\bar{\pi}}) + b_1 E_t(\pi_{t+1} - \bar{\pi}_{t+1}) + \gamma U_t$$

(see, for example, Blanchard and Gali (2010) for a model that motivates a similar role for unemployment in the NKPC). Again, we incorporate a conventional policy rule with a time-varying inflation goal

$$r_t = \rho r_{t-1} + (1-\rho)[\mathrm{a}_\pi(\pi_t - \bar{\pi}_t) + \mathrm{a}_u(\mathrm{U} - \bar{\mathrm{U}}) + \bar{\mathrm{r}}]$$
$$\bar{\pi}_t = \bar{\pi}_{t-1} + \varepsilon_t$$

The model is closed with a reduced form equation for the unemployment gap.[2] The model is estimated on a sample from 1970 to 2009, so this must be considered an in-sample exercise. The estimated values of the parameters imply a modest degree of forward-looking behavior, with $\tilde{\rho}$ =0.55 and b_1=0.44. The trend inflation variable (time-varying inflation goal) is taken from Cogley and Sbordone's code, as replicated and extended in Barnes, Gumbau-Brisa and Olivei (2009). We examine an alternate case in which the backward-looking/expectations parameters are set to 0 and .99, respectively. We simulate the model at a quarterly frequency, starting from the beginning of the NBER recession troughs for the past five recessions. Figure 4 displays the results.

In general, the hybrid model more closely conforms to the historical post-recession experience than the forward-looking model. The exception is the 1975 recession, in which inflation drops more rapidly than the hybrid model would suggest. Of interest is the current recession, in which the purely forward-looking model suggests that well-anchored expectations that weight the central bank's inflation goal, coupled with the expectation of a large but improving unemployment

[2] The reduced-form equation for unemployment is adjusted to be consistent with sensible steady-state values for itself, the rate of inflation, and the funds rate. The latter two are pinned down by the target rate of inflation, and by the sum of the target rate and the equilibrium real rate of interest, respectively. Note that for estimation purposes, the "trend inflation" Phillips curve is augmented to allow for the effects of the relative price of oil, although this term plays relatively little role in the simulations displayed below.

gap, suggest immediately rising inflation. To date, this model has not tracked the generally declining trajectory of inflation.

One should exercise caution in using such stylized models to guide policy decisions. But as far as they go, these exercises suggest that in general and today in particular, a hybrid model that puts some weight on lagged inflation may be a better guide in predicting the evolution of inflation than the purely forward-looking model. The hybrid model suggests continued gradual decline in inflation for the next several years.

U.S. inflation dynamics in the post-WWII period are unlikely to be adequately summarized by a stable reduced form. Sources of instability in the inflation process have been documented extensively in the literature. These instabilities have made inflation difficult to forecast over time. The current very low inflation environment is also relatively new, and it is possible that the dynamics of inflation at 2 percent or lower will turn out to differ fundamentally from the past. Potential changes in inflation dynamics associated with a low inflation environment are difficult to test at this point, as data for this low inflation regime are still scarce. While some studies have tried to address the issue of whether the inflation – unemployment tradeoff changes at different levels of inflation, it has been difficult to usefully exploit non-linearities of this kind for forecasting purposes (see Stock and Watson, 2009).

With these caveats in mind, it is still valuable to revisit the one period in U.S. post-WWII history other than the most recent years when inflation was very low. This period covers the years going from 1954 to 1963.[3] Then, inflation averaged about 1.6 percent and fluctuated within a relatively narrow range. The private consumption expenditures deflator's maximum year-over-year increase over this period was 3¼ percent. The conduct of monetary policy in those years was recently reevaluated in favorable terms by Romer and Romer (2002). After examining the F.O.M.C. records of the time, Romer and Romer claim that the Federal Reserve "… showed the same overarching concern about inflation that is the hallmark of post-Paul Volcker monetary-policy orthodoxy." In particular, they argue that policy tightening in response to increases in expected inflation was more aggressive than in the late 1960s and in the 1970s. Moreover, these years are potentially interesting not just because inflation was low, but also because there were some relatively large fluctuations in the unemployment rate.

It is apparent from Figures 5 and 6 that over the period 1954:Q1 to 1963:Q4 there were some run-ups in inflation during expansions, followed by declines in inflation immediately after recessions. This pattern is consistent with a Phillips curve-type relationship. Figure 7 plots the unemployment rate against the 4-quarter-ahead inflation rate, as measured by the PCE deflator. While not especially tight, the data show a negative tradeoff between inflation and unemployment. The data also suggest the presence of a non-linearity, with inflation increasing at a faster rate for a given decline in the unemployment rate when the unemployment rate is at low levels. This non-linearity mimics the one found by Phillips (1958) in his original study of the relationship between wage inflation and the unemployment rate in the United Kingdom. In the U.S. context, such a feature of the data is the result of a few quarterly observations and, as a result, it should be taken with caution.

---

[3] We exclude the early 1950s from the analysis because of the 1951 and 1952 price control amendments associated with the Korean War.

The observations over the period 1954 to 1963 suggest a modest tradeoff between inflation and unemployment overall, and especially so at high levels of the unemployment rate. This is confirmed by estimating the following relationship between inflation and the unemployment rate:

, (1)

where standard errors for the estimates are in parenthesis.[4] In the equation, is the 4-quarter-ahead inflation rate expressed as a deviation from the long-run trend of inflation as of time *t*.[5] We assume that trend inflation is equal to the average rate of inflation (1.6 percent) prevailing over the period. With inflation equal to trend, the equation estimates the equilibrium unemployment rate at about 5¼ percent. The log-relationship in (1) allows the tradeoff between inflation and unemployment to differ at different levels of the unemployment rate. This specification is very similar to the relationship estimated in Phillips (1958) original article. The estimated nonlinearity in U.S. data over the period 1954 to 1963 is fairly small. A decline in the unemployment rate from 5 to 4 percent raises annual inflation by seven-tenths of one percent. The same one percentage point decline in unemployment, this time from 8 to 7 percent, increases inflation by four-tenths of one percent. We note that the estimated root mean squared error (RMSQE) from (1) is roughly two-thirds of the RMSQE from estimating a completely backward-looking standard linear Phillips curve over the same sample. In other words, over this sample period an "accelerationist" Phillips curve specification does not perform particularly well.

Suppose we apply the estimated relationship in (1) to the current situation. With an unemployment rate averaging 9.6 percent in 2010:Q3, core PCE inflation over the period 2010:Q3 to 2011:Q3 should fall another percentage point, to roughly one half of one percent. Such a projection is conditional on trend inflation not moving over the forecast horizon. Note that the projected magnitude of the decline in inflation is fairly close to the range of estimates (eight-tenths of one percentage point in a comparable exercise) provided by Stock and Watson (2010). Stock and Watson examine the relationship between inflation and unemployment by looking exclusively at recession episodes over the post-1960 sample. In addition to a different sample period, labor market slack is defined differently and enters the specification linearly. Still, there are some more fundamental similarities between the two approaches. Most notably, inflation in both exercises is a function of long-run (trend) inflation and labor market slack. Whether this is an accurate reduced-form representation of the inflation process is an issue that will be considered more closely in the next section.

While it may seem far-fetched to project the near-term evolution of inflation based on a Phillips curve relationship estimated over the period 1954 to 1963, we note that recent movements in inflation fall well within the range experienced during that period. Figure 8 depicts the unemployment rate vis-à-vis the 4-quarter-ahead inflation rate, expressed as a deviation from long-term inflation. The chart superimposes observations from the period 1954:Q1 to 1963:Q4 with observations spanning the period 2003:Q1 to present. As before, trend inflation in the earlier period is given by the sample average inflation rate. For the most recent period, inflation is measured by the core PCE index, and trend inflation is the 10-year inflation expectation from the Survey of

[4] Standard errors are corrected to account for serial correlation in the estimated error term in (1)

[5] In equation (1), 4-quarter-ahead inflation is computed as: $\pi^4_{t+4} = \left(\frac{JC_{t+4}}{JC_t} - 1\right)$, where JC is the PCE deflator.

Professional Forecasters. Merging these two periods yields the following estimate of the nonlinear Phillips curve:

, (1')

which is very similar to the estimate obtained in (1). The slope of the curve is somewhat shallower, but the implications in terms of the near-term evolution of inflation are not materially different. The projected decline in inflation over the period 2010:Q3 to 2011:Q3 is now eight-tenths of one percent. Using the estimates in (1'), the RMSQE of the forecast for inflation over the period 2003:Q1 to present is 22 percent smaller than the RMSQE from a naïve forecast à la Atkeson and Ohanian (2001), where the forecast for 4-quarter-ahead inflation is given by the current value of 4-quarter inflation. It is also 13 percent smaller than the RMSQE from forecasting 4-quarter-inflation by just positing an expected value equal to the current level of long-term inflation expectations. At least when measured against these benchmarks, the estimates in (1') do appear to have some content.

The modest further decline in inflation implied by the relationship in (1) given the current level of the unemployment rate is the result of two main features. First, the process for inflation is anchored by trend inflation (or long-run inflation expectations). In this respect, the specification in (1) generates dynamics that are similar to the purely forward-looking New Keynesian Phillips curve considered in the previous section. As already mentioned there, there are other specifications of the inflation process that do find some support in the data and which entail a much less benign inflation outlook, as the next section will show. Even within the confines of (1), stability of long-term inflation is a crucial factor in limiting the decline in inflation, as no feedback is assumed from the current low readings of inflation to long-run expected inflation. If long-run expectations were to decline, then the implied decline in inflation would be larger. In this regard, it is important to note that during the period 1954 to 1963, the recessionary episodes were fairly short-lived. For example, in the recession of 1957-58, the unemployment rate was back to 5 percent in the second quarter of 1959. With relatively short downturns, it may not be unreasonable to assume little change in long-run expected inflation, other things equal. But in a situation in which the unemployment rate may stay persistently above the equilibrium rate for a long time, such an assumption may prove inaccurate.

The other feature at play in generating a small projected additional decline in inflation is the non-linear tradeoff between inflation and unemployment. In his original article, Phillips (1958) argued that when the unemployment rate is low and labor demand is high, firms will "bid wage rates up quite rapidly." Conversely, Phillips conjectured that when the unemployment rate is high wages will fall "only very slowly" because workers are "reluctant to offer their services at less than the prevailing rates." Wage developments over the period 1954 to 1963 in the U.S. do suggest some non-linearity. Figure 9 depicts wage inflation as measured by compensation per hour in the nonfarm business sector against the unemployment rate. There is a tendency for the wage rate to increase when the unemployment rate is low, but relatively little downward pressure when the unemployment rate is high. Again, one should not infer too much from this chart because of the very limited number of observations. In addition, workers' reluctance to be paid less than the ongoing wage may be inversely related to the persistence of unemployment.

The more recent data on compensation are somewhat more difficult to decipher. Figure 10 shows the relationship between compensation per hour in the nonfarm business sector (expressed as a deviation from long-run expected inflation) and the unemployment rate over the period 1997:Q1 to present. There is some tendency for compensation to accelerate at low levels of unemployment, although the dispersion in the data is large. There are just a few observations associated with high unemployment, but it is worth noting that the most recent year-over-year change in nominal compensation stands at 0.8 percent, with compensation showing quarterly declines in the first two quarters of this year. The extent to which wage dynamics may provide (or not) a cushion to the fall in inflation in present circumstances will be examined in section 4 of the paper by looking at recent micro-level wage data.

We close this section by noting that the nonlinearity in the inflation unemployment tradeoff that appears to be present in the 1954 to 1963 period may have been at work in other circumstances as well. Figure 11 shows the relationship between the unemployment rate and the 4-quarter-ahead core PCE inflation rate less the 10-year inflation expectations. The figure merges the periods 1964:Q1 to 1969:Q4 and 1981:Q1 to 1997:Q4. The SPF 10-year expectation measure was not available for the earlier period, and it is constructed by fitting the estimated evolution of the 10-year SPF over the period 1981 to 2006 to that earlier sample.[6] The estimated process for the series is highly persistent, with some additional influence coming from actual core inflation, the stance of monetary policy, and the output gap. Figure 12 shows the same type of scatter plot, but over the period 1973:Q1 to 1980:Q4. Both figures are suggestive of a sacrifice ratio increasing at higher levels of the unemployment rate. For this type of relationship to emerge, it is crucial to consider the deviation of inflation from long-run expected inflation. Through this lens, recessions are periods when inflation reliably falls, but the change in the inflation gap when the unemployment rate is high appears to be smaller than the change in the inflation gap when the unemployment rate is low. Still, it is important to stress that over the periods represented in Figures 11 and 12, long-run inflation expectations were moving, too. The change in the inflation gap may well be muted at high levels of the unemployment rate, but if long-run expectations are moving considerably so will actual inflation.

3. **The Japanese example**

As suggested above, the US has relatively little experience with inflation below two percent over the past 50 years. As a consequence, it is difficult to garner much evidence that would help policymakers gauge whether inflation may behave differently as it approaches zero. However, Japan provides a potentially useful example of a developed country that has experienced a prolonged recession and an extended period of very low (indeed negative) inflation. While one must be cautious about drawing parallels too closely between the US and Japan, Japanese experience may provide some clues as to how US inflation might evolve in the years to come.

We begin by examining key macroeconomic data for Japan over the past 30 to 40 years. Figure 13 displays some of the data central to our analysis. The top panel shows both headline and "core" CPI inflation for Japan, where core is defined as all items excluding food and energy, along with a measure of the output gap. Estimates of the output gap for Japan are hard to come by. Some estimates suggest that output returned above potential in the mid- to late-1990s, although this seems implausible. The output gap presented here was kindly provided by Sachihiro Hayashi at the Cabinet Secretariat in Japan. Unlike other published estimates, it suggests what we consider a plausible

[6] This series is constructed by the Board of Governors of the Federal Reserve System in the context of the FRB/US model. The mnemonics for this series in FRB/US is "PTR."

history for the Japanese output gap, and performs reasonably well in an array of estimates presented below.

In some of the Phillips curves estimated below, we also use estimates of real marginal cost (proxied by the labor share of income, as is conventional in the Phillips curve literature), the relative price of imported energy goods, and the exchange rate. These data are readily available and obtained from standard sources.[7] Of more interest and less readily available are measures of inflation expectations. The bottom panel of figure 13 displays both short-run (one year) and long-run (six to 10 years) inflation forecasts, taken from the Consensus Forecast database. These data are available only back to the second half of 1989, and are collected semi-annually. Other horizons (2-5 years) are included in the dataset, as are forecasts for a limited set of other macro variables. We use the output growth forecasts for forecast horizons one to five years to create estimates of the expected output gap.

Note immediately a few features of the data:

1. Inflation has clearly fallen noticeably and persistently below zero
2. This occurred despite long-run inflation expectations that have remained around one percent for the past 15 years
3. However, inflation did not "spiral downward," despite the presence of an (estimated) negative output gap throughout most of this period.
4. While the 6-10 year forecast has been wrong (on the optimistic side) for many years, the one-year forecast tracks actual inflation reasonably well. Note that in the bottom panel, the forecasts are displayed on the date the forecast is made.

Modeling inflation in Japan

Reduced-form models: Old-style Phillips curves and the Phillips correlation

Consider a standard backward-looking Phillips curve, as discussed above for the US. Inflation depends on lagged inflation, usually with a coefficient of one, and on the gap between unemployment and the NAIRU

$$\pi_t = \pi_{t-1} - a(U_t - \bar{U})$$

Lagged inflation serves as a proxy for expected inflation, or for frictions in the price-setting mechanism. The logic of this specification implies that inflation will continue to rise (fall) as long as unemployment is below (above) the NAIRU. This dynamic has been dubbed the "accelerationist" Phillips curve, as it posits a relationship between the *change* in inflation (and thus the second-difference or acceleration in the price level) and the gap. Examining the top panel of figure 13, it would seem difficult to reconcile such a specification with Japanese inflation experience of the past 30 years. Despite a persistent output gap since 1990, inflation appears to have been bounded below by -2 percent.

More formally, estimated Phillips curves show a marked deterioration in fit in recent decades. The following table provides a simple illustration of this point. The first set of columns displays the estimated backward-looking Phillips curve for Japanese core inflation from 1970 to 1989. As the estimates suggest, the unit sum restriction on lagged inflation (three quarterly lags) is

[7] See the data appendix for details.

not rejected, and the output gap displays a sizable and significant coefficient sum. The second set of columns show the estimated shift in these parameters from 1990 to the present. The effect of the level of the output gap is essentially eliminated, with great precision. The table suggests other significant shifts as well, but this change is economically quite important, and we return to it in more detail below.[8]

Table 2

Simple test for shift in Japanese Phillips curve, 1970-2009

Variable	Coefficient	*p*-value	Estimated shift, 1990-present	*p*-value
Lagged inflation	1	-	-	-
Output gap	0.39	0.00	-0.38	0.00
Rel. price of imported energy goods	2.19	0.008	-0.18	0.003
p-value for unit sum restriction	0.35		0.0051	

Figure 14 displays the in-sample and out-of-sample fit for a conventional Phillips curve over a variety of sub-samples from 1980 to the present. The Phillips curve includes three lags of core inflation, the change in the relative price of imported energy goods, and the estimate of the output gap displayed in figure 13. As the figure indicates, forecasts made employing a Phillips curve estimated through the mid- or late-1990s consistently and significantly under-forecast the level of inflation in ensuing years. The implication is that the persistent negative output gap would have implied more pronounced declines in inflation, but these evidently did not occur. Note that the trajectory of the inflation forecast is consistent with a continued downward spiral—as long as the output gap remains negative, inflation continues to fall in the forecast.

Why does the Phillips curve miss so badly during this period? As suggested by the shift test above, it appears that the basic accelerationist nature of the relationship is missing in this period. As evidence in favor of this hypothesis, consider the next figure (15), which compares the forecast for 1995-2005 made by the Phillips curve estimated from 1980-1994 with a forecast that imposes a zero sum on the output gap coefficients, leaving the rest of the estimated coefficients the same. Similarly, the cyan line in figure 14 shows the fit when data from the past 15 years are used to estimate the Japanese Phillips curve. Here, the estimated coefficients on the output gap sum to 0.016, and are not significantly different from zero.

If, as these simple results suggest, one can couch the relationship between inflation and output as a change/change equation, one could equivalently posit a level/level equation. That is,

$$\pi_t = a\tilde{y}_t$$

could be consistent with the changes relationship observed in the more recent sample

[8] At some econometric risk, we splice the OECD's estimated output gap for 1970-1979 with the Cabinet Office's gap from 1980 to the present. We allow for shifts in the pattern of lagged inflation coefficients as well, but these are not individually significant. They are constrained to sum to zero to preserve the unit sum constraint on lagged inflation. Note that unknown breakpoint tests of the Japanese Phillips curve suggest a break in the relationship in the late 1970s, well before the time of the asset price collapse that preceded Japan's great recession.

$$\pi_t - \pi_{t-1} = a(\tilde{y}_t - \tilde{y}_{t-1})$$ [9]

Such a levels specification bears the uncomfortable implication that inflation and output (or unemployment) exhibit a long-run trade-off. In most developed economies, there is little evidence of such a long-run trade-off. However, as figure 16 suggests, a levels relationship is not entirely out of the question for the past twenty years in Japan—that is, inflation remains low as long as the output gap is negative, and turns positive as the output gap does. The relationship in the figure is not tight, but it may be suggestive. Similarly, as discussed in the previous section on post-recession inflation in the US, the data might suggest a purely forward-looking rational expectations model of inflation in which inflation depends on the average (discounted) future output gaps, so that as the level of the expected output gap increases, inflation would rise. This possibility is examined in more detail below.

Models that include survey expectations

The concept of "anchored" expectations plays an important role in discussions of the expected path of US inflation. As illustrated in the preceding section, sufficiently anchored expectations can act as an offset to the downward pull from a significant output gap and/or declining marginal output costs. The following subsections examine the role of expectations in the inflation process in Japan, using both surveys of professional forecasters and rational expectations models as alternative means of measuring expectations. The goal is to derive implications from the Japanese experience that may cast light on the future course of US inflation.

Returning to figure 13, the bottom panel suggests that long-run inflation expectations (the red line) are pretty well anchored, fluctuating relatively little around their average of 1.2% over the past 15 years (shown in the dashed red line). While expectations are anchored, they appear to place more faith in the Bank of Japan than might be warranted given the constraints under which it labors: Since the inception of the survey in 1989, long-run expectations have always foreseen positive inflation in the 6-10 years ahead. That forecast has had the wrong sign for most of the past 15 years, yet even the most recent reading, in the aftermath of a prominent drop in Japanese output, expects inflation to recover to about one percent on average 6 to 10 years from now.

While the accuracy of the long-run forecast is of interest, of more relevance is the role that these expectations play in influencing realized inflation. To be sure, anchored long-run expectations have not prevented Japanese inflation from falling below zero and remaining there for the last dozen years. But interestingly, inflation has not *continued* to decline, as suggested by the accelerationist Phillips curve and discussed above. This observation raises the possibility that anchored long-run expectations have kept inflation from dropping further than they might have in the presence of a large output gap.

To examine the influence on realized inflation of both short- and long-run expectations, we estimate both VAR's and simple Phillips curves that examine the relationship among inflation, expectations, marginal cost, and important relative price shifts. The VAR results are summarized in

[9] If the true underlying relationship were a levels relationship, then a relationship estimated in changes would display an error term with a moving average error term with near-unit coefficient. We don't observe strong evidence for this in our estimates.

the table below; we focus on tests of exclusion restrictions for key variables in the inflation equation.[10]

Table 3 *p*-values for exclusion tests from inflation equation in VAR(2)						
Specification/estimation range	**1-yr. forecast**	**LR forecast**	**Output gap**	**Marginal cost**	**Exchange rate**	**Energy price**
1990-2009:2	0.04	0.62	0.15	0.35	0.31	0.31
1997-2009:2	0.27	0.61	0.59	0.87	0.69	0.36
Full sample, drop LR forecast	0.01	-	0.20	0.30	0.30	0.27
Full sample, drop exchange rate	0.02	-	0.22	0.30	-	0.35
Full sample, drop output gap	0.00	-	-	0.20	-	0.12
Full sample, drop marginal cost	0.00	-	0.14	-	-	0.32

The key results from the VAR estimates are that (a) for the full sample, the one-year forecast consistently enters significantly, while the long-run forecast never does; (b) neither the relative price of imported energy goods nor the exchange rate enters significantly; (c) the output gap and marginal cost bear only a very weak relationship to core inflation in this reduced form, contrary to standard theory.[11]

We now examine more constrained but still simple Phillips curves of the form

$$\pi_t = a\pi_t^{1y} + ba\pi_t^{LR} + c\pi_{t-1} + d\tilde{y}_{t-1} + es_t + fx_t$$

Where s denotes marginal cost and x represents relative price variables (energy prices and the exchange rate variable). The sum of the coefficients on survey expectations and lagged inflation are constrained to sum to one, although we test the significance of that constraint throughout. In preliminary tests, we find that marginal cost often enters these regressions significantly; the output gap never does, so we exclude it from the results presented below. In addition, the relative price variables are rarely significant, so we exclude them from the analysis. We present both full-sample (the table below) and rolling regression estimates (figure 17).

Note that now-common specifications for inflation typically express inflation relative to "trend inflation," see Cogley and Sbordone (2009). If inflation exhibits a trend component, perhaps attributable to a time-varying inflation goal, then it is appropriate to write the inflation model so that all inflation terms are expressed as deviations from the trend term. In the context of these datasets, the long-run inflation expectations could serve as reasonable proxies for trend inflation.

10 The standard criteria for lag length selection are ambiguous about using 2 or 3 lags. We choose two in this table, recognizing the limited degrees of freedom available for the semi-annual data.

11 Of course, these VARs are reduced-form equations, so the precise role that lagged output or marginal cost plays in a more structural model of inflation remains to be determined.

While the theory supporting such specifications is compelling, the data for this sample is much less so. Consider a trend-inflation version of the equation above

$$\pi_t - \pi_t^{LR} = a(\pi_t^{1y} - \pi_t^{LR}) + (1-a)(\pi_{t-1} - \pi_{t-1}^{LR}) + d\tilde{y}_{t-1} + es_t + fx_t$$

This equation implies restrictions on the way in which long-run expectations enter: Re-arranging this equation, it can be shown that the coefficients on long-run inflation in periods *t* and *t-1* should be *1-a* and *–(1-a)* respectively. A test of this restriction fails to reject the restriction, developing a *p*-value of 0.11. However, the coefficients for the long-run expectation enter insignificantly, both individually and jointly. This likely explains the inability to reject the opposite-sign restriction implied by the trend inflation model. [12] Altogether, it seems well justified to abstract from the issues implied by trend inflation for Japan over this period. The results that follow employ the simpler specification.

Table 4
Full-sample estimates of Japanese Phillips curve with survey expectations, 1990-2009

Coefficient	Estimate	*p*-value
1-yr. expectation	0.77	0.00
LR expectation	0.017	0.92
Lagged inflation	0.21	0.041
Marginal cost	0.16	0.007
R^2 : 0.79 *p*-value for unit sum restriction: 0.14		
Excluding long-run survey expectation		
Coefficient	**Estimate**	***p*-value**
1-yr. expectation	0.79	0.00
Lagged inflation	0.21	0.037
Marginal cost	0.16	0.005
R^2 : 0.79 *p*-value for unit sum restriction: 0.31		

The estimates in the top panel of table 4, along with the rolling regression results in Figure 17 suggest that the long-run expectation plays little or no role in the evolution of Japanese core inflation. We re-estimate these equations excluding the long-run expectation. The results are presented in the bottom panel of table 4 and in Figure 18.

The results suggest a strong role for the one-year expectation, both for the full sample and for the sub-samples represented in the rolling-regression results of figure 18. The *p*-value for the one-year expectation is always near zero. Lagged inflation plays a small but moderately significant role throughout almost all of the sample. For the last dozen years, the role of lagged inflation becomes insignificant, the marginal cost variable gains significance, and the *p*-value of the unit sum restriction drops toward zero, suggesting this restriction is violated in the most recent decade.

[12] The relative stability of the 6-10-year inflation expectation over the past twenty years also suggests that this will add little to the model. Equations estimated in this deviation form produce results that are very similar to those in tables 4 and 5. Specifically, in Table 4, the estimated coefficients on the one-year expectation and lagged inflation are 0.92 and 0.08 respectively; the coefficient on marginal cost is 0.16. In table 5, the estimated effect of the rational expectations term is once again zero, and the weights on the one-year expectation and lagged inflation are 0.7 and 0.3, respectively. The marginal cost estimate is a bit higher than the OLS estimate at 0.22.

While the backward-looking accelerationist Phillips curves of the preceding section are poor predictors for the Japanese inflation experience of the past 20 years, these survey expectations models, using real marginal cost as a driving variable, achieve more success. Figure 19 displays the in-sample fit for inflation for the full-sample estimates in table 4 (the red line) and the out-of-sample fit (the blue stars) for the last five years of the sample, for this model. The out-of-sample fit suggests the model is reasonably stable (given the limited number of semi-annual observations available to test this hypothesis), and certainly far more stable than the Phillips curve results presented in Figure 14 above.

Because the success of this model rests heavily on the inclusion of the one-year survey expectation of inflation, it is important to understand better what information is incorporated in the one-year survey. Reduced-form regression models of the one-year expectation reveal that one-year expectations are well-explained by the following regression equation:

$$\pi_t^{1y} = a\pi_{t-1}^{1y} + b\pi_{t-1} + c\tilde{y}_t + d$$

The estimated equation using semi-annual data from 1990:H1 to 2009:H2 is (HAC standard errors in parentheses)

$$\begin{array}{ccccccccc} \pi_t^{1y} & = & 0.34\pi_{t-1}^{1y} & + & 0.23\pi_{t-1} & + & 0.18\tilde{y}_t & + & 0.55 \\ & & (0.13) & & (0.12) & & (0.03) & & (0.11) \end{array}$$

The one-year expectation moves with lagged inflation and with the current output gap, and exhibits very modest persistence, as evidenced by the significant but small lagged dependent variable. The in-sample fit of this equation, which is quite tight, is displayed in the bottom panel of figure 19.

The implication of this simple depiction of one-year expectations is that expectations need not be well-tied to the monetary authority's inflation goal, nor to the long-run expectations of inflation as measured by the Consensus forecast survey. Expectations respond with a lag to realized inflation, and to the output gap. As the output gap improves, inflation improves, but with a lag, given the dependence on its own lag and on realized inflation. How well such an expectations mechanism, coupled with the inflation equation estimated above, can explain the Japanese inflation experience in a more fully fleshed-out model remains to be seen.

Structural models

While the old-style backward-looking Phillips curve fits the Japanese inflation data quite poorly, a model with survey expectations achieves some success. How well would the current generation of forward-looking models explain Japanese inflation behavior?

In this section we examine two tests of the validity of a forward-looking inflation model of the NKPC variety. First, we use the Consensus forecast data to test one of the key implications of the canonical Calvo/NKPC model. Second, we construct a simple rational expectations model that encompasses both the NKPC model—with and without indexation—and a survey expectations-based model. An interesting aspect of this last exercise is that it allows us to examine the extent to which the survey expectations mimic the expectations implied by the rational expectations NKPC model for Japanese data.

Consider the simplest version of the purely forward-looking Calvo/NKPC inflation model

$$\pi_t = \beta E\pi_{t+1} + \gamma \tilde{y}_t$$

where the variables are as defined above. Since the work of Galì and Gertler (1999), most models use a proxy for real marginal cost instead of the output gap, but here we use the output gap, because forecasts of marginal cost are not available in the Consensus Forecasts dataset. As is well known, one can solve this model forward to obtain

$$\pi_t = \gamma \sum_{i=0}^{\infty} \beta^i E_t \tilde{y}_{t+i}$$

so that the NKPC implies that inflation is in essence proportional to the discounted sum of future output gaps (apart from a contemporaneous shock, which we abstract from here for convenience).

If this model is approximately correct, then we should find a reasonable correspondence between realized inflation and a forecast of the output gap over the next several years. We construct the weighted average output gap over the next five years using the Consensus forecast dataset. Combining these with the estimate of the current output gap employed in this paper, we can compute the implied expected average output gap over the next five years, and examine its relationship to realized inflation.[13]

Figure 20 displays the core inflation rate (semi-annually) along with the average expected output gap over the current and ensuing five years, computed as described above and in footnote 4. Because a key prediction of this model is that inflation will rise (fall) as soon as expected average output rises (falls), one can informally assess the model by inspection of the figure. Consider the periods beginning in 1994, 1999 and 2003. At each of these junctures, the expected output gap consistently improved—it was still expected to be negative, but the forecasts for growth were sufficiently above (estimated) potential growth that the output gap was expected to increase. However, at these times, the inflation rate either fell for several years (1994 and 1999) or was essentially flat (2003), despite a dramatic rise in the expected gap in the five years from 2004 to 2009. This suggests that the basic mechanism implied by this model likely does not hold strongly in Japanese data. The figure suggests, however, that when the expected output gap turns decidedly positive, inflation may begin to rise, as in 1997 and 2006.

More formally, the estimated fit of the model—which we obtain by an OLS regression of core inflation on the expected average output gap—is displayed in figure 21. This exercise confirms the conclusions drawn by inspection of the previous figure: The simple model predicts that inflation will rise in the periods identified above, and particularly dramatically in the mid 2000's, but inflation simply doesn't behave in that way.[14]

Finally, we estimate a model that allows the data to choose among a purely forward-looking NKPC with rational expectations, a model that uses survey expectations as the expectations proxy, or a model in which expectations are proxied by lagged inflation. The inflation equation is

[13] This computation also requires an estimate of potential output growth for the period in question. We use the Cabinet Office's estimated potential output series, subtracting the estimate of potential growth in period t from the annualized forecast of GDP growth for periods *t+1 …t+5* to compute the increment to the output gap. We cumulate this increment over each of the five years ahead to obtain output gaps for the current and next five years, then average those to obtain the average output gap expected to prevail over the next five years.

[14] The regression model controls for changes in the relative price of imported energy goods, and for the change in the broad nominal trade-weighted exchange rate.

$$\pi_t = aE_t\pi_{t+1} + b\pi_t^{1y} + (1-a-b)\pi_{t-1} + cs_t$$

If *a=1*, *b=0* then the rational expectations NKPC is a reasonable approximation for Japanese inflation data. If *b=1, a=0,* then the model with one-year survey expectations serves as a better depiction of Japanese inflation. With *a=b=0*, the model depends only on lagged inflation, as in an old-style Phillips curve model. A host of intermediate combinations are of course possible.[15]

Because we need to generate model-consistent expectations of inflation for period *t+1*, and because we need to be concerned about the possible simultaneous determination of marginal cost (and potentially of the survey variable) and inflation, we close the model with the reduced-form equation for one-year inflation expectations as described above, and with VAR equations for marginal cost and the output gap, estimated over the full sample from 1990 to 2009. Thus the full model that we employ for estimation is (error terms are suppressed)

$$\begin{aligned}
&\pi_t = aE_t\pi_{t+1} + b\pi_t^{1y} + (1-a-b)\pi_{t-1}^T + cs_t + c_0 \\
&\pi_t^{1y} = d\pi_{t-1}^{1y} + e\pi_t + f\pi_{t-1} + g\tilde{y}_t + g_0 \\
&\pi_t^T - \pi_{t-1}^T = j(\pi_{t-1}^T - \pi_{t-1}) + k(\pi_{t-1} - \pi_{t-2}) \\
&s_t = \sum_{i=1}^{2} \mathrm{B}_i X_{t-i} \\
&\tilde{y}_t = \sum_{i=1}^{2} \Gamma_i X_{t-i}
\end{aligned}$$

where π_{t-1}^T is the lagged value of overall inflation (which performs better than lagged core inflation in initial regressions), and *X* is the vector of variables in the 2-lag VAR (inflation, output gap, marginal cost). We hold the estimated VAR coefficient Β and Γ at their OLS values and estimate the remainder via maximum likelihood, conditional on their values. The results are presented in table 5 below.

Table 5
Estimation of hybrid model for Japanese inflation, 1990-2009

Coefficient	**Estimate**	**Standard error (BHHH)**	***t*-statistic**
a	0.11	0.12	0.94
b	0.67	0.11	6.1
(1-a-b)	0.23	0.11	2.0
c	0.17	0.055	3.1
c_0	-1.1	0.34	-3.2
d	0.37	0.1	3.7
e	0.29	0.1	2.8
f	0.083	0.089	0.93
g	0.059	0.036	1.6

[15] In preliminary estimation exercises, the lag of *total* inflation, rather than core, enters significantly in the inflation equation. Thus we allow inflation to respond to lagged total inflation, and link total inflation to core inflation via an estimated error-correction equation.

g_0	0.28	0.13	2.1
j	-0.33	0.3	-1.1
k	-0.78	0.32	-2.4

The estimates find the weight on the forward-looking (model-consistent or RE) component to be small and imprecisely estimated. There appears to be a small role for indexation or lagged inflation, while two-thirds of the weight on the expectations terms falls on the one-year expectation. This result mirrors the results in the reduced-form equations presented above.[16]

In sum, Japanese inflation is best modeled not as a backward-looking accelerationist Phillips curve, nor as a forward-looking NKPC, nor by the hybrid version of the NKPC. Instead, inflation depends on the one-year survey expectations, which evolve according to the dynamics described above, and do not replicate either old-style or new-style Phillips curves.

But how much does the US look like Japan?

To this point, this modeling is suggestive, in the sense that an economically developed country like Japan clearly can experience a sustained bout of deflation, despite long-run expectations that are clearly anchored above zero. More important to the evolution of inflation is the short-run (one-year) expectation, which in turn is influenced by recent realized inflation and the output gap.

How well would such a model fit the US in recent years? Surprisingly, the answer appears to be quite well. The upper panel of the following table presents estimates of a Phillips curve that closely parallels the estimated model in table 4 for Japan.[17] Note that in initial regressions, the long-run (10-year average) expectation from the SPF survey enters with a very small coefficient (0.01), and is imprecisely estimated (standard error 0.25), also paralleling the evidence from Japan. Thus we exclude it from the regressions in table 6. The issues that arise from a consideration of so-called "trend inflation" models arise here as they did in the estimates for Japan presented above. The data for the US also fail to reject the restriction implied by the trend inflation model, with a *p*-value of 0.15. However, also like Japan, the long-run (10-year) expectations are estimated extremely imprecisely, and are not close to significant either individually or jointly. As a consequence, we also abstract from trend inflation for the US for this sample. The estimates that follow employ the simpler, non-trend-inflation specification outlined above.[18]

The bottom panel displays the determinants of the one-year expectation for US inflation, which is pivotal in determining US inflation, just as it is in Japan. The parallels are striking. The coefficients on the one-year (SPF) survey expectation and lagged inflation are extremely close to those estimated for Japan. The coefficient on marginal cost is smaller, but still likely plays an

---

[16] The specification appears well-behaved in most dimensions. For example, the estimation residuals are approximately white, with Q(12) *p*-values for the five errors of .60, .72, .89, .61 and .51 respectively. The in-sample fits are all tight. As noted above, a version of this model that expresses the inflation variables as deviations from long-run expectations does not differ qualitatively from the version presented in table 5.

[17] See the data appendix for details on data construction and sources.

[18] A version of the model in which we express inflation as the deviation from trend inflation—in this case, as a deviation from the SPF 10-year inflation expectation—produces qualitatively similar results. Specifically, the coefficients on the one-year, lagged inflation and marginal cost are 0.73, 0.27 and 0.057, remarkably close to the estimates in Table 6. Again, the relative stability of the SPF 10-year expectation suggests that this adjustment is close to subtracting a constant from the inflation measures, which of course should not much affect the results.

important role. The one-year expectations similarly depend on lags of realized inflation and the output gap.[19] Figure 22 displays the fitted values for both of these regressions, with the post-2006 out-of-sample fit for the core inflation regression depicted in the blue stars in the top panel.[20] There is no obvious sign of out-of-sample degradation.

Table 6

Estimate of inflation equations for US, paralleling Japanese specifications

1990:Q1-2010:Q2

$$\pi_t = a\pi_t^{1y} + b\pi_{t-1} + cs_t + c_0$$
$$[a+b=1]$$
$$\pi_t^{1y} = \sum_{i=0}^{4} d_i \pi_{t-i} + \sum_{i=0}^{2} e_i \tilde{y}_{t-i} + e_0$$

Coefficient	Estimate	Standard error	*p*-value
	Core inflation equation		
a	0.70	0.11	0.00
b	0.30	0.11	0.010
c	0.052	0.028	0.067
c_0	-0.22	0.093	0.022
	One-year expectation equation		
$\sum d_i$	0.66	0.032 (20.6)*	0.00**
$\sum e_i$	0.037	0.032 (1.2)*	0.00**
e_0	0.95	0.087	0.00

*t-stat. for sum of coefficients

**p-value for joint significance of contemporaneous and lagged values

In parallel with the exercises conducted for Japan, we estimate a version of the model for US data that allows the rational expectation of next period's inflation to enter the inflation equation. The model is the same as in table 5, estimated on US data. The results are summarized in table 7 below. Broadly speaking, the results are quite similar to the OLS estimates above, which is not surprising given the absence of a rational expectations effect in the equation.[21] Inflation is well-

[19] One can improve the fit of the equation by including the lagged one-year expectation, but the estimated coefficient in this case is about 0.7, which seems like an undue reliance on lagged expectations.

[20] The estimates in the table are OLS estimates. GMM estimates of the same model yield very similar results, with the sum of coefficients on the inflation and output gap terms insignificantly different from the OLS estimates. The *J*-statistic fails to reject the overidentifying restrictions. Full information estimates appear in the following table. The out-of-sample fit is generated by estimating the regression from 1990 to 2006, and then dynamically simulating the model with these estimated coefficients for 2007-2010:H1.

[21] Note, however, that the maximum likelihood technique employed in these estimates takes full account of the simultaneity implied by the presence of the contemporaneous inflation and output terms in the equation for one-year inflation expectations. The estimates for *a, b* and *c* are constrained to be greater than or equal to zero, as it is difficult to develop a structural interpretation in which any of these expectations proxies would enter negatively. Estimates for the model in which inflation variables are expressed as deviations from the 10-year expectation deliver remarkably similar results. The estimates for *a, b* and *c* are 0.00, 0.78 and 0.52 respectively. In the one-year expectation equation, the sum of the coefficients on realized inflation is smaller, and the sum of the coefficients on output larger than in the estimates presented in Table 7.

modeled as depending on the one-year survey expectation, with literally no role for the rational expectation that is consistent with this model, and a modest dependence on lagged inflation that may reflect indexation or rule-of-thumb price-setting. The reduced-form evolution of the one-year inflation expectation is quite similar to that of Japan, in that one-year expectations respond sluggishly to realized inflation and the output gap. It is of some interest that for the US, the change in the output gap may be as important an influence as the level of the output gap in influencing one-year expectations, as evidence by the small sum of the lag coefficients, but very high joint significance of the lags of output.

Table 7
Estimation of hybrid model for US inflation
1990:Q1-2010:Q2

Coefficient	**Estimate**	**Standard error (BHHH)**	***t*-statistic**
a	0.00	0.018	0.056
b	0.74	0.13	5.6
(1-a-b)	0.26	0.034	7.7
c	0.048	0.039	1.2
c_0	-0.19	0.060	-3.2
$\sum d_i$	0.69	0.047 (14.6)*	0.00**
$\sum e_i$	0.056	0.038 (1.5)*	0.00**
e_0	0.89	0.99	0.90

*t-stat. for sum of coefficients
**p-value for joint significance of contemporaneous and lagged values

The dynamics of inflation for both the US and Japan may be summarized as follows:

1. Inflation is not well-modeled by an old-style accelerationist Phillips curve, or by a forward-looking, rational expectations New-Keynesian Phillips curve.
2. Inflation may be sensibly modeled as depending on the one-year (survey) expectation of inflation with a weight of between 2/3 and 3/4, and lagged inflation with a weight of between 1/3 and 1/4.
3. Inflation also depends on marginal cost.
4. The one-year expectation responds sluggishly to current and lagged realized inflation rates, and to current and lagged output gaps.

Of interest is what are the implications of such a model for the evolution of inflation in the US in current circumstances, in a very low-inflation environment when monetary policy is constrained by the zero lower bound. It is conceivable that slowly-moving one-year expectations, combined with a large and persistent output gap and depressed marginal cost, could lead to a very slow adjustment of inflation towards its long-run target. We next simulate a model that incorporates these dynamic properties of inflation, augmented by a conventional monetary policy rule in which the policy rate is constrained not to fall below zero.

The model comprises the inflation and expectation equations outlined above, plus a conventional policy rule that imposes the zero lower bound on the short-term policy rate r_t[22]

$$r_t = \max\left\{\rho_r r_{t-1} + (1-\rho_r)[a_\pi(\pi_t - \overline{\pi}) + a_y \tilde{y}_t + \overline{r}], 0\right\}$$

where $\overline{\pi}$ is the central bank's (assumed constant) inflation goal, and the parameters of the policy rule are set to $a_\pi = 2, a_y = 1.0$, a bit more aggressive than the canonical Taylor (1993) parameters, but consistent with empirical estimates that include more recent years. The inflation goal is set to two percent, and the equilibrium real rate is two percent, so that the equilibrium nominal policy rate is four percent. The model is closed with a "hybrid" I—S curve of the form

$$\tilde{y}_t = \omega_b \tilde{y}_{t-1} + \omega_f E_t \tilde{y}_{t+1} - \sigma(r_t - E_t \pi_{t+1} - \overline{\rho}) + u_t$$

$$u_t = \rho_y u_{t-1} + \eta_t$$

Here $\overline{\rho}$ is the equilibrium real rate of interest, and σ is set to 0.05, consistent with many estimates in the DSGE literature. We eliminate the constants from the regression specifications for inflation and expected inflation from table 7 above; the equilibrium levels of inflation and expected inflation are given by the inflation target in the policy rule.[23]

The model is simulated with initial conditions that are intended to mimic macroeconomic conditions at the beginning of this year, *viz*

- Inflation is one percent;
- The funds rate is just above zero;
- The output gap is a considerable seven percentage points;
- Short-term inflation expectations are also one percent.

The evolution of the key variables in this simulation depends importantly on the specification of the output equation. If one assumes that the shock *u* is uncorrelated over time, and that output is purely forward-looking, then the output gap jumps immediately above zero, inflation begins to rise, and the policy rate is generally not constrained by its zero lower bound. However, in more realistic depictions of output—either a highly-correlated error term or a substantial weight on lagged output, ω_b --the combination of a high dependence on short-term expectations and the somewhat sluggish adjustment of expectations to current and lagged conditions implies the outcome depicted in figure 23.[24]

This outcome is striking, as it mimics to a great extent the historical experience of Japan, using a model that fits US inflation experience of the past 20 years quite well. The model suggests a prolonged period for which inflation is negative and the policy rate is pinned at the zero lower bound, as the output gap recovers gradually. At first, inflation falls with one-year inflation expectations, which decline as a result of the large and negative output gap. Over time, as the output

---

[22] The zero lower bound is imposed by using the Heaviside function H(x) = 1 if x >= 0, = ε if x < 0. Because the solution procedure utilizes analytic derivatives, it is convenient to use a function with a well-defined derivative. The derivative of the Heaviside function is the Dirac function.

[23] The one-year inflation equation is modified so that the sum of the coefficients on lagged inflation is one. This modification, together with the specifications for the policy rate and output, ensures that the steady-state for the model implies an inflation rate that converges to the central bank's target rate.

[24] At the time of writing, most private forecasters envision a protracted period of elevated unemployment that extends several years into the future. This would normally imply a significant output gap over the same period, even with due consideration to the possibility that structural unemployment has risen and/or potential output has fallen.

gap improves, inflation remains low because of the effect of lagged inflation on the one-year expectation. While inflation and expectations ultimately return to the two percent inflation goal in the model, the simulation suggests that return to long-run equilibrium could take nearly a decade. Perhaps Japan and the US aren't so different after all.

What we learn from the Japanese example

While we should apply lessons from Japan to the US experience with caution, the Japanese data may serve as a cautionary tale for US policymakers. First, long-run expectations, while reasonably well-anchored in both the US and Japan, do not appear to exert a strong influence on the evolution of core inflation over the past 20 years. It may be the case that anchored long-run expectations have prevented inflation from spiraling downward in Japan, but it is difficult to develop empirical evidence in favor of that proposition.

Second, one-year expectations do appear to act as significant determinants of inflation in both the US and Japan. In both countries, those expectations respond gradually to the estimated output gap, and to recent inflation. Taken together, these influences can lead to a very slow adjustment of inflation to output and marginal cost: As the output gap widens, one-year expectations fall, pulling down actual inflation. Even as the output gap improves, the influence of recent inflation on expectations continues to depress expectations, in turn slowing inflation on its eventual return towards the central bank's inflation goal. In Japanese data, and in the simulation for the US, this can take a decade or more.

Because the process that generates one-year expectations is a reduced form, the models estimated on Japanese and US data do not have the status of more theoretically-grounded, micro-founded models of inflation. But the models appear to fit historical data quite well, appear reasonably stable (given the limitations to testing stability inherent in our relatively short samples), and also empirically dominate purely forward-looking models, hybrid rational expectations models, and old-style accelerationist Phillips curves for the samples examined. Thus, their implications should be taken as serious cautionary tales for current US policy.

Finally, as section 2 emphasizes, the current crop of theoretical models do not necessarily imply that well-anchored expectations—even perfectly well-anchored expectations—preclude a significant drop in inflation in the presence of very low costs and a very large output gap. The extent to which inflation drops in these models depends critically on the extent to which price-setters are forward-looking in the strict sense implied by rational expectations in these specifications. Historical evidence on the evolution of inflation following recessions suggests that this very forward-looking characterization of inflation is less likely to hold that a characterization in which inflation exhibits a significant backward-looking component.

It is worthy of note that while the models examined in section 3 empirically generate the "inflation-floor" outcomes that characterize recent Japanese economic history, they leave open the fundamental puzzle of why inflation would behave as if it had a floor or lower bound. The next section examines a possible explanation for the presence of such a floor.

4. Wage Rigidity and Its Effects on Inflation

As the previous section highlights, there is mounting evidence that the inflation process may change as the rate of inflation approaches zero. Until the recent experience in the US, Japan represented the best test case study of this possible effect. Since Japan's real estate collapse, short run inflation expectations and the output gap appear to affect its inflation rate, but the role of longer-run expectations is difficult to discern. The US has been flirting with low levels of inflation since the late 1990s, with inflation near zero twice in the past decade, but never (to date) dipping below zero. In contrast to the Japanese evidence, many economists in the US suggest that the relatively stable behavior of US inflation may be attributed to the "anchoring" effect of fairly stable long-run inflation expectations. Several hypotheses have been forwarded for the stability of long-run expectations, such as increased Fed credibility, a smaller effect of the output gap on inflation, inattentiveness to inflation on the part of price-setters when inflation is near zero, or just good luck. This section examines a more traditional explanation, downward nominal wage rigidity.[25]

Speculation about nominal wage rigidity has a storied past in macroeconomics, both theoretically and empirically. If present, wage rigidity would become more problematic as inflation fell toward zero. When nominal wages and inflation are expanding robustly, significant reductions in the real wage are possible without declines in the nominal wage. As inflation approaches zero real productivity-adjusted wages can only fall through nominal wage declines.[26] If nominal wages are rigid downward even in the face of significant output gaps, inflation may become rigid downward. If the stickiness of wages is well known, inflation expectations may also become rigid downward at low rates of inflation. Hence, nominal downward wage stickiness can help produce a very different inflation dynamic. This section examines some preliminary evidence on this channel.

Results from Data on Individuals

Essential to this explanation is the idea that workers refuse to accept, or firms refuse to provide, nominal wage cuts. How much evidence is there that this frequently accepted idea is true? The evidence at the level of individual wages is fairly convincing, but the implications for inflation are more complicated. Following many researchers, and to illustrate the results that have been published over the years, we use the Panel Study of Income Dynamics (PSID) to illustrate the answer most frequently found in this literature. Because this is one of the most comprehensive surveys of individuals, many examinations of downward wage rigidity began here. The PSID is, now, a biannual survey that asks the same group of households questions about their employment, income, wages, wealth, and a host of other questions. The nominal wage of the household head can be culled from the data, as can whether the worker has stayed in the same job.[27] Figure 24 provides the distribution of wage growth for a representative year for the cross-section of household heads who have stayed in the same job. The year is a more recent one since a binding nominal wage floor would be more visible at a low level of inflation.

[25] The attribution of inflation's stability to the supposed anchor provided by long-run inflation expectations (in both Japan and the US), even as actual inflation remains below that expectation for long stretches of time, is a profession of faith in the anchoring power of the central bank's long-run inflation objective. Interestingly, adherents to this faith quickly turn into apostates once the inflation rate rises *above* the long-run expectation for any length of time.

[26] With productivity growth, a firm's costs and prices may fall even without cutting nominal wages. In the extreme, the lower bound on inflation would look something like the negative of the growth rate of productivity.

[27] Often determining the hourly wage and whether they have changed jobs within an institution is problematic.

Two elements of the distribution in figure 24 bear on the presence of absence of downward wage stickiness. The most obvious is the shape of the distribution around zero. It is very difficult to look at figure 24 and not conclude that there is a rather significant discontinuity at zero. Excluding zero, the mode occurs around 6 percent and its frequency is just over 4 percent. Zero wage changes represents close to 14 percent of the sample. That 14 percent is roughly 14 times the frequency to *either* side of the zero change. While a number of possible biases in the data have been examined (integer bias, reporting bias, etc.), wage data from the PSID indicate that something hinders wages from declining. The second piece of evidence is the size of the tail to the left of zero. Comparing the mass of the distribution in the tail to the left of zero to the size of the tail an equal distance from the mean on the positive side is another measure of downward rigidity. The tail to the left of zero is roughly one half the size of the tail on the other side of the distribution.

Researchers have used these results to emphasize somewhat different conclusions. Some see the skewness and the spike at zero in the wage change distribution as convincing evidence that nominal wages are sticky downward. Others point to the fairly significant number of nominal wage declines as evidence that, although not perfectly flexible, nominal wages can and do fall. One complicating issue, of course, is that no one knows what these distributions *should* look like. Standard economic theory suggests that the distribution of the prices and marginal products in these jobs and industries should affect the distribution of their nominal wages. Those distributions need not be normal; if they are skewed to the right, then nominal wage distributions may tell us more about the distribution of productivity growth by industry and job than they do about nominal wage stickiness. These complications are non-trivial and lightly addressed in the extant literature. Still, we will take it as given that the spike at zero is fairly convincing evidence of downward wage rigidity, at least for those workers who remain in the same job.

Workers who change jobs are not included in the distribution, in part because their productivities could have changed along with their jobs. Examining the distribution of wage changes for job-changers, in figure 25, produces a very different story. The year chosen might be critical here since people switch jobs for different reasons. Workers leave jobs to take better jobs, and they lose jobs and find other work. The wage experience of these two groups might differ. In the aggregate unemployment data, job losers tend to dominate job leavers, particularly during recessions and the early stages of a recovery, but these aggregate data do not include job-to-job changes. If job losers dominate the job changers in the job-shifters sample in the PSID, then one might expect wages to fall. In fact, the distribution of wage changes for job-shifters is far different. There is still a large spike at zero wage change, explained only in small part by minimum wage workers, but the distribution looks more uniform on each side of zero. The sample of job shifters for any one year in the PSID is very small, so caution should be taken in placing too much emphasis on these numbers.

Still, the difference in these two distributions might carry important implications for inflation dynamics. For modeling inflation, the relevant question is not whether individual workers are reluctant to accept cuts in their nominal wages. The question is whether firms face impediments to cutting their labor costs when wages and inflation are near zero. For example, firms may be able to fire and hire workers to shift the wage composition of their employees towards lower wages, even if the wages of workers who remain with the firm are resistant to decline.[28] The resulting decline in labor costs could allow the firm to reduce (or reduce the rate of increase in) goods prices. As a result, wages could be completely rigid downward for people who stay in the same job, but as long

[28] At this point we are abstracting from possible differences in the productivities of job shifters versus job stayers.

as wages can fall for new entrants and job losers, the wage bill as a whole can decline. The fall in labor costs would help put downward pressure on prices, so as long as the labor market was not moribund, labor costs and prices could continue to decline even when inflation approaches zero.

Hence, the potential damping of downward inflationary pressure is likely better measured by changes in the *wage bill*, for a given level of employment. Thus the PSID data (or any data on individual wage changes) may not be the best data to examine whether wages contribute to an inflation floor. For this reason, we focus in this section on data for labor costs to employers. Looking at employer data also avoids a host of problems with the PSID data highlighted throughout the literature (Akerlof, Dickens, and Perry (1996); Bound (1994); Gottschalk (2005); McLaughlin (1994); …). Many of the well-known biases and errors in the reporting of own wages are avoided when the data are collected from the employer.

Employer data

There are two large sources for employer data: The compensation data collected by the BLS for the Employment Cost Index, the National Compensation Survey (NCS), and the data collected by the BLS for the Occupational Employment Statistics (OES) survey. Several researchers have used establishment data to examine wage rigidities. Each of these measures has its drawbacks and strengths. Wilson (1999) examined compensation data for a small number of establishments. Lebow, Saks, and Wilson (2003) expanded the sample by utilizing the NCS data, which cover the entire compensation package for a job. The drawback is that the sample is relatively small. In the NCS, the BLS surveys firms for a random selection of about four people in each establishment and obtains comprehensive information only for those four. Information about the wage bill of the establishment or the industry is not collected.

The OES, on the other hand, is a more complete survey of workers in these establishments. The survey collects the number employed in each job. Unfortunately, many of the details about the compensation package are omitted. For example, the OES does not collect the exact wage; it asks for the wage interval for each worker. Essentially the survey provides a sequence of wage intervals and asks how many workers in a given job are in each interval. Not only does it fail to get the exact wage, but because of the large number of establishments and individuals, it surveys establishments over a 3 year cycle and uses a weighted average to calculate the wages and employment in a given job. As a result, this measure should be slow to register wage changes.

Figures 26 presents the distribution of wage changes by job in the OES, and figure 27 presents the distribution of wage changes for employed workers in the OES. A job is defined as one of approximately 800 occupations in one of 300 industries. The wage provided is for each occupation in a given industrial classification; for example, the wages of "research assistants in central banking" might be included.

The distributions shown in figures 26 and 27 provide a stark contrast to those found at the individual level. Two years are presented to show that the recent data are not an anomaly due to the recession. Where the PSID revealed significant evidence of a discontinuity at 0, the OES jobs data, figures 26a and 27a, seem to show none at all. Because the jobs that are experiencing an actual wage decline may have a disproportionately small number of workers, figures 26b and 27b show the employment frequency of each wage change. While the variance of the distribution declines when measured by employment, its basic balance remains about the same. The negative tail is substantial, it looks like the positive tail, and there is no large spike at zero.

How do we explain these two apparently conflicting results?

The previous literature viewed the fact that the establishment data only provide the wage of the job, not the person as a major drawback. The problem, given their interest in the wage flexibility of people on the job, was that decreases in the wage for a given profession could be due to either the wages of those employed in that profession declining or the composition of those in the job changing. The firm could be replacing high wage workers in a given profession with lower wage workers in that profession, thus lowering the average wage in that job. For tests of nominal wage rigidity at the individual level, this composition effect is a major drawback to using establishment data. However, since this paper is only concerned with labor costs and their effects on prices, it matters little if the wage bill declines because the wage fell or because lower wage workers were brought in to do the work. The wage bill comparisons are valid as long as the firm is changing the composition of its labor force so that productivity-adjusted wages are declining. If this were not the case, the firm could increase profits by re-allocating into higher wage workers whose productivity more than compensated for their higher wages. Our assumption is simply that firms generally adjust the composition of their work force in a way that does not increase productivity-adjusted wages.[29]

It would be of interest to decompose a decline in wages in a profession into the decline in the workers' wages for a fixed worker composition and the shifts in the composition of the workers in a given job. Since the BLS has the number of workers in a given profession in each wage interval for all the establishments, they have data that might shed a great deal of light on that issue. By examining the distributions of workers in a given profession for a given establishment, and how they change over time, we might be able to detect whether compositional effects play a significant role in these labor cost reductions. Given the lags in the BLS research application process, that work will come later.

Preliminary Findings

In the absence of perfect data, this section explores the available data for evidence that is consistent or inconsistent with the hypothesis that compositional shifts allow firms to continue to cut costs and prices when inflation gets low, even though an individual's wages at the firm may remain downwardly rigid. One obvious question concerns the leeway firms have to shift costs with compositional changes; is there room for much average wage movement given the wage distribution in a given job? If the spread of the wages within the firm paid to people doing the same job is high, then the firm might have more flexibility to lower the wage bill with shifts in the composition of workers. The BLS provides the average wage in each quartile of workers in almost every job.[30] Figure 28 provides the distribution of the wage ranges within a job, an occupation/industry cell, between the upper quartile wage and the lower quartile wage. The figure reveals that these professions have a great deal of room to lower the average wage for a given job if they changed the

[29] This assumption may not hold for every change a firm makes in its worker/wage composition. For example, a firm that changes the mix of output may not always alter its workforce so as to lower productivity-adjusted wages. But for the low-inflation-recession periods that we are focused on, the assumption that firms are attempting to lower productivity-adjusted wage seems plausible.

[30] These distributions are so condensed in some jobs that no range is provided.

composition of workers in that job. The modal difference in wages paid to the highest and lowest paid groups within a job in an industry is 6 dollars an hour. Certainly some of this difference must be explained by regional factors, or important differences in human capital. But given that the average wages for each job seem to decline fairly often, and individual wages do not, such compositional shifts are something profit maximizing firms seem comfortable doing.

Unfortunately, the reason for these wage declines is unclear, and thus their effect on the prices firms charge is unclear. These wage declines could be due to falling demand for that occupation or that industry or rising supply of the labor force in that job. If it were a rising supply of labor, or a secular decrease in demand, one would expect a secular change in the wage, which might produce a sustained drop in the wage of these workers. Figure 29 provides the distribution of the number of times a job experiences a decline in wages. Over the 7 year sample, from 2002 to 2009, only 20 percent of the jobs experienced no declines in average wages. Most jobs experienced one or two years of declining wages. These data are consistent with a conclusion that only a small share of negative wage changes in the OES is due to secular declines in an industry or increases in labor supply to that profession.

The distribution in Figure 29 might raise doubts about the data. The relative frequency of wage declines in a job seems high, which makes one wonder if it is an artifact of the survey methodology. However, the sample is constrained to a time period where inflation was historically low and twice flirted with zero. However, there are details of the survey methodology that might explain this large percent of jobs experiencing at least 1 wage cut. The BLS surveys a third of the establishments each year and averages over the three year period, so it includes all the establishments (with assumptions made about wage inflation for the non-current group). Perhaps the annual surveys are biased toward higher or lower cost firms in the sample. A high one this year might be offset by a low one next year. Yet, averaging over 3 years should take care of most of that variation, and the size of even the annual survey is quite large – about 400,000 establishments.

To examine the effects of these two deficiencies in the OES data, we consider the government sector, for which data are collected in a census every year, not averaged over three years, and for which the survey collects the exact wages, not wage intervals. If either of these two drawbacks to the non-government part of the survey caused the appearance of negative wage changes, we would not expect to see them in the government data. Figure 30 presents the wage changes by profession in the government sector. The distribution looks similar to that for the economy as a whole. More complete data are required before we can be confident about the causes of these wage declines. In the meantime, the similarity of the government distribution with the distribution for the economy as a whole suggests that quirks in the survey methodology are not the origin of the frequent declines in a job's average wage.

Several preliminary regressions are estimated in order to clarify the interaction between wage declines for a given job and the employment changes in that job. The basic story is that during recessions firms that successfully cut labor costs must be firms that successfully shift the composition of labor in a given job from high wage workers to low wage workers.[31] Those firms and jobs will experience stronger employment growth than firms or professions that cannot create this shift. The hypothesis is that firms with more employment growth in a job will experience more

[31] Firms could also substitute labor between jobs rather than within jobs, but that would not necessarily affect the average wage in a given job, although it could lower the wage bill.

flexible nominal wages as the compositional shifts occur. To examine if the data are consistent with this hypothesis we estimate,

$$\dot{W}_{it} = \alpha_0 + \alpha_1 \dot{L}_{it} + \alpha_2 \operatorname{Re} cess_t + \alpha_3 \operatorname{Re} cess * \dot{L}_{it} + \alpha_4 Spread_t + \mathrm{B} X_t + \varepsilon_{it}$$

where Recess is a dummy for whether the economy is in an employment recession, W_{it} is the wage growth of job i at time t, L_{it} is the employment growth in that job, Spread is the inter-quartile range of the wages for that job as a percentage of the mean wage for that group, and X_t is a vector of year dummies.[32] If during recessions firms increase relative employment in lower-wage jobs, the regression should develop a negative coefficient on the interaction term. The results are presented in Table 8. Although the regression is crude, the results suggest that employment growth may be capturing the effects of compositional shifts, particularly during recessions. During recessions, above-normal employment growth appears concentrated in occupation/industry pairs that make wage changes more negative, $\alpha_3<0$. This result suggests that compositional re-allocations are used by firms to lower wages during recessions.

Table 8
The effect of employment growth on wages in each job, with recession effect

Variable	Coefficient	Std. Error	*p*-value
$L_{i,t}$	-0.000054	.00021	0.80
Recession	-1.25	.043	0.000
Recession* $L_{i,t}$	-0.0018	.00068	0.010
$Spread_t$	0.0095	.0012	0.000
dummy09	0.0029	.41	0.94
dummy08	1.01	.042	0.000
dummy06	-0.26	.045	0.000
dummy05	-0.88	.049	0.000
dummy04	-0.78	.048	0.000
_cons	2.82	.061	0.000
Number of obs: 222,934 R^2 : 0.0071			

Finally, we look to see if these changes in a sector's wage bill are correlated with price changes in the sector. We use prices from the CPI and PPI for about 180 industries at the 4-digit NAICS level. The OES data on employment and wages in the corresponding industry are used to create the wage bill for that industry. The basic regression estimated relates the inflation rate in an industry to the growth rate of the wage bill in that industry.

$$\pi_{jt} = \alpha_0 + \alpha_1 \dot{L}_{j,t-1} + \alpha_2 \dot{WB}_{j,t-1} + \alpha_3 \pi_{jt-1} + \beta X + \varepsilon_{jt}$$

where π_{jt} is the change in the inflation rate in industry j at time t, $L_{j,t-1}$ is the change in employment growth rate in that industry at time t-1, $WB_{j,t-1}$ is the change in the growth of the wage bill in that industry at time t-1, and X is a vector of control variables including year dummies and some industry dummies. Because the rate of inflation may also differ across industries due to unobserved

[32] The spread is included because the wider the range, the more room the firm has to lower the average wage of a job with compositional changes.

differences in trend productivity, we difference the industry level inflation equation to difference out the effect of (presumably slow-moving) trend productivity differences. The very tentative results in Table 9 provide modest support for the notion that declines in the wage bill due to reductions in the average wages reduce pressure on prices.

Table 9 Inflation and OES Wages			
Variable	**Coefficient**	**Error**	***p*-value**
C	-1.16	.31	.000
$L_{i,t}$	-.26	.14	.058
$WB_{i,t-1}$	.27	.13	.039
$\Pi_{i,t-1}$	-.66	.048	.000
Number of observations: 682 R-squared = .22			

Summary

Other things equal, downward nominal wage rigidity would tend to offset other downward pressures on inflation, which could alter the dynamics of inflation as average wage growth gets closer to zero. Data on individuals suggest there is nominal wage rigidity on the downside for individual job-stayers. Workers who lose their jobs face a different reality. They, and new entrants, often find work at a reduced wage, even if it is in the same job but a different firm. The question at hand is how important the re-allocation affect is during periods of low inflation. While our data to examine this question are imperfect, the apparent flexibility of wages in the establishment data suggests that this re-allocation is occurring, particularly during recessions. Evidence from the OES is consistent with the idea that wage bills are more flexible than individual wage rates, which could eliminate one potential reason to expect a floor near zero on inflation.

5. Conclusions

The recent decline of inflation in the US raises a number of questions of some relevance to policymakers. Will inflation continue to fall, or rise towards the central bank's implicit inflation goal? Will well-anchored inflation expectations mitigate or completely offset other disinflationary forces? If inflation falls, will it behave as if it has a lower bound, as has been the case for Japan? If so, why? To what extent is the experience of Japan a useful guide for current US experience? What models best serve to capture the dynamics of US inflation in current circumstances?

The paper provides only partial and tentative answers to these questions. In sum, the paper suggests that while inflation has not fallen as quickly as some empirical models would have suggested, it is quite common for inflation to fall in the aftermath of a recession. How far one should expect it to fall depends in part on how forward-looking one believes price-setters to be. The evidence over most postwar recoveries suggests that a purely forward-looking model of inflation, which would imply a gradual return towards the central bank's target once output and marginal cost are expected to improve, derives less support than a hybrid model that includes some backward-looking or indexation behavior. If one accepts such a hybrid model as a reasonable approximation to the behavior of inflation, then modest additional declines in inflation are likely. The extent of the decline may well be smaller than the large estimated output or unemployment gap would suggest, as the paper develops some evidence that the sensitivity of inflation to activity gaps is typically smaller in the wake of large recessions.

The paper examines the example of Japan, as it has (sadly) a much longer experience with below-potential output and negative inflation. The evidence from Japan suggests that a traditional accelerationist Phillips curve does not provide a good description of inflation's behavior in recent decades. While long-run expectations have been relatively stable and positive throughout, they show little or no relationship to the evolution of inflation over the past two decades. The paper develops a model that ties inflation to short-term inflation expectations, which in turn depend on realized inflation and the output gap. It is of some interest and concern that a very similar model appears to fit recent US inflation experience quite well. On the positive side, the model suggests that US inflation might be subject to a lower bound, much as it has been in Japan. On the negative side, there is little in the Japanese data that would give us confidence that the US will not experience deflation.

With regard to the lower bound on inflation, the paper examines evidence regarding the extent to which downward nominal wage rigidity—with respect to a firm's wage bill, rather than to individual wages—might slow or stop the decline of inflation. While our data are not completely up to the task, the evidence so far suggests no obvious downward rigidity in the wage bill of the firm, despite the significant evidence suggesting downward rigidity in individual wages. Thus one may not be able to take comfort in the buffer to disinflation provided by this type of downward nominal wage rigidity.

Overall, we take this evidence as a cautionary tale for policymakers. While one should not conclude that inflation must fall, or that deflation is inevitable, the evidence that anchored expectations or wage rigidities will halt the decline of inflation is also weak. Thus the risk of further declines in inflation is a risk that merits serious attention.

References

Akerlof, George A., Dickens, William, Perry, George L. (1996), The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity*, Studies Program, The Brookings Institution, vol 27(1996-1), pp. 1-76.

Atkeson, A., Ohanian, L.E., 2001. "Are Phillips curves useful for forecasting inflation?" Federal Reserve Bank of Minneapolis Quarterly Review **25**, pp. 2–11.

Barnes, M., F. Gumbau-Brisa, D. Lie, and G. Olivei (2009). "Closed-Form Estimates of the New Keynesian Phillips Curve with Time-Varying Trend Inflation," FRB Boston Working 9-15, November

Blanchard, O. and J. Galí, (2010) "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment," *American Economic Journal: Macroeconomics* **2** (April), pp. 1–30

Bound, J., Brown, C., Duncan, G., and W. Rodgers (1994), "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics*, University of Chicago Press, vol. 12(3), pp. 345-68.

Calvo, Guillermo A. (1983). "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics,* **12**, pp. 383 398.

Cogley, T. and A. Sbordone (2008). "Trend inflation, indexation, and inflation persistence in the New Keynesian Philips curve." *American Economic Review,* **98**, No. 5, pp. 2101-2126.

Fuhrer, J., Kodrzycki, Y., Little, J. and G. Olivei (2009), "The Phillips curve in historical context," in **Understanding Inflation and the Implications for Monetary Policy A Phillips Curve Retrospective**, proceedings of the Federal Reserve Bank of Boston's 2008 annual economic conference, MIT Press, Cambridge, MA, pp. 1-68.

Fuhrer, J. and G. Olivei (2010), "The Role of Expectations and Output in the Inflation Process: An Empirical Assessment," Public Policy Brief No. 10-2, Federal Reserve Bank of Boston.

Galí, Jordi and Mark Gertler (1999). "Inflation dynamics: A structural econometric analysis," *Journal of Monetary Economics*, **44,** pp. 195-222.

Gottschalk, Peter (2005), Downward Nominal Wage Flexibility –Real or Measurement Error?" *Review of Economics and Statistics*, 87(3), pp. 556-568, August.

Lebow, D., Saks, R., and B. Wilson (2003), "Downward Nominal Wage Rigidity: Evidence from the Employment Cost Index," *Advances in Macroeconomics*, Berkeley Electronic Press, vol. 3.

McLaughlin, Kenneth, J. (1994), "Rigid Wages?" *Journal of Monetary Economics*, vol. 34(3), pp. 383-414, Dec.

Phillips, A.W. (1958). "The Relation between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957," *Economica*, **25**, No. 100 (November), pp. 283-299.

Romer C. and D. Romer (2002), "A Rehabilitation of Monetary Policy in the 1950's," *American Economic Review Papers and Proceedings,* **92** No. 2, (May), pp. 121-127.

Stock, J. and M. Watson (2009). "Phillips Curve Inflation Forecasts," in **Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective,** proceedings of the Federal Reserve Bank of Boston's 2008 annual economic conference, MIT Press, Cambridge, MA.

Stock, J. and M. Watson (2010), "Modeling Inflation After the Crisis," Prepared for the Federal Reserve Bank of Kansas City Symposium, "Macroeconomic Policy: Post-Crisis and Risks Ahead," Jackson Hole, Wyoming, August 26-28

Wilson, B. (1999), "Wage Rigidity: A Look inside the Firm," *Finance and Economic Discussion Series* 1999-22, Washington: Board of Governors of the Federal Reserve System.

Data Appendix for sections 2 and 3

Japanese data		
Variable	**Definition and source**	**Mnemonic**
Core inflation	Consumer price index less food and energy, 400 times the log change in the price index, Min. of Intern. Affairs & Communic.	π_t
Total inflation	Consumer price index, 400 times the log change in the price index, Min. of Intern. Affairs & Communic.	π_t^T
Output gap	As computed by the Japanese Cabinet Office	$\tilde{y}_t$
Real marginal cost	Labor share of income (equivalently, nominal unit labor costs divided by the price level), OECD	s_t
Exchange rate	Broad nominal exchange rate, JP Morgan	-
Relative price of imported energy	Petroleum, coal and gas import price index, log difference between index and Japanese total CPI, Bank of Japan	-
1-yr. inflation forecast	Consensus Forecasts, one-year ahead forecast, *t* is date forecast is made, frequency semi-annual.	π_t^{1y}
6-10-yr. inflation forecast	Consensus Forecasts, 6-10-year ahead forecast, *t* is date forecast is made, frequency semi-annual.	π_t^{LR}
US data		
Core inflation	Consumer price index less food and energy (BLS), annualized quarter-to-quarter percentage change	π_t
Total inflation	Consumer price index (BLS), annualized quarter-to-quarter percentage change	π_t^T
Output gap	One hundred times the log difference between real GDP (BEA) and the HP-filtered trend in real GDP (λ=1600)	$\tilde{y}_t$
1-yr. inflation forecast	Survey of Professional Forecasters, four-quarter-average inflation forecast (FRB Philadelphia), *t* is date forecast is made	π_t^{1y}
10-yr. inflation forecast	Survey of Professional Forecasters, average over next ten years inflation forecast (FRB Philadelphia), *t* is date forecast is made	π_t^{LR}
Real marginal cost	Labor share of income (100 times log nominal unit labor cost less log of the implicit price deflator, nonfarm business)	s_t
All data, with the exception of the Japanese output gap, the SPF and Consensus inflation forecasts, are obtained from Haver Analytics.		

Figure 1
Inflation tends to fall for several years following the end of a recession

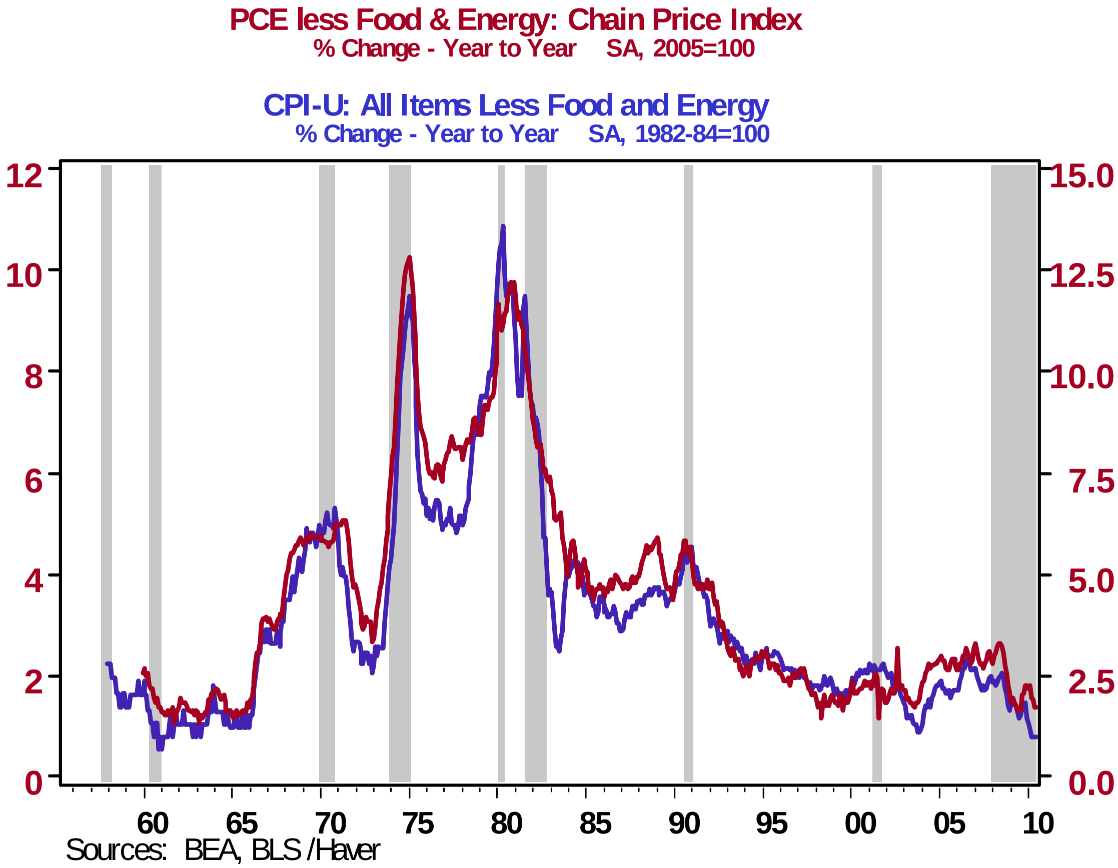

Figure 2
Inflation following a large recession, hybrid model

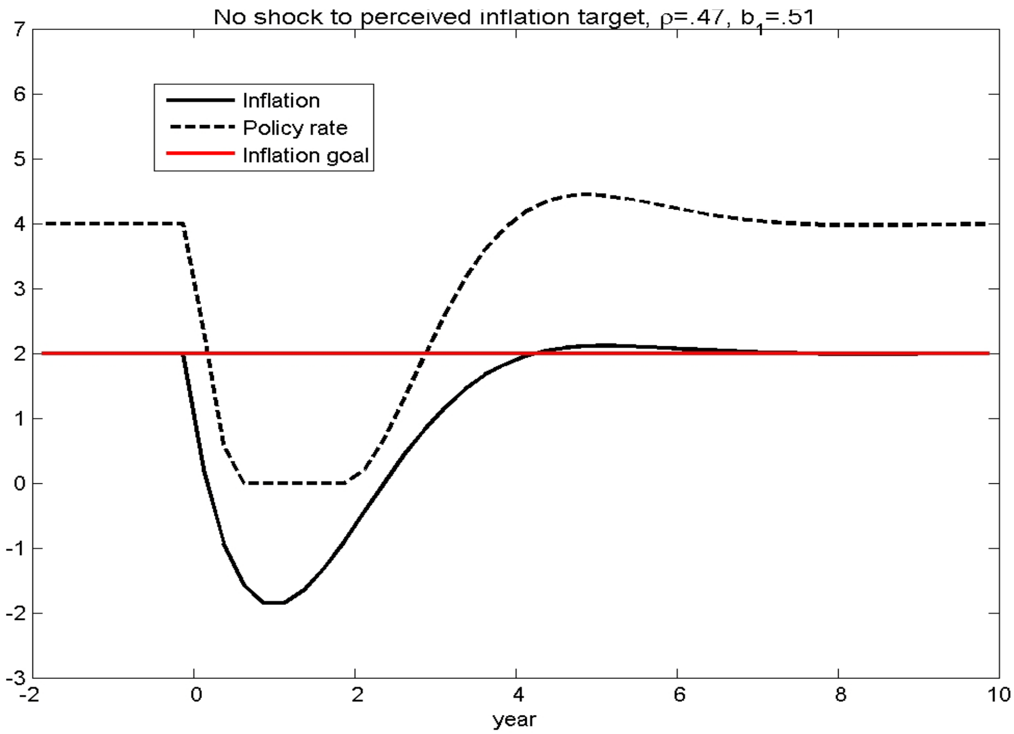

Figure 3
Inflation following a large recession, forward-looking model

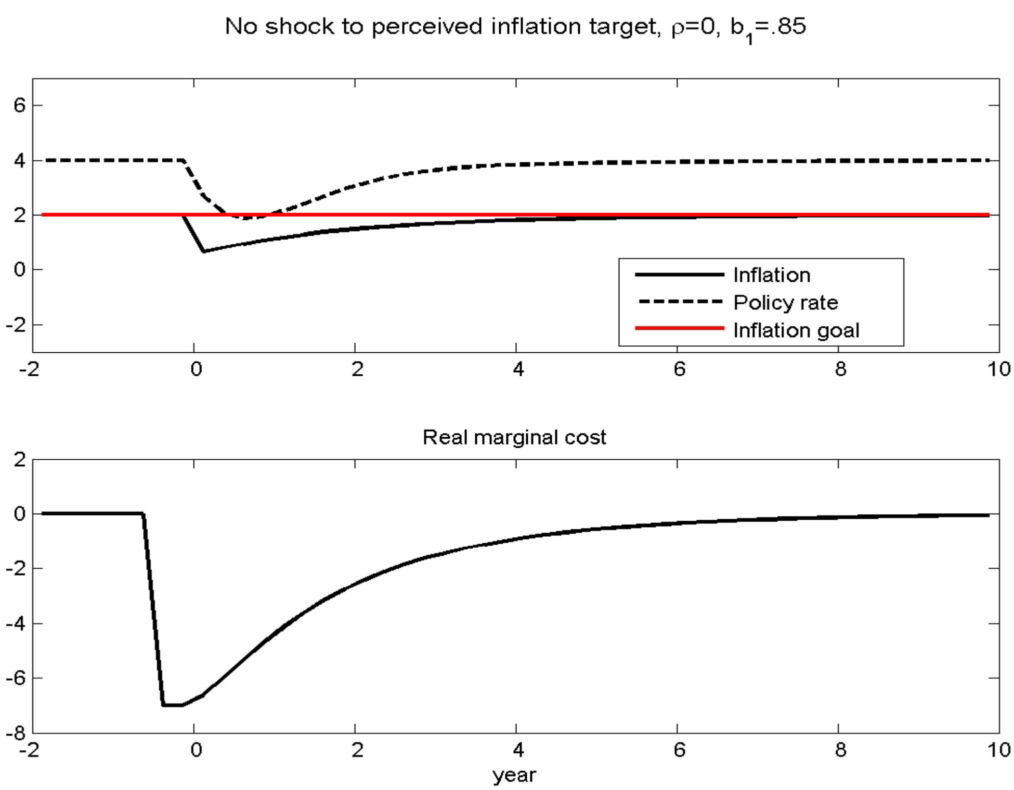

Figure 4
Inflation following post-war recessions

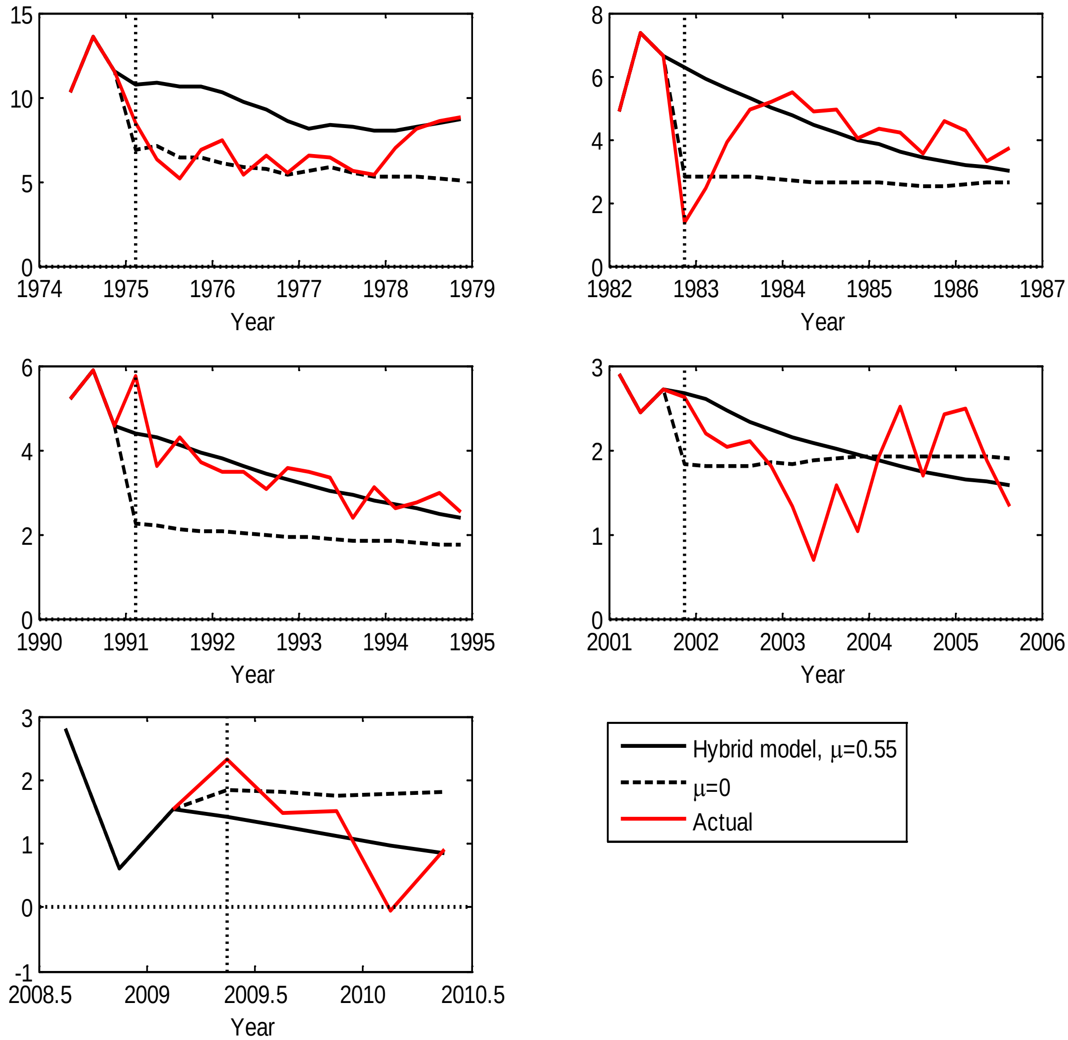

Figure 5

Consumer Price Inflation

% change, year-to-year, 1954:Q1 to 1964:Q4

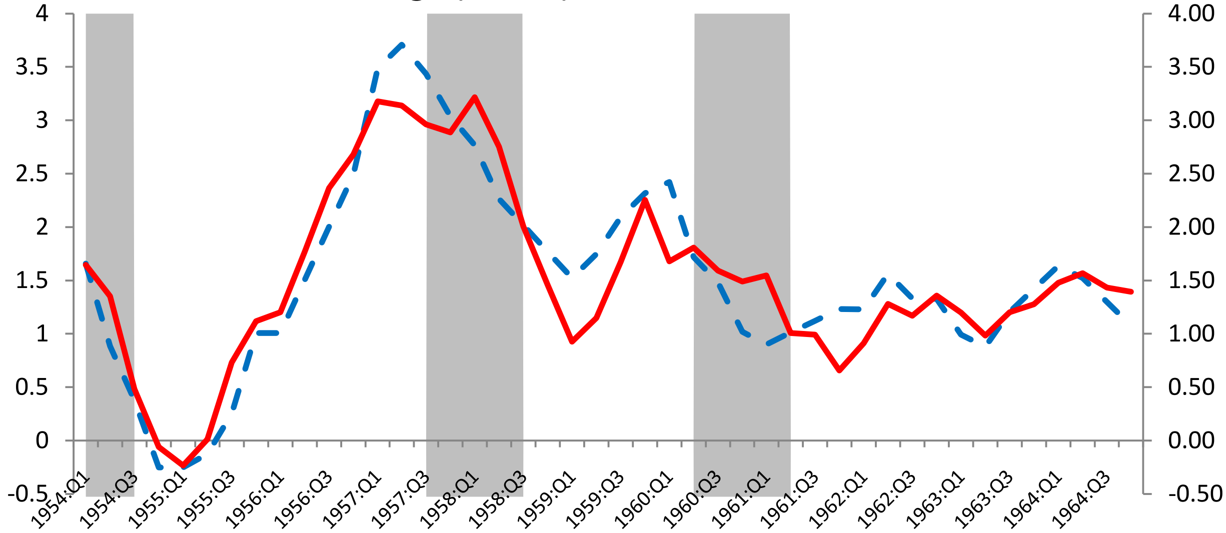

Figure 6

Civilian Unemployment Rate

1954:Q1 to 1964:Q4

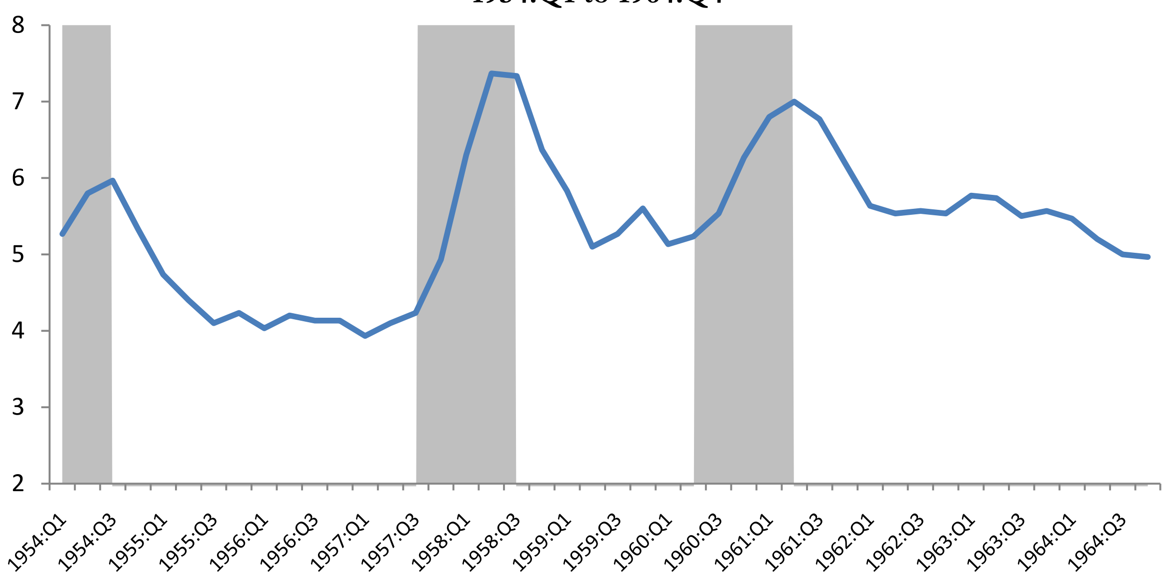

Figure 7

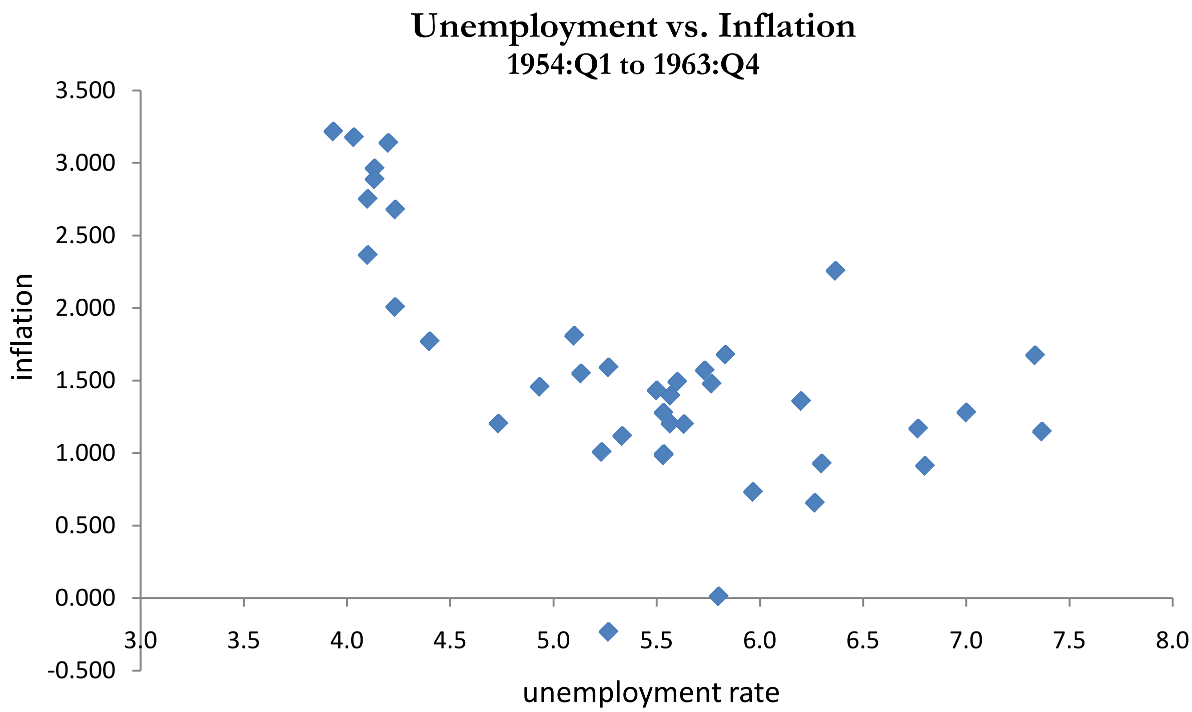

Figure 8

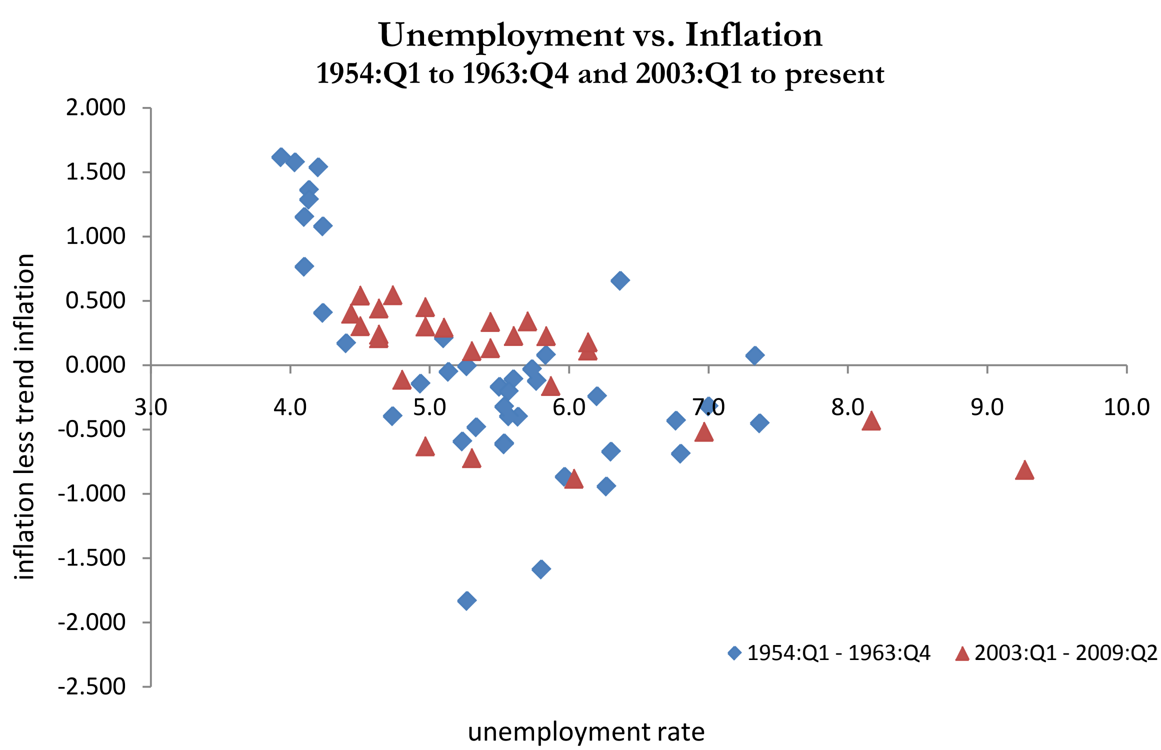

Figure 9

Figure 10

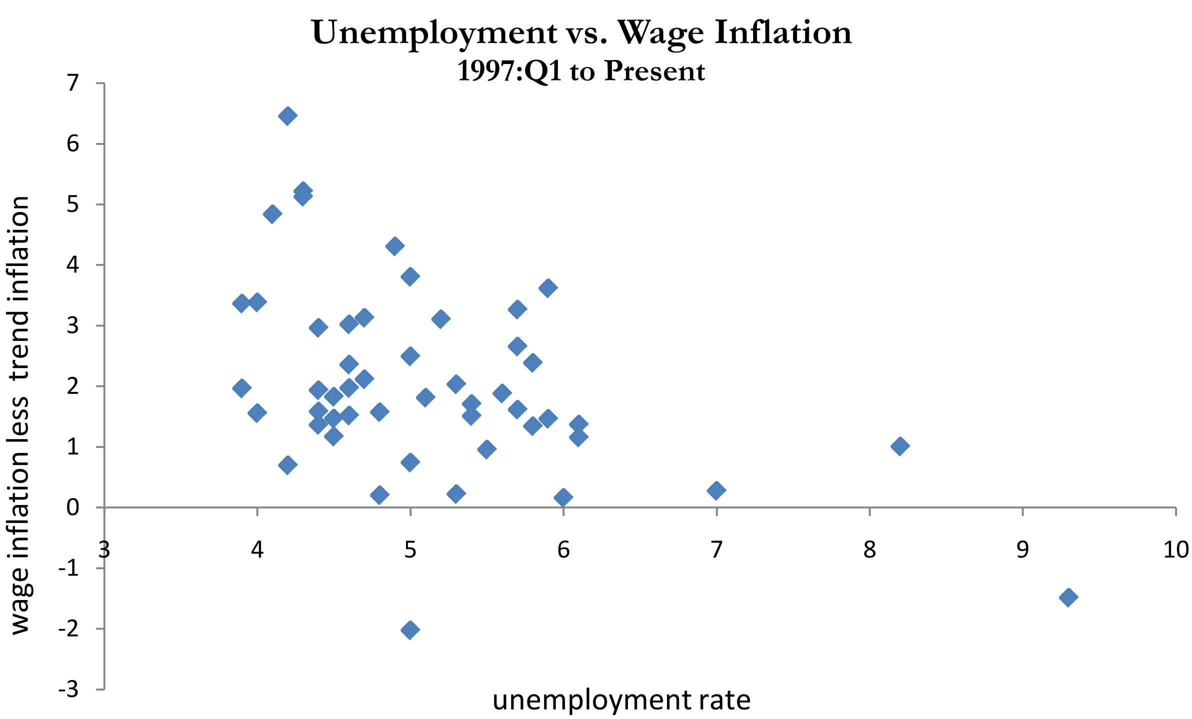

Figure 11

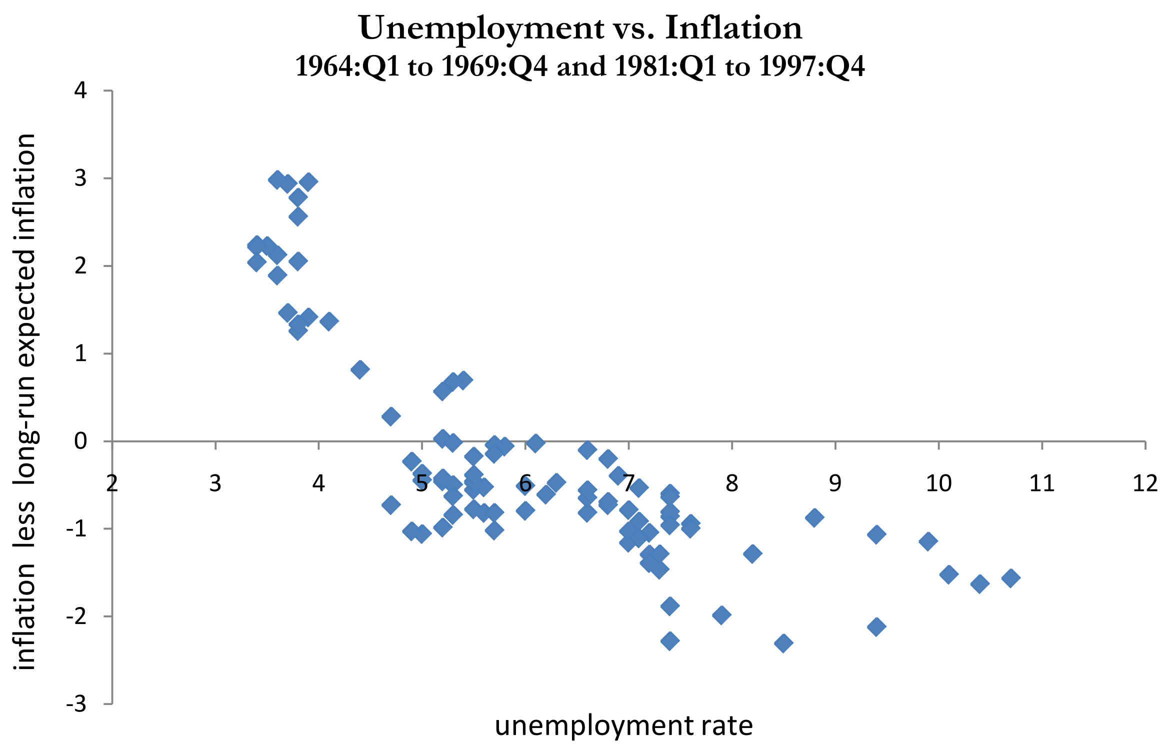

Figure 12

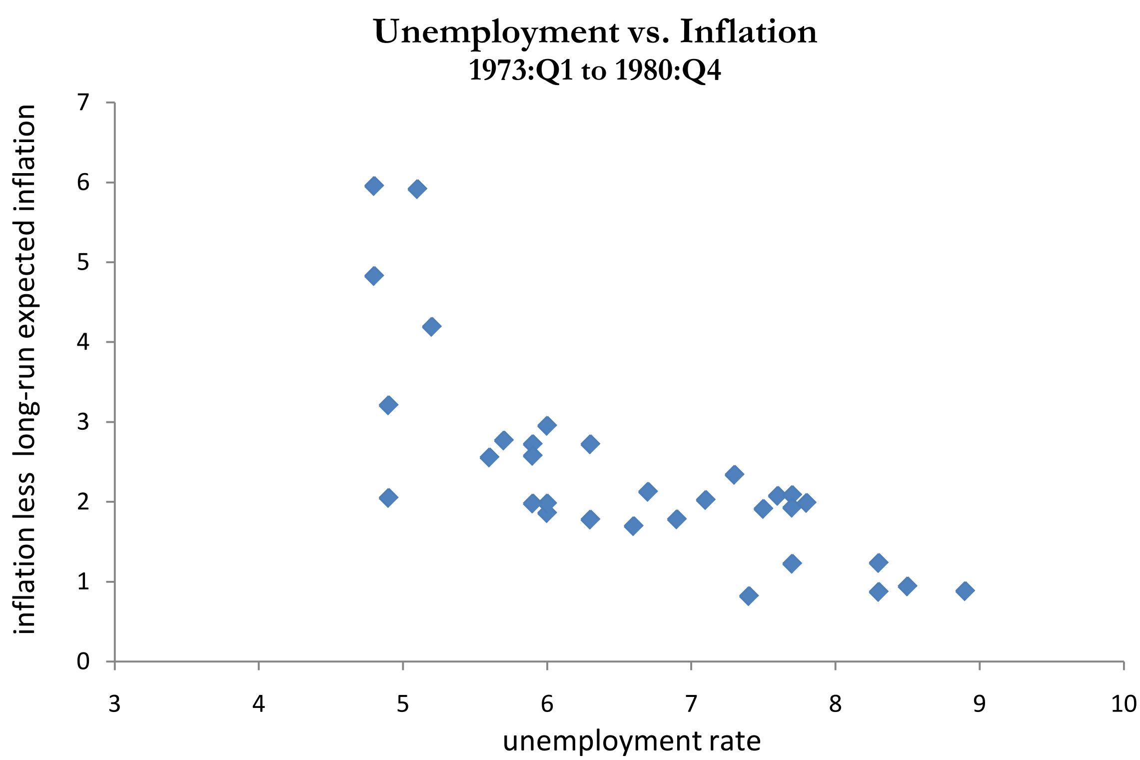

Figure 13

Key Japanese data

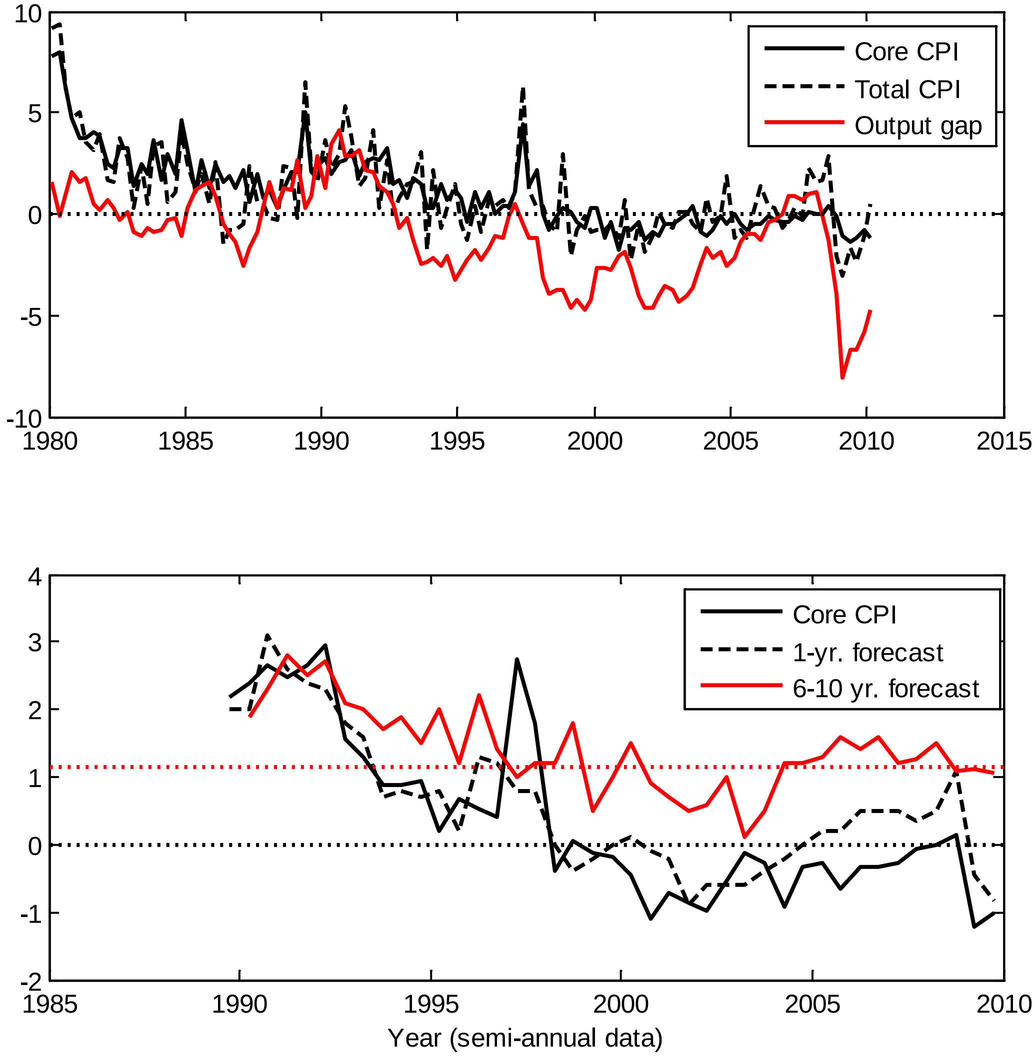

Figure 14
Fit of conventional Phillips curve, 1980-2010

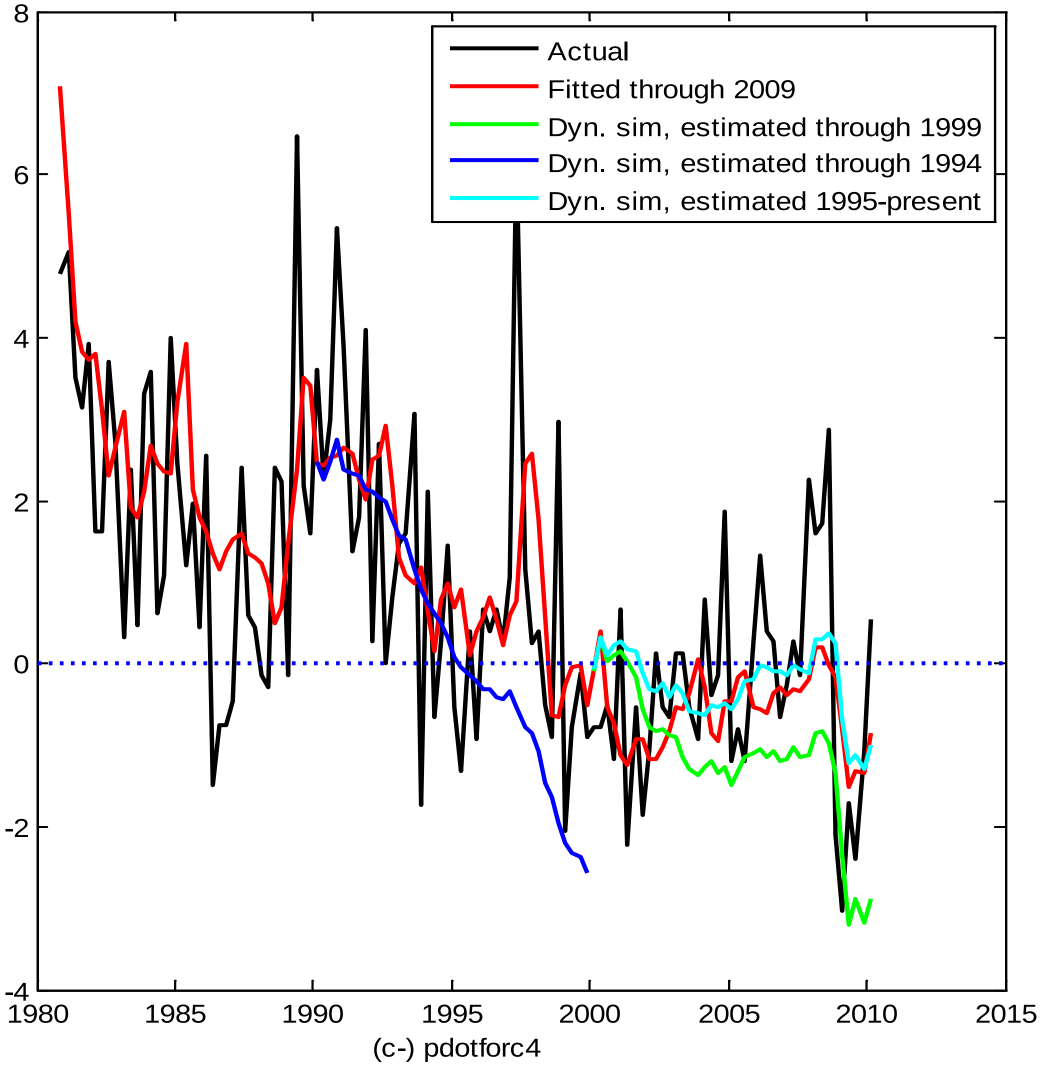

Figure 15

A Non-accelerationist Japanese Phillips curve?

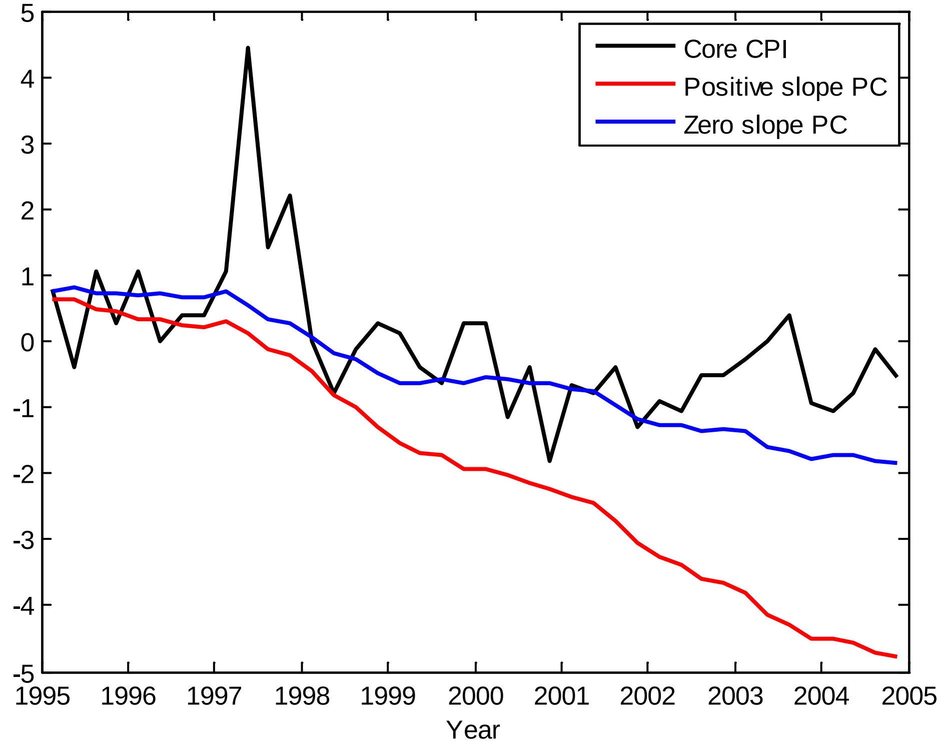

Figure 16

A long-run trade-off between output and inflation?

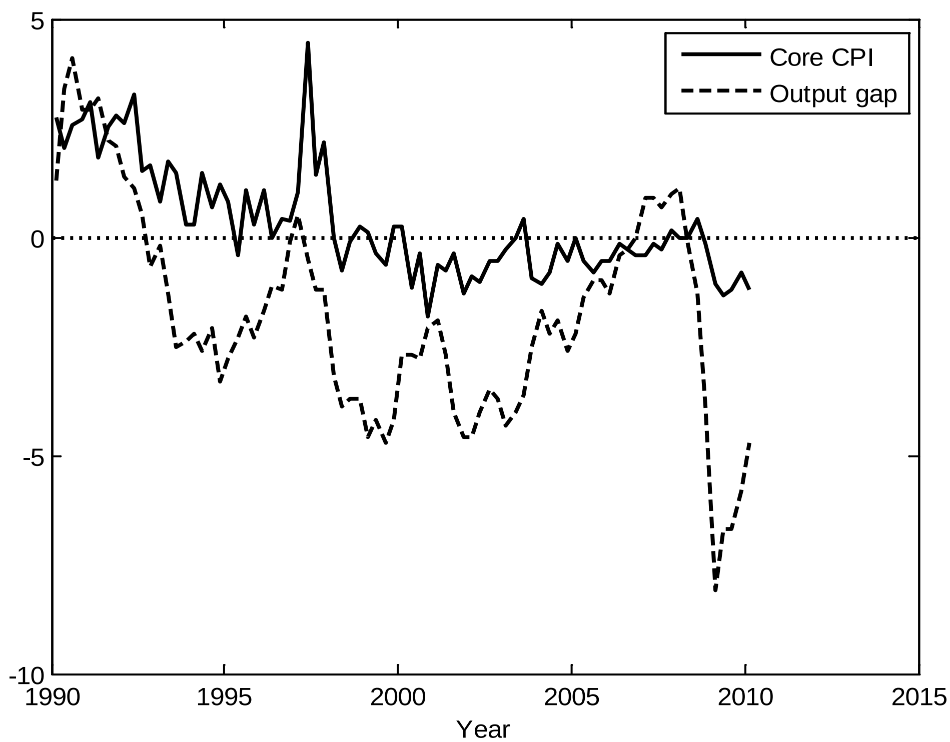

Figure 17
Japanese Phillips curve with survey expectations

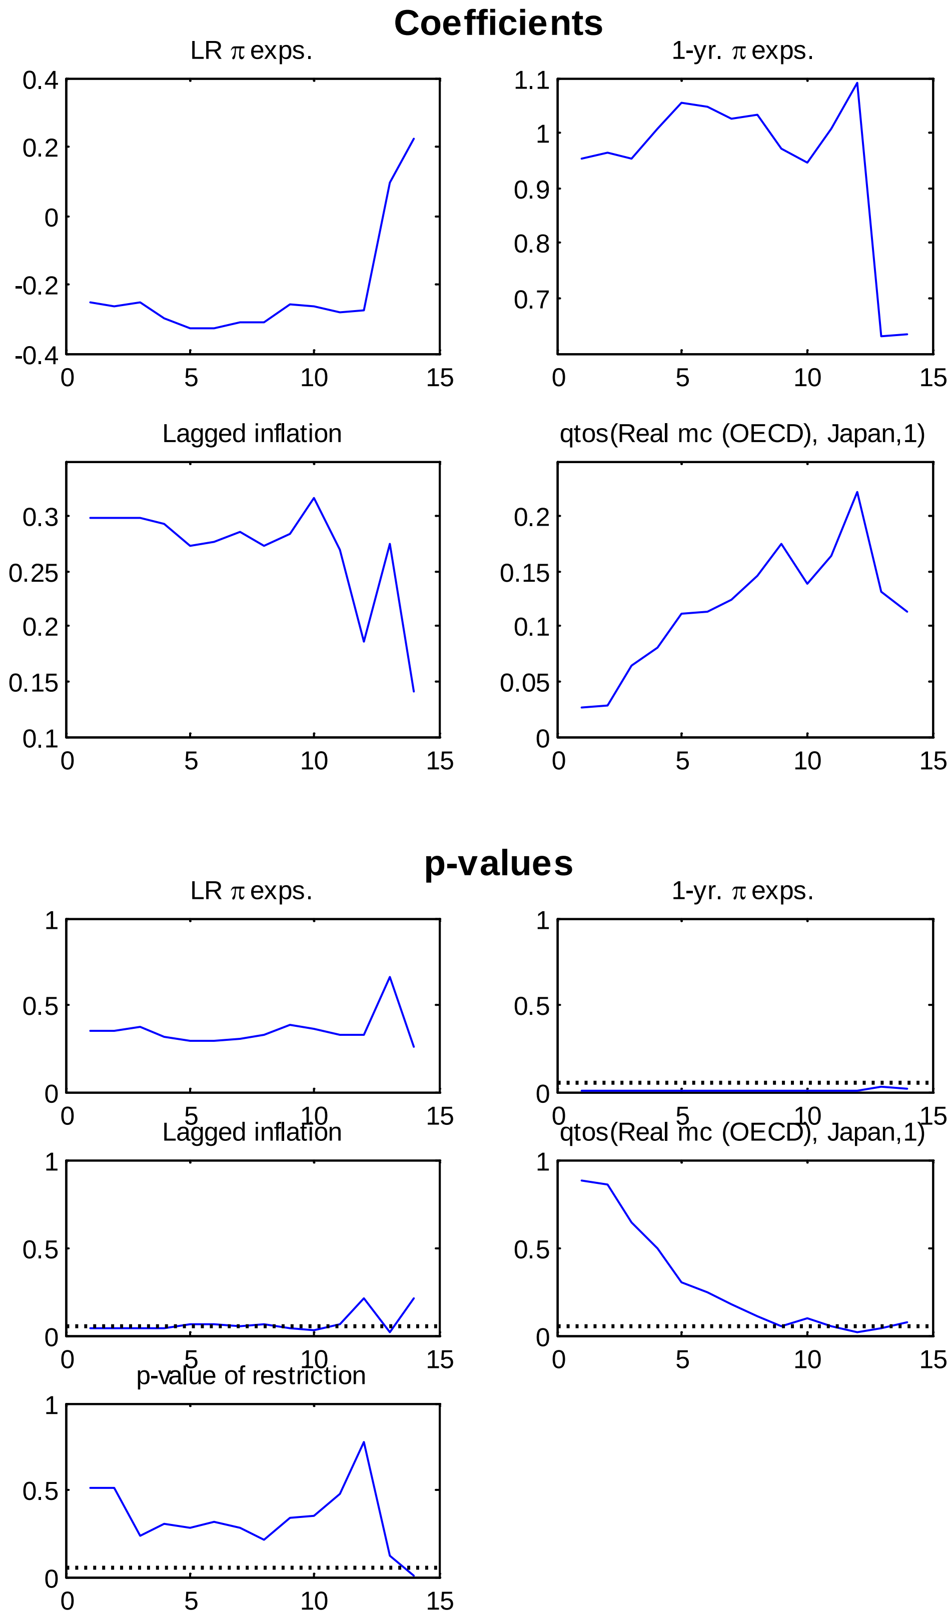

Figure 18
Japanese Phillips curve with one-year survey expectations

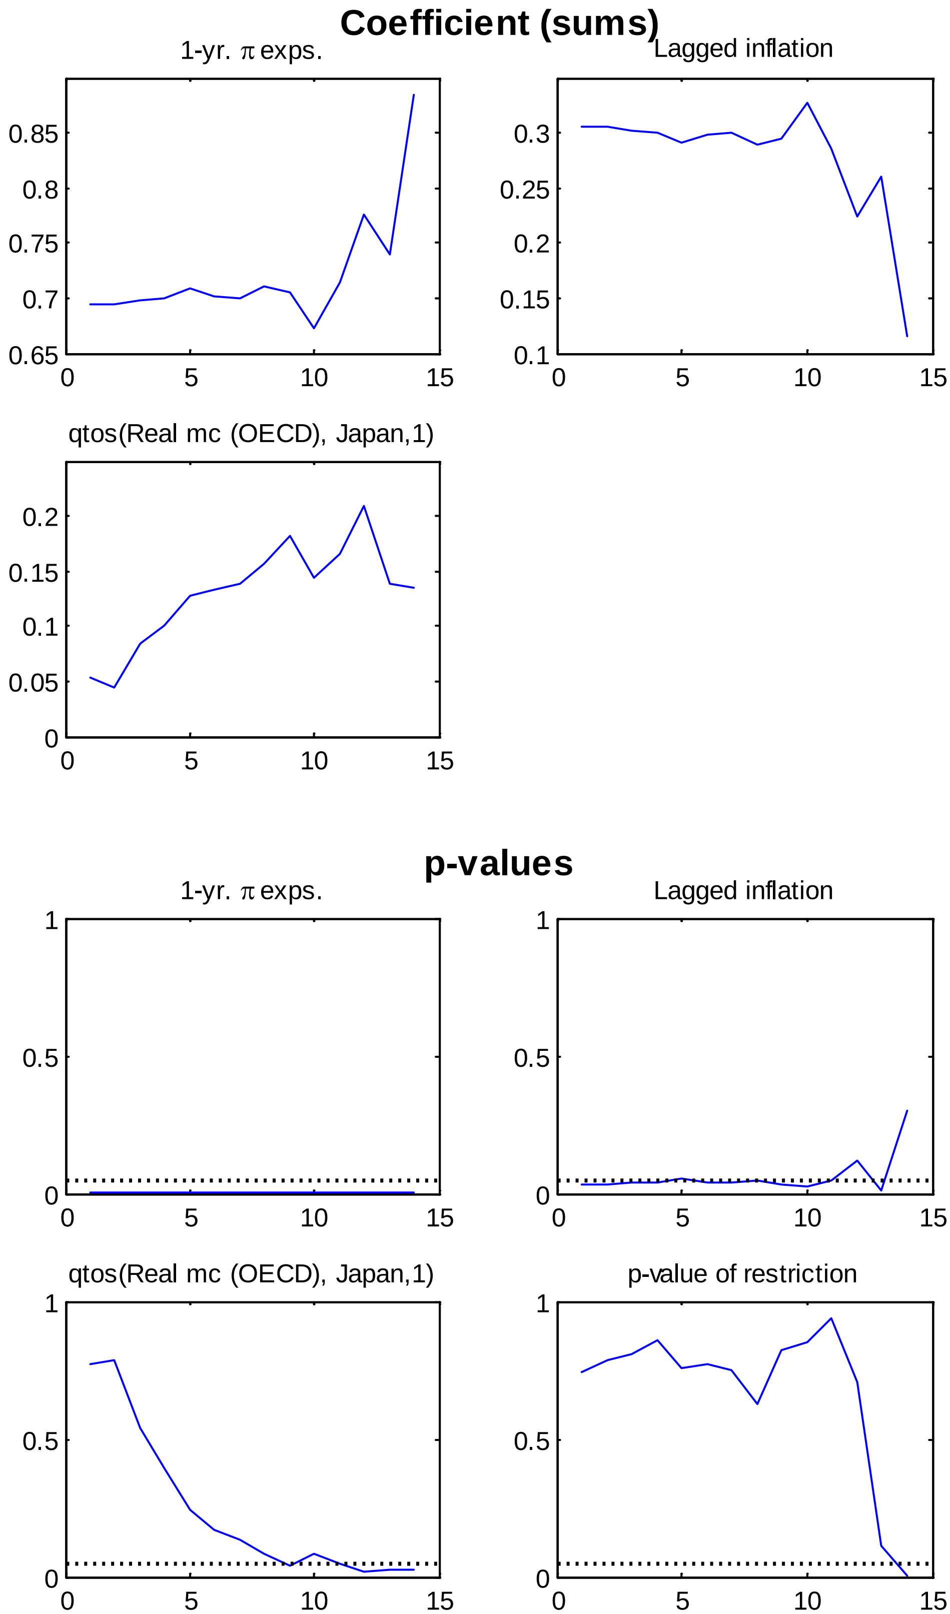

Figure 19
Japanese Phillips curve with one-year survey expectations

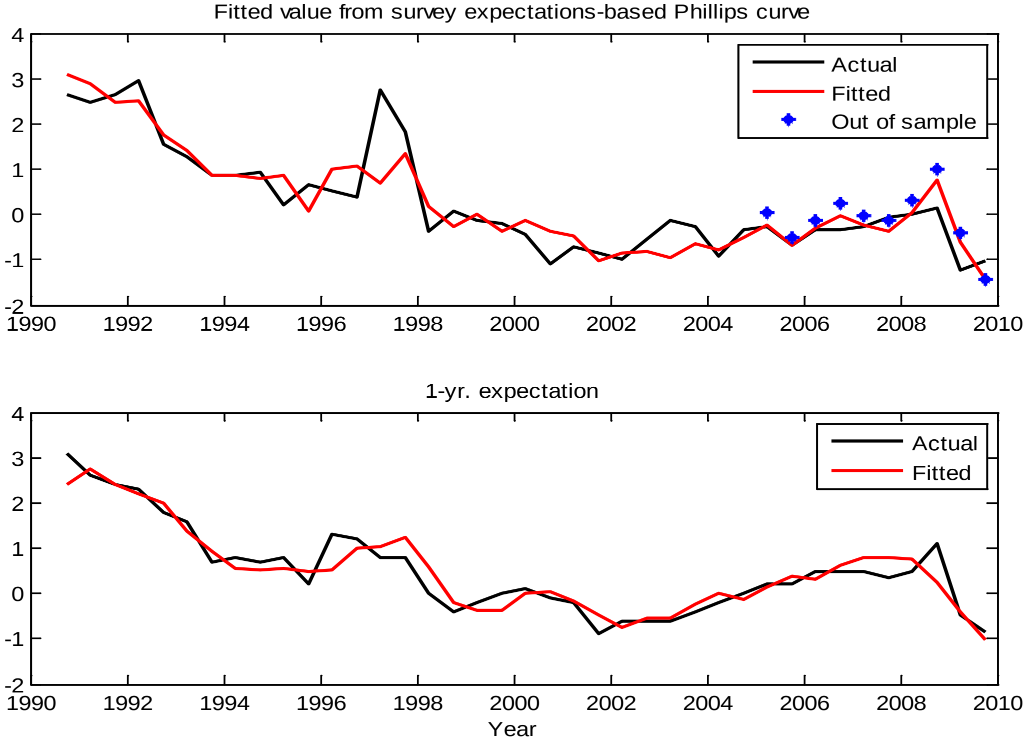

Figure 20

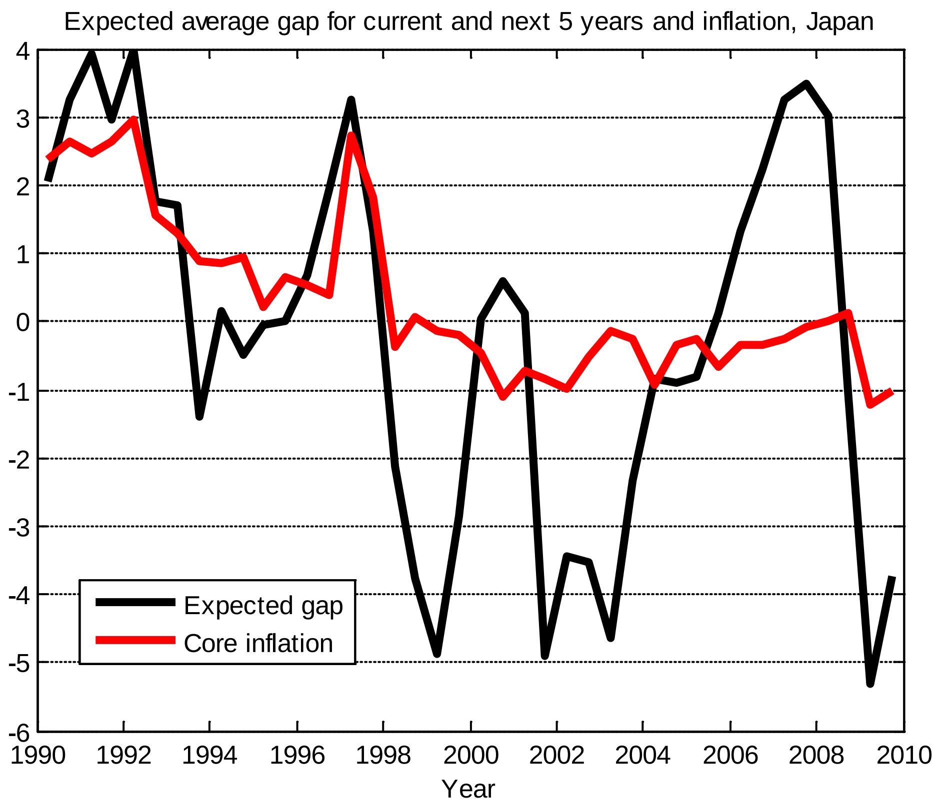

Figure 21

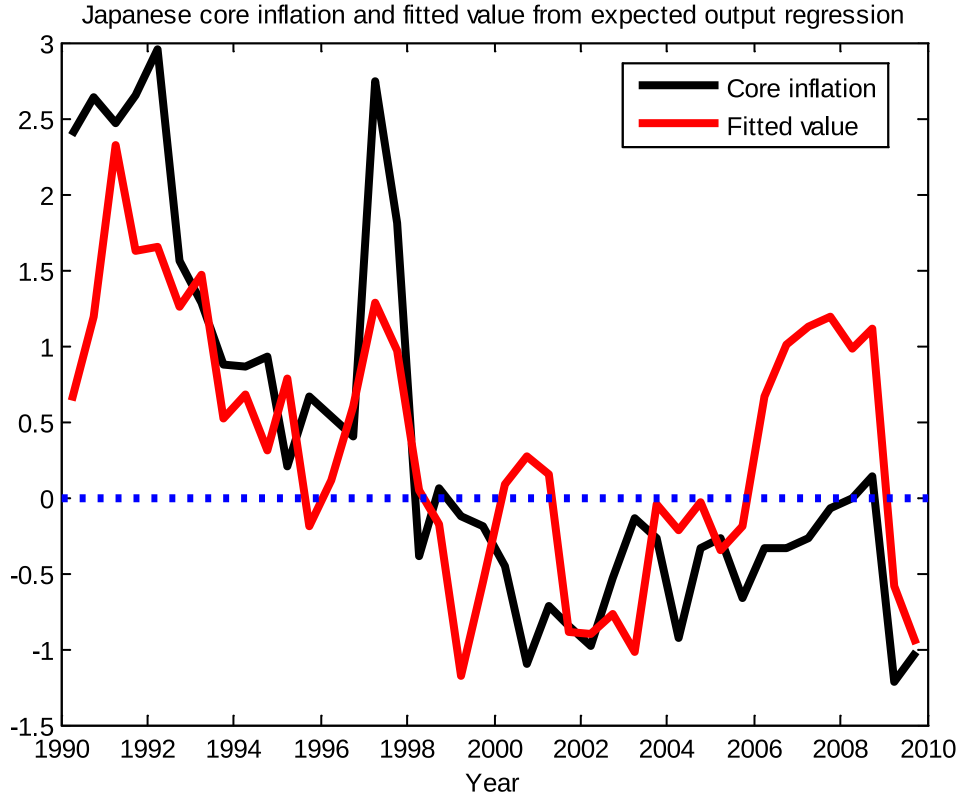

Figure 22

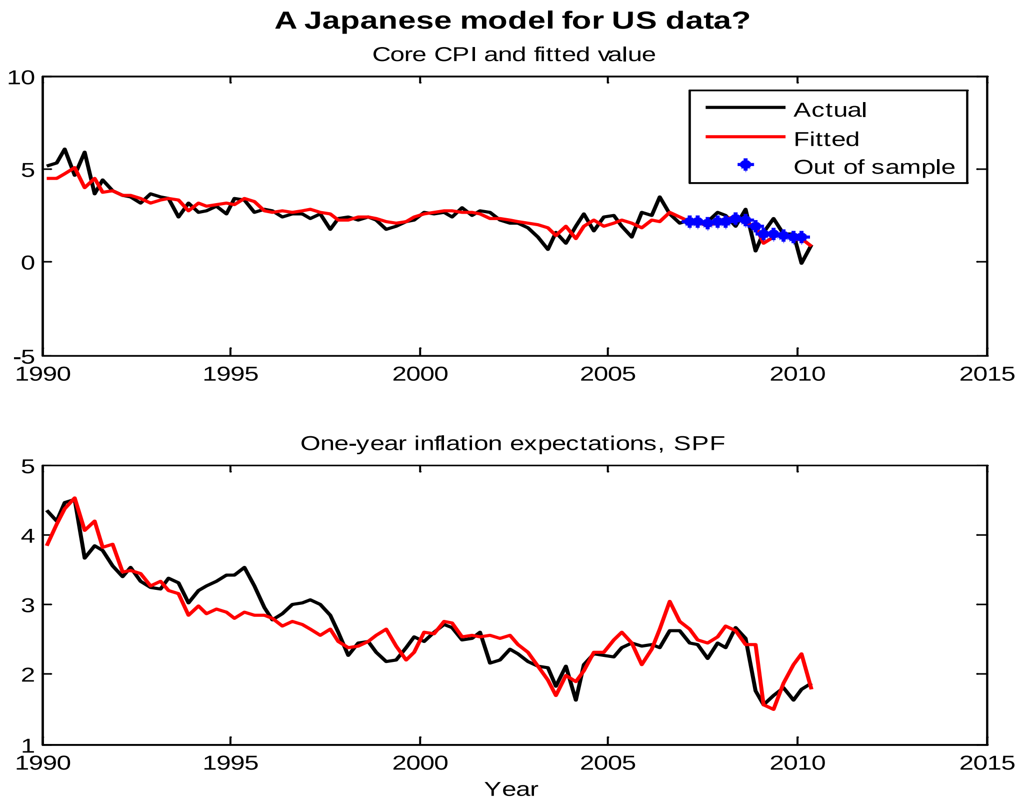

Figure 23

Simulation of model with 2010 initial conditions

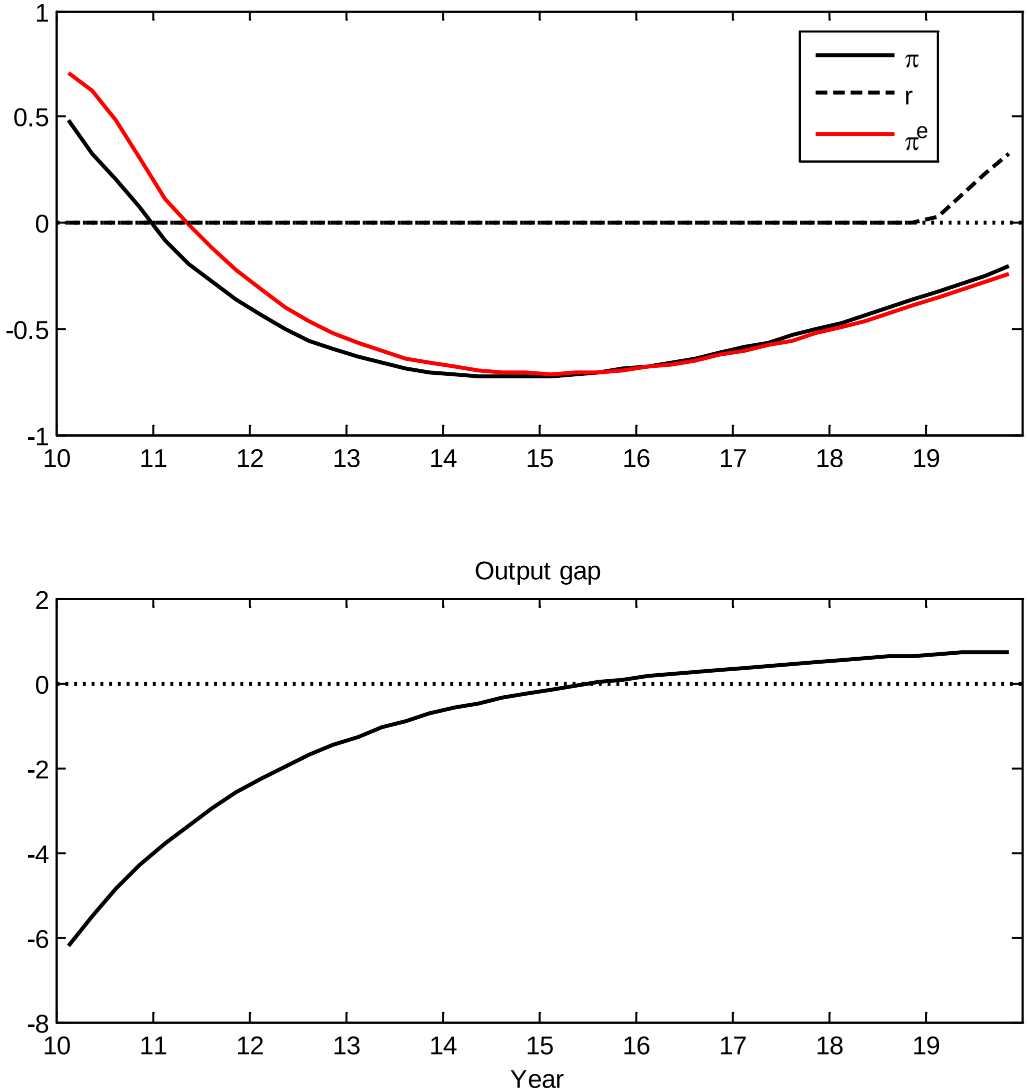

Figure 24

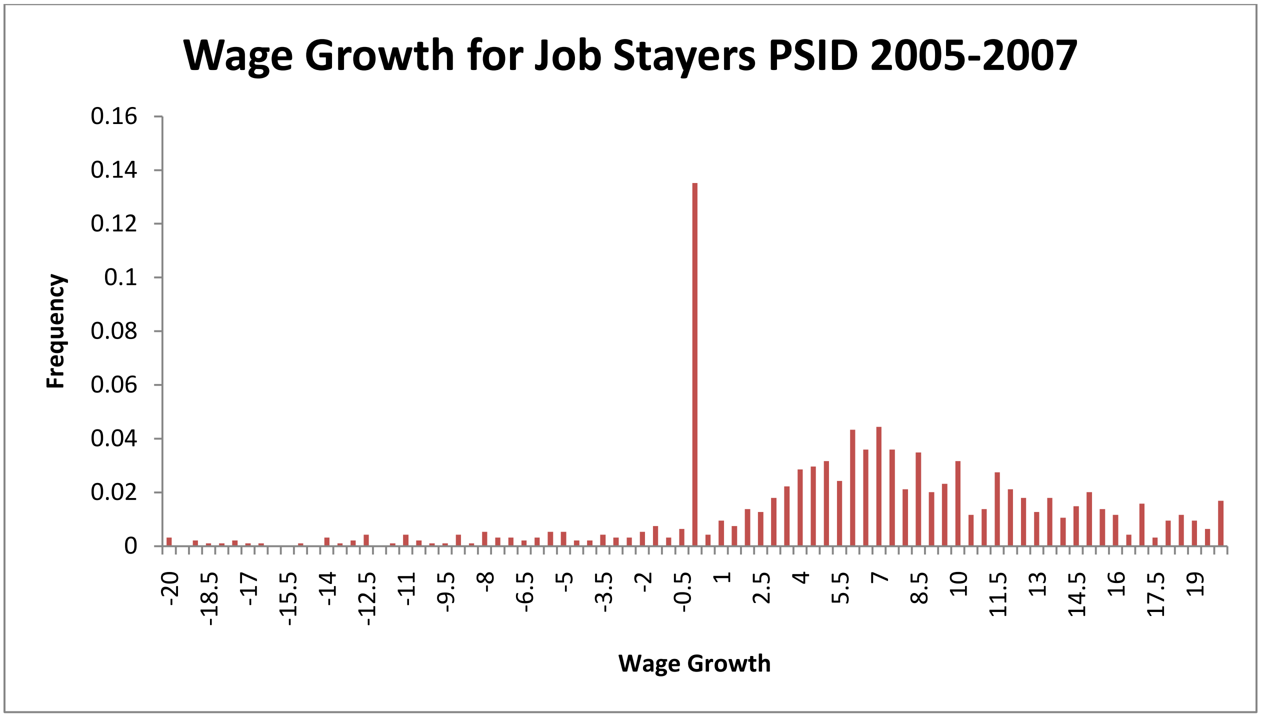

Figure 25

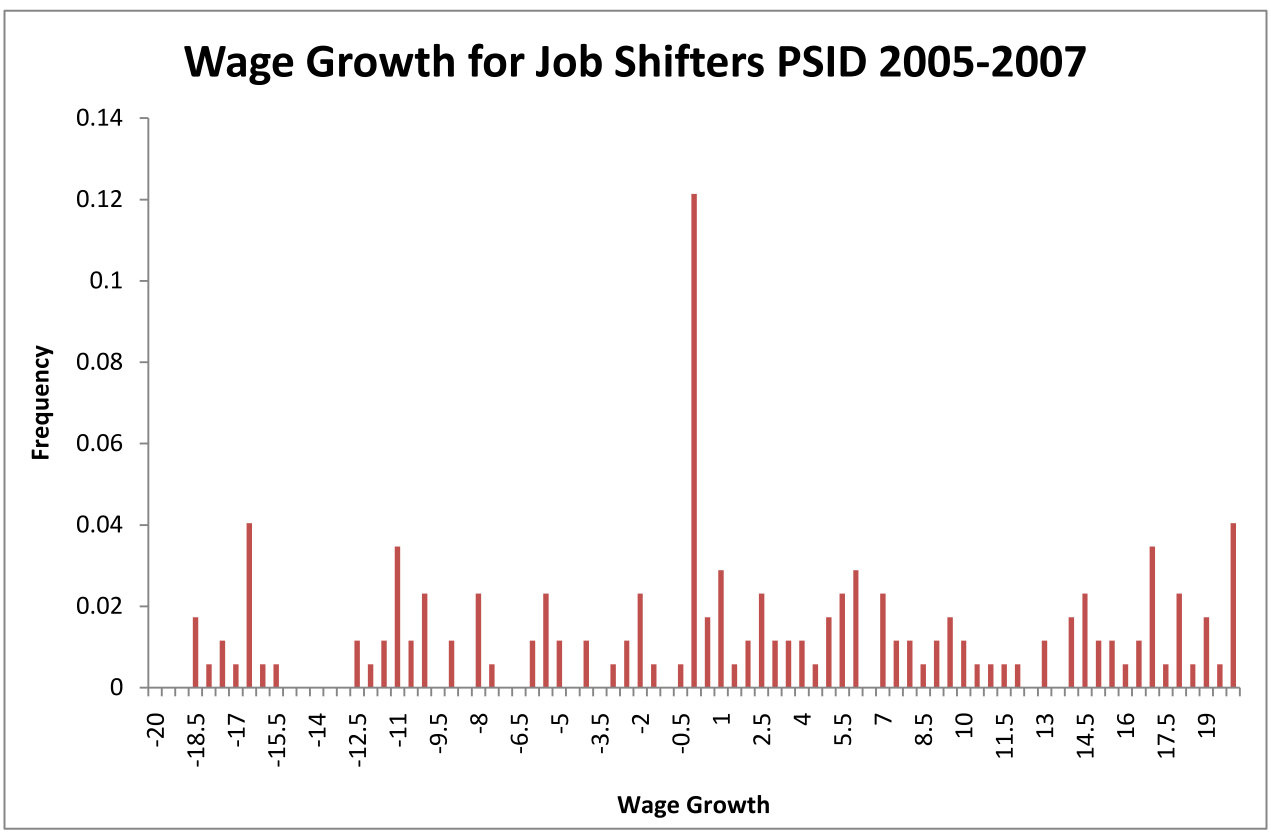

Figure 26

FREQUENCY OF WAGE CHANGES FOR JOBS IN THE OES

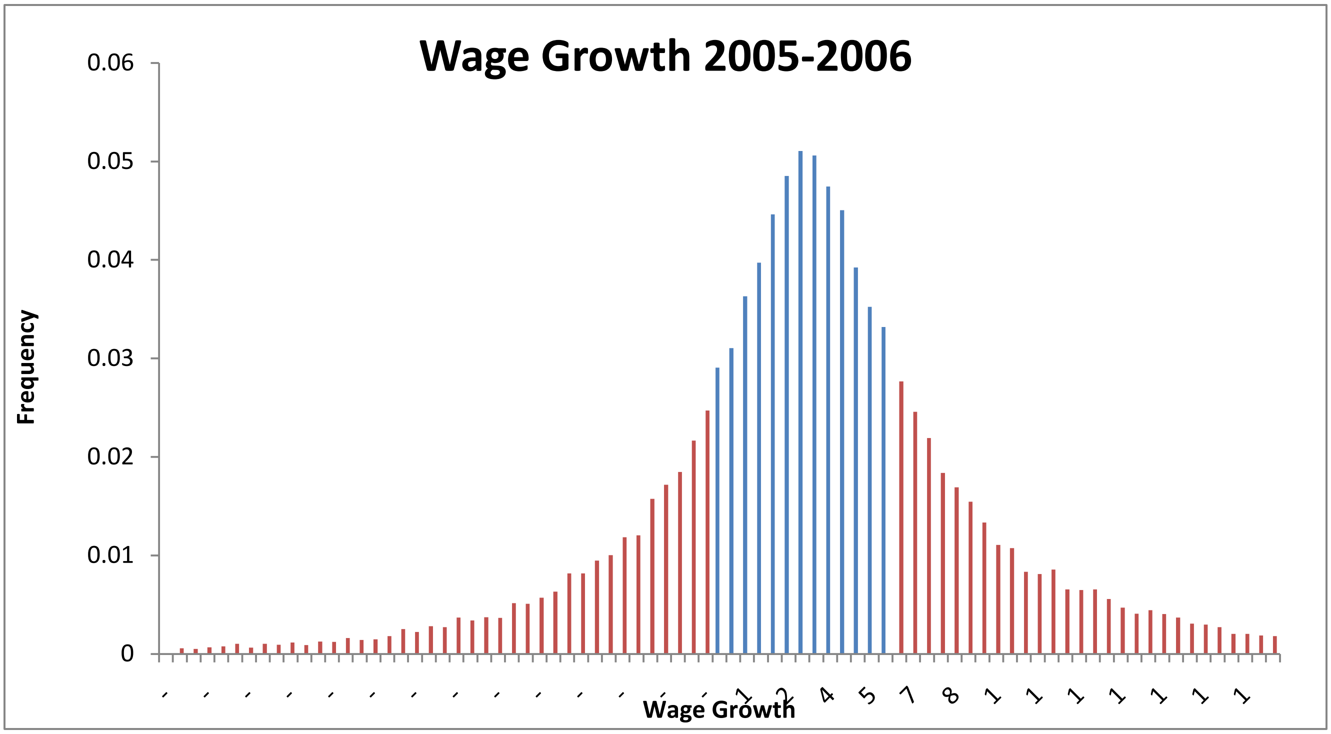

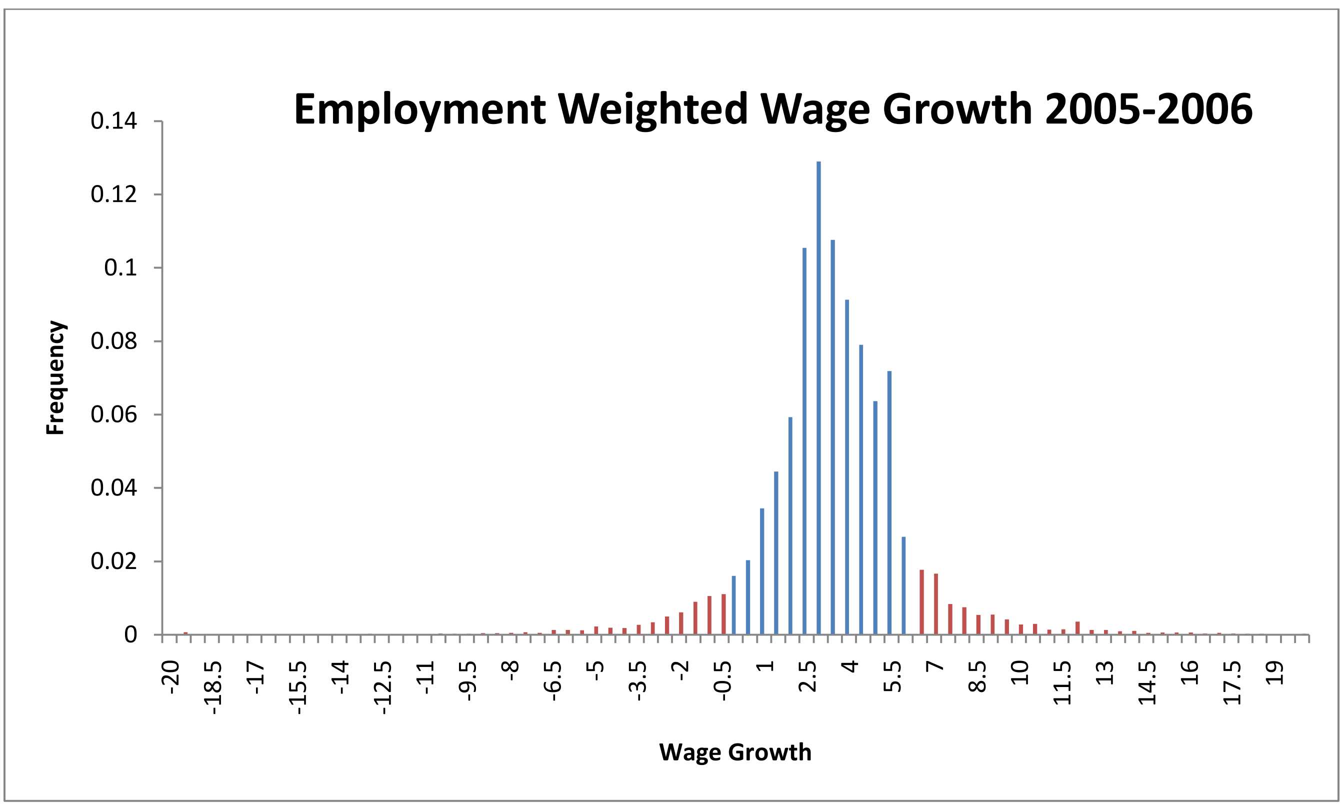

Figure 27

FREQUENCY OF WAGE CHANGES IN OES, 2008-09

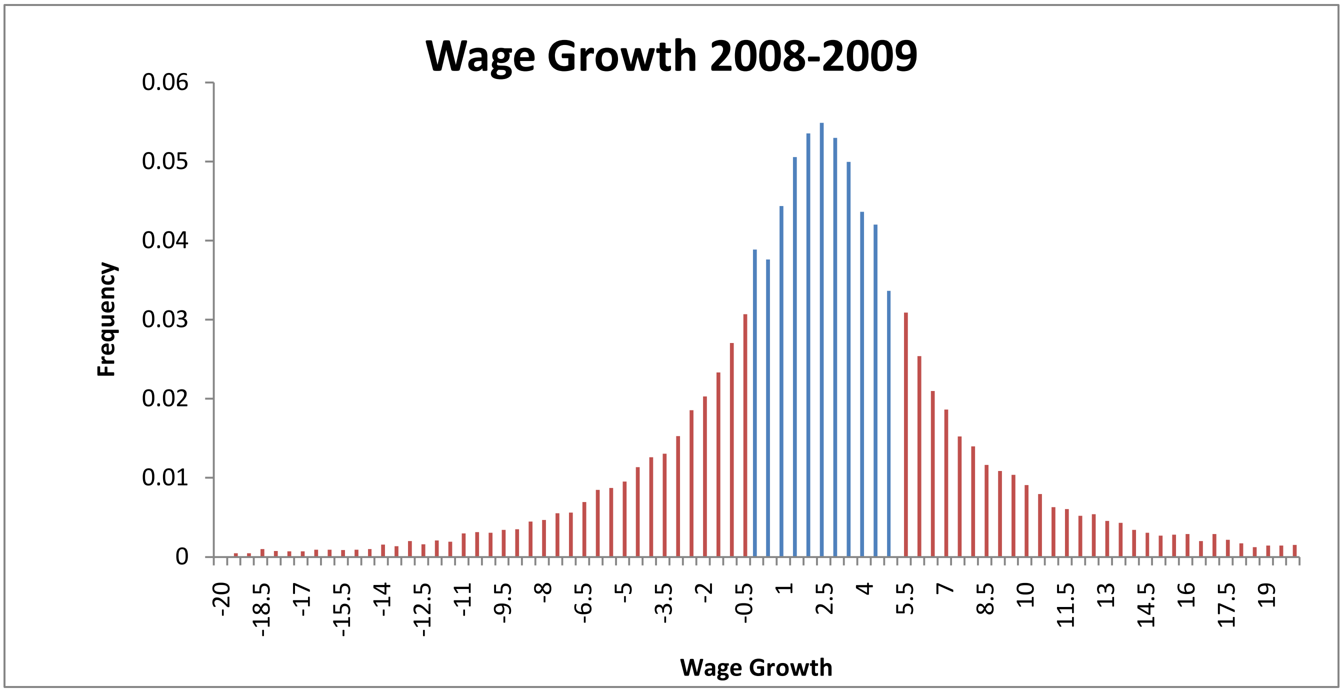

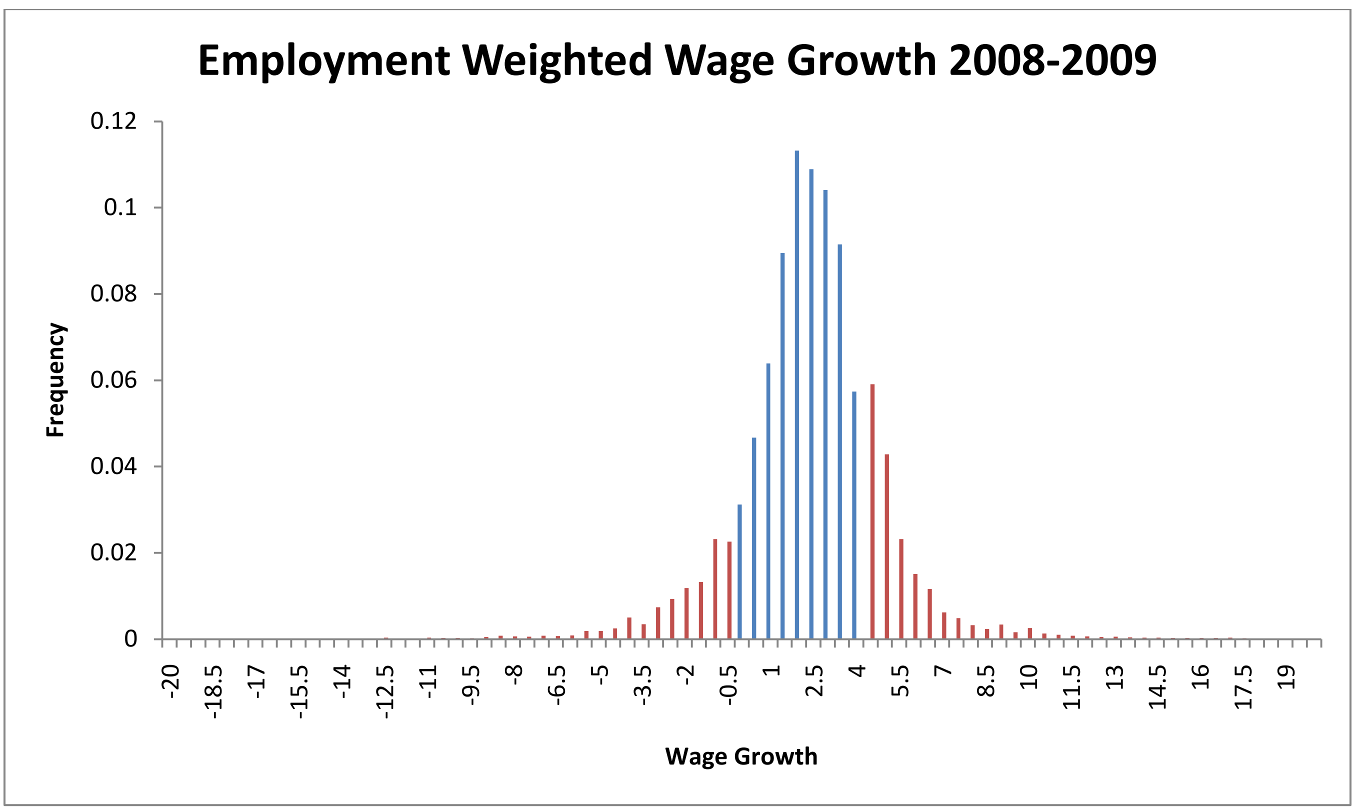

Figure 28

Range of Wages in Jobs

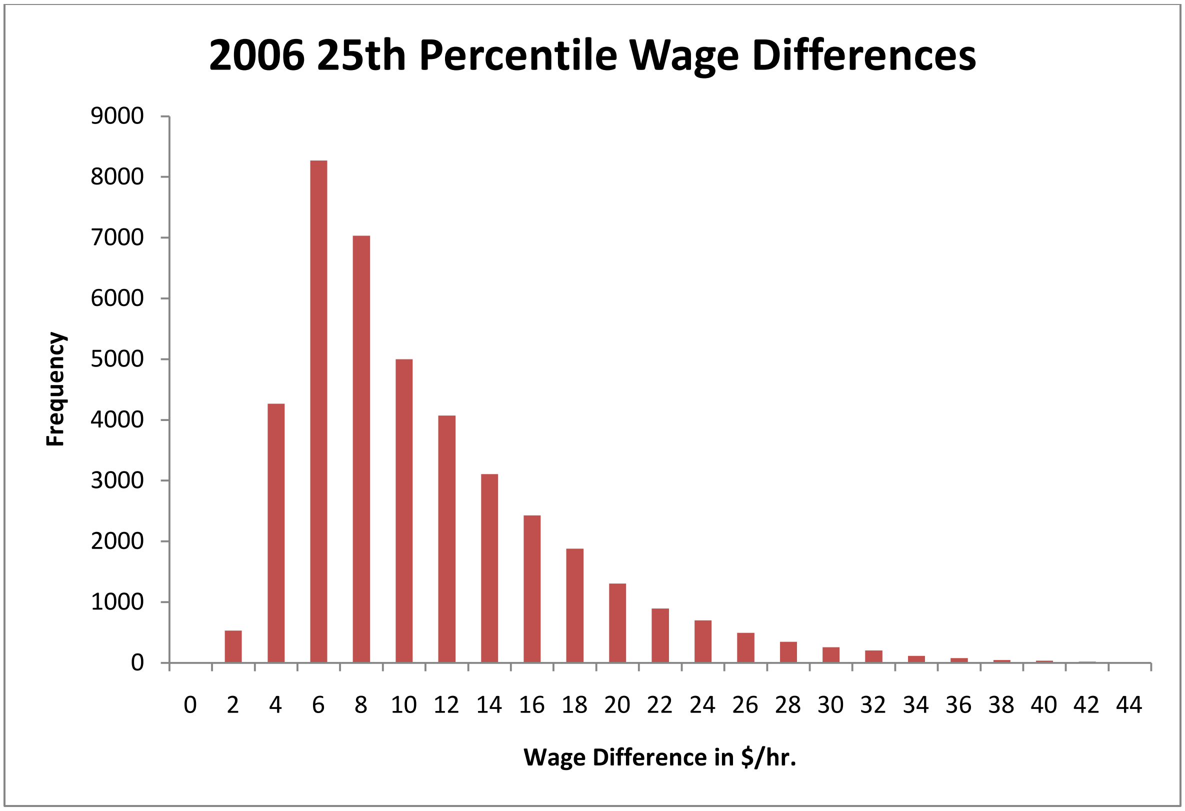

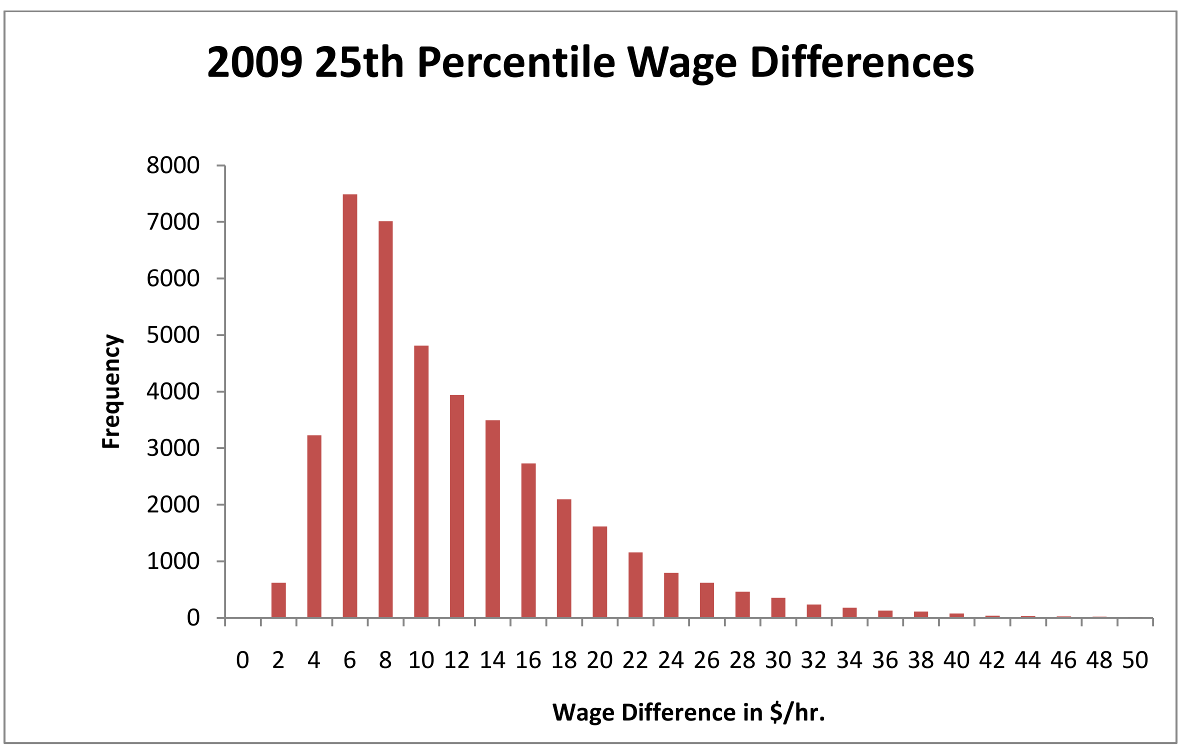

Figure 29

Distribution of negative wage changes

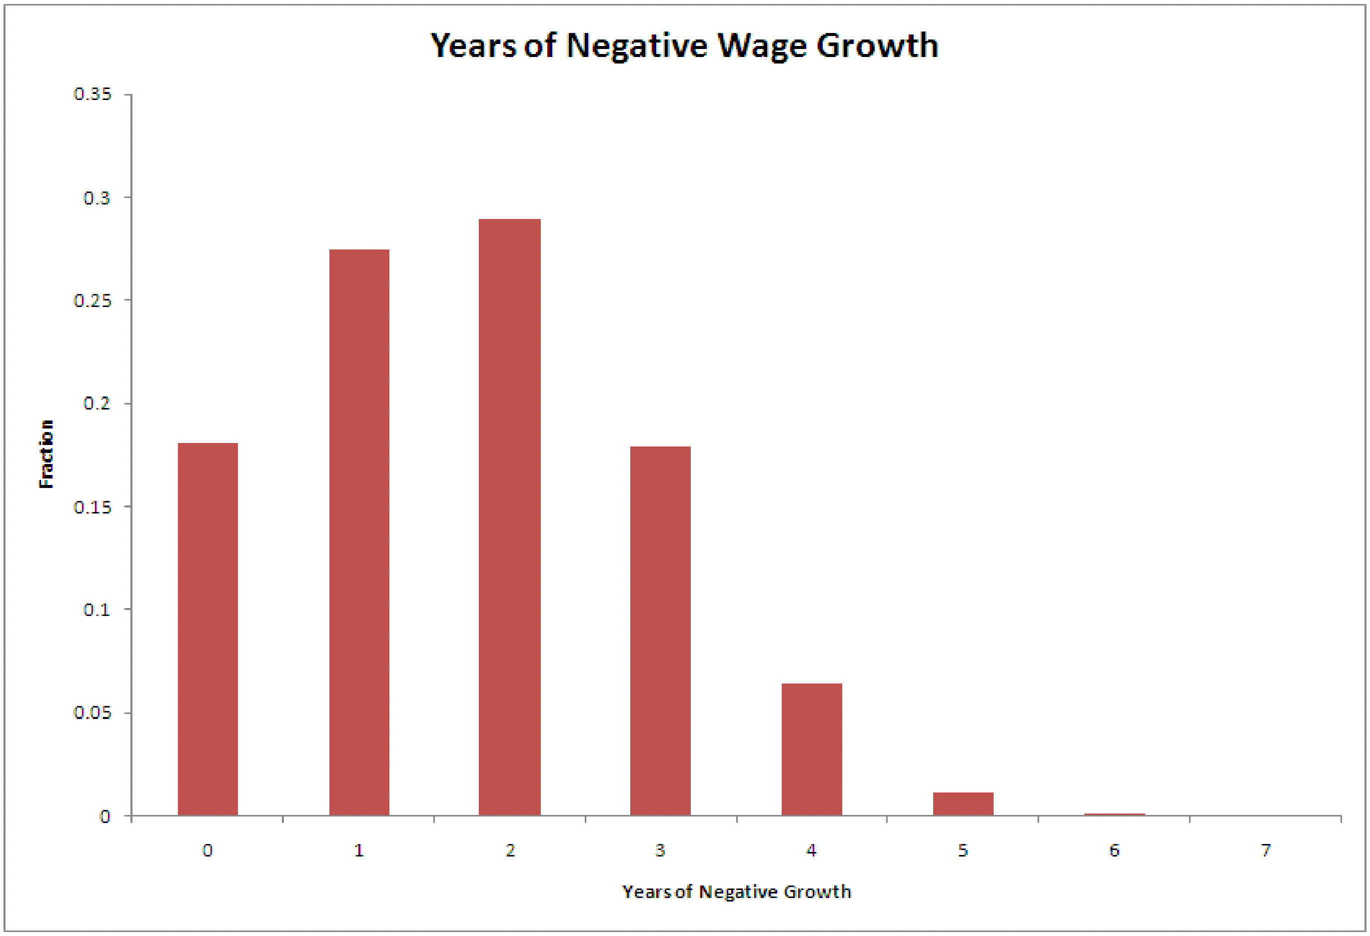

Figure 30

Wage Growth in Government Sector