

A Principal Components Approach to Estimating Labor Market Pressure and Its Implications for Inflation

Michelle Barnes, Ryan Chahrour, Giovanni Olivei, and Gaoyan Tang

We build a summary measure of labor market pressure that captures the common movement among a variety of labor market series. Obtained as the labor market series' first principal component, this measure explains a large portion of the variability of the underlying series. For this reason, it is a good summary indicator of labor market pressure. We show that the unemployment rate gap has tracked this summary measure closely over the past 35 years. At times, however, the summary measure and the unemployment rate gap have sent somewhat different signals. In terms of relying on the principal components summary measure vis-à-vis the unemployment rate gap for explaining inflation, we argue that the recent evolution of wage inflation is more consistent with the evolution of the summary measure than with the unemployment rate gap. This is because over the past two years the principal components summary measure has been suggesting less labor market pressure than the unemployment rate gap. Over the past 35 years, however, there is little systematic evidence favoring the summary measure of labor market pressure over the unemployment rate gap as a predictor of inflation.

This version: December 21, 2007

Michelle Barnes is Senior Economist and Policy Advisor at the Federal Reserve Bank of Boston. Her email address is michelle.barnes@bos.frb.org. Ryan Chahrour is a second-year student in economics at Columbia University. His email address is rc2374@columbia.edu. Giovanni Olivei is Senior Economist and Policy Advisor at the Federal Reserve Bank of Boston. His email address is giovanni.olivei@bos.frb.org. Gaoyan Tang is a Senior Research Assistant at the Federal Reserve Bank of Boston. Her email address is gaoyan.tang@bos.frb.org.

The authors thank Kathy Bradbury and Christopher Foote for many helpful suggestions on this brief.

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

The views expressed in this brief do not necessarily reflect the official position of the Federal Reserve Bank of Boston or the Federal Reserve System.

Introduction

The unemployment rate is often taken as the summary statistic for the current state of the labor market. At times, however, other measures of labor market activity send signals about labor market conditions that differ from the unemployment rate. For example, in the second half of 2003 gains in payroll employment were relatively muted (75,000 net new jobs per month on average), while the unemployment rate declined by 0.6 percentage points, a fairly significant improvement. Policymakers are known to follow many series, and are keenly aware of the conflicting signals that different series sometime provide. This raises the practical issues of how to usefully summarize the information from a large number of series and, in the present context, of determining whether the unemployment rate gap can indeed be taken as a sufficient summary statistic of labor market pressure.

In this brief we show that a single factor which captures the common movement among several different labor market series—the first principal component of the series—explains a large portion of the variation among the several series. As such, the first principal component is a good summary indicator of the labor market series themselves. We find that over time the unemployment rate (expressed as the deviation from its long-run trend) has tracked this summary indicator closely. Thus, while the concept of “labor market pressure” may not be perfectly represented by a single observable series, the unemployment rate usefully summarizes a large portion of the information contained in several labor market variables that indicate tight or loose employment conditions.

While the co-movement between this first principal component factor and the unemployment rate gap has been high over time, there have been occasional divergences between the two series. Under these circumstances, it is interesting to ask which of the two approaches to estimated labor market pressures provides a better signal of inflationary pressures. We show that recent wage inflation behavior appears more consistent with the evolution of the principal component factor than with the evolution of the unemployment rate. However, there is little statistical evidence to favor

one series over the other. During the past 35 years, reliance on the unemployment rate gap rather than on the labor market pressure factor does not produce consistently worse outcomes for predicting inflation.

The remaining portion of the brief is structured as follows. In Section I we estimate the summary measure of labor market pressure, and compare its evolution to the evolution of the unemployment rate gap. In Section II we discuss the robustness of our findings by reconstructing the summary measure using different information sets. Section III focuses on the recent evolution of the summary measure of the gap in labor activity vis-à-vis the evolution of the unemployment rate gap. Section IV discusses some implications for predicting inflation within a standard Phillips curve framework when using the summary measure of labor market pressure instead of the unemployment rate gap. Section V concludes.

I. A Summary Statistic of Labor Market Pressure and its Relationship to the Unemployment Rate Gap

In this section we consider the information on labor market activity contained in a relatively parsimonious set of labor market variables to construct a summary statistic of labor market pressure based on principal components analysis. The twelve chosen series incorporate information about the labor market drawn from different sources. (These variables and the principal components summary indicator constructed from them are listed first in the Appendix, and are denoted “PC” there and in the figures.) We consider information from the two labor market surveys produced by the Bureau of Labor Statistics, the Current Population Survey (from which the unemployment rate is constructed), and the Current Employment Statistics Survey (which, among other things, records business establishments’ payrolls and hours of work). The series we take from these two labor market surveys are closely followed by policymakers; it is significant that one of these is the unemployment rate. We also consider the job-finding

rate, an indicator of the ease with which people are finding jobs.¹ The flip side of unemployment is job vacancies, and to proxy for vacancies we use the Conference Board's help-wanted advertising index. The other variables are initial claims for unemployment and consumers' appraisal of the state of the labor market.² In all, these twelve series provide a wide array of information about labor market activity, including some information more specific either to the supply or to the demand side of the market.

Virtually all of the labor market series we consider display low-frequency trends, due to demographic and other factors. For example, payroll employment will depend on the growth of the labor force, which, in turn, depends on demographic trends. The unemployment rate will exhibit long-term fluctuations that are, among many other things, the result of changes in the age composition of the labor force. The allocation of time between work and leisure has changed over the last few decades, with attendant consequences for the number of market work hours reported per week.³ To highlight business-cycle fluctuations, each series is expressed as the deviation from an estimated low-frequency trend which we estimate by means of a Hodrick-Prescott filter (Hodrick and Prescott 1997).⁴

The Hodrick-Prescott filter is a symmetric two-sided filter, and it optimally decomposes a time series into a trend component and a cyclical component. The filter will be efficient at the "center" of a "long" series, but the estimated decomposition can become imprecise near the endpoints of the series.⁵ Because of the filter's inefficiency near endpoints, for the purposes of the discussion in this section we consider the behavior of the unemployment rate gap and of a summary indicator of labor market pressure up to 2002:Q4.⁶ Of course, from a policy standpoint there is considerable

¹ See Shimer (2005). The job-finding rate is constructed with information from the Current Population Survey.

² This last series is taken from the Conference Board's Consumer Attitudes survey.

³ See Aguiar and Hurst (2007).

⁴ The filter has smoothing parameter 10^5 for quarterly data.

⁵ See Mise, Kim, and Newbold (2005).

⁶ Since the endpoint for each of the series considered (and for the estimated Hodrick-Prescott trend) is 2007:Q2, we are excluding from the analysis the most recent 18 quarters. We adopt a similar trimming at the beginning the sample.

interest in analyzing the most recent data, notwithstanding the uncertainty concerning the estimation of more recent trends. We devote Section III to discussing the current situation.

When considering a large number of different series, it is necessary to summarize the information contained in the collective series in a useful manner. We follow the recent literature⁷ and estimate the common movement in the chosen labor market series using a principal components approach. From the original variables, this method generates a new set of variables—called principal components—with each principal component being a linear combination of the original variables. The main feature of the principal components is that these are orthogonal to each other, which essentially means that each principal component contains no redundant information and measures a different driving force in the original data. There are as many principal components as original variables, and each component summarizes a decreasing portion of the total variance of the original data. This method usefully summarizes the information contained in the original variables when a large fraction of the original variables' total variance can be captured with the first few principal components.

Figure 1 compares the unemployment rate gap and the first principal component extracted from the twelve labor market series we consider over the period from 1973:Q1 to 2002:Q4.⁸ The first principal component summarizes the original series quite well, in that by itself the component explains 75 percent of the variability in the original data. The figure shows that the unemployment rate gap closely tracks this principal component, meaning that historically the unemployment rate gap has provided a good summary of overall labor market conditions. Yet there have been occasional divergences between the unemployment rate gap and the labor market pressure indicator. The unemployment rate gap peaked higher than the first principal component during the

⁷ See, among others, Stock and Watson (2002).

⁸ Figures 1 through 5b display the series as percentage deviations from a trend. To compare the principal component series to the unemployment rate gap, we rescale these summary indicators to have the same mean and standard deviation as the unemployment rate gap over the sample periods under consideration.

mid-1970s and early 1980s recessions. In the late 1990s, the principal component indicated a tighter labor market than the unemployment rate. Still, the correlation between the two series, at 0.96, is very high.

II. Robustness

We examine the robustness of our findings along several dimensions. First, we compare the unemployment rate to the summary indicator of labor market pressure as measured by the first principal component of the twelve series when the indicator is estimated in a more “real time” and dynamic fashion. Specifically, instead of estimating the summary indicator over the whole sample period, we estimate the summary indicator at each point in time over a 10-year rolling sample.⁹ This procedure lets the weights given to the original series change over time, possibly making the first principal component a better summary of labor market pressure at each point in time. Figure 2 shows the evolution of the unemployment rate gap and of the dynamically-estimated first principal component over the period 1982:Q1 to 2002:Q4. The picture shows that, even in this case, the unemployment rate gap tracks the summary indicator closely.¹⁰

We have already mentioned that a principal component is a linear combination of the underlying variables. Given the close correspondence between the unemployment rate gap and the estimated first principal component, it is tempting to think that the principal component gives a preponderant weight to the unemployment rate—one of the underlying twelve variables—in the linear combination. It turns out, however, that this is not the case: the unemployment rate has about 10 percent weight in the construction of the first principal component. Altogether, the six series belonging to the Current Population Survey receive a weight of 50 percent. Figure 3 compares the first principal component constructed with all twelve series—the same depicted in Figure 1—

⁹ In other words, the first principal component is estimated at each point in time t on the basis of ten years of data available up to time t . This means that no information from time t onward is used in constructing the principal component, except for future information contained in the estimated trend for each of the underlying series.

¹⁰ The correlation between the two series is 0.97.

with the first principal component constructed when using only the six series that do not belong to the Current Population Survey. Both of the two first principal component measures look very similar, with relatively small divergences that do not persist over time. Thus, the unemployment rate gap remains a good summary of labor market pressure, even when the unemployment rate series is not included in the information set upon which the principal component is based.

Our main analysis considers a parsimonious set of labor market variables, but it is interesting to check the robustness of the findings by using a larger information set. For this purpose we consider, in addition to our twelve original variables, the larger set of labor market variables in Stock and Watson (2002).¹¹ (These variables and the summary indicator constructed with them are denoted “S&W” in the Appendix and figures.) Stock and Watson use a very large set of variables—including labor market activity variables—to extract principal components that in turn are used to predict different measures of activity and inflation.¹² Essentially, our expanded information set encompasses more disaggregated information about labor market activity from the Current Population Survey and the Current Employment Statistics Survey. As with our original variables, the additional variables are expressed as deviations from an estimated trend. Figure 4 compares the first principal component estimated with the original twelve variables—the same series depicted in Figure 1—with the first principal component estimated on the expanded information set. This new principal component explains 68 percent of the variability in the underlying series. From Figure 4, it is evident that the two indicators track each other closely. Hence, the new estimated principal component closely resembles the unemployment rate gap as well.

Overall, the results in this section confirm the close relationship between various principal-component-based summary indicators of labor market pressure and the

¹¹ Four of our original twelve variables are included in Stock and Watson’s larger informational data set.

¹² Bernanke, Boivin, and Elias (2005) use the same set of variables as Stock and Watson to construct principal components, and then model the effects of monetary policy on the principal components.

unemployment rate gap.

III. Recent Evolution of the Unemployment Rate Gap and the Summary Measure of Labor Market Pressure

We now extend the analysis to the most recent years, increasing our sample to include data from 2003: Q1 through 2007: Q2. The comparison between the unemployment rate gap and the principal-component-based summary indicator of labor market pressure is particularly interesting in the current situation because some of the data from the most recent quarters have conveyed a less sanguine picture of the labor market than has the unemployment rate. For example, the percentage of civilians unemployed for 27 weeks or longer has been high and the job-finding rate low relative to the unemployment rate. Consumers' appraisal of the labor market has not been as favorable as in the past, given the relatively low level of the unemployment rate.

As mentioned earlier in Section I, the analysis is complicated by the added uncertainty about a series' estimated trend near its endpoints. Because of this uncertainty, we consider two different assumptions in regards to the trend. In the first case, we simply take the trend beyond 2002:Q4 as estimated by the Hodrick-Prescott filter. In the second case, the trend over the most recent period is kept flat at the estimated 2002:Q4 value for all the series whose steady-state does not depend on growth in the labor force.¹³ For the series that do depend on labor force growth (such as non-farm payroll employment and total hours), over the period 2003:Q1 to 2007:Q2 the trend is assumed to grow at a rate consistent with the growth rate of the labor force during that same period.¹⁴

¹³ The only exception is the help-wanted index series, for which the trend beyond 2002:Q4 is still estimated by the Hodrick-Prescott filter. Keeping the trend flat at the 2002:Q4 value over the most recent period would have been problematic, since changes in the search process, brought by the increased reliance on the Internet as a job advertisement tool, are most likely to have decreased the trend level of jobs advertised in newspapers.

¹⁴ The average annual growth rate of the labor force over the period 2002:Q4 to 2007:Q2 was 1.15 percent.

Using the information set composed of the original twelve series, under the two different assumptions about the evolving trends for the underlying variables over 2003:Q1 to 2007:Q2, Figures 5a and 5b compare the evolution of the unemployment rate gap and the first principal component measure of labor market pressure. As the figures show, the pattern for the estimated first principal component is hardly affected by using alternative de-trending methods. The unemployment rate gap over this period suggested more labor market pressures than the principal component. Since the beginning of 2006, the unemployment rate gap has moved somewhat beyond full labor market capacity, while the summary indicators continued to suggest a labor market at full employment. The difference between the summary indicator and the unemployment rate gap depends on the de-trending assumption used for the unemployment rate. When the recent trend for the unemployment rate is estimated via the Hodrick-Prescott filter (the 2007:Q2 estimate for the unemployment rate's trend value is about 4.75 percent), this difference is relatively small. When the trend for the unemployment rate is kept flat at its 2004:Q4 value (estimated to be 5 percent) over the most recent period, the difference starts to become more significant by historical standards. These findings indicate that the unemployment rate gap may have overstated the extent of labor market tightness since 2005, though the extent of the overstatement is uncertain.

IV. Implications for Wage and Price Inflation

Ultimately, one reason why policymakers are interested in accurate measures of labor market tightness or slack is that knowledge of the present state of the labor market helps improve the assessment of current and future inflationary pressures. Here, we show some implications for wage and price inflation stemming from the behavior of the summary measure of labor market pressure vis-à-vis the unemployment rate gap. For this purpose, we use a standard Phillips curve framework that relates wage or price inflation to labor market pressure, wage or price inflation expectations, and a set of other control variables. The Phillips curve specification takes the following form:

$$\Delta x_t = (1 - \alpha_0) \sum_{i=1}^j \alpha_i \Delta x_{t-i} + \alpha_0 \Delta x_t^f - \beta y_t + \boldsymbol{\gamma} \mathbf{z}_t + \varepsilon_t ,$$

where Δx is the quarterly change in wages or prices, Δx^f is a proxy for wage or price inflation expectations, y is the measure of labor market pressure, \mathbf{z} is a vector of control variables¹⁵, and ε is an error term. The sum of the coefficients α_i , $i = 1$ to j , is constrained to unity. The equation says that inflation at time t depends on a weighted average of past inflation and inflation expectations, on current labor market conditions, and on a set of control variables. For the wage Phillips curve specification, the proxy for wage inflation expectations is given by the sum of estimated trend inflation and estimated trend labor productivity growth in the non-farm business sector.¹⁶ The set of controls in the wage Phillips curve specification include changes in the minimum wage and in employers' social insurance premiums. For the price Phillips curve specification, the inflation expectations proxy is given by 10-year inflation expectations,¹⁷ while the set of controls includes the change in food and in energy prices, and the change in the price of non-oil imported goods.

As a measure for wages, we use the Employment Cost Index (ECI) for private industry workers. The measure for prices is given by the deflator for personal consumption expenditures. Both the wage and the price Phillips curve specifications are estimated over the period 1980:Q1 to 2003:Q4. While we discuss findings based on this sample period only, the results do not appear sensitive to different estimation periods.

Based on the Phillips curve specification when using the unemployment rate gap and the summary measure of labor market pressure as alternative variables for y in the specification, we compare predictions for wage and price inflation. Specifically, we show

¹⁵ The control variables for the price Phillips curve specification consist of energy and food price changes, non-oil import price changes, and a term capturing the impact of the wage-price controls during the Nixon presidency. The control variables for the wage Phillips curve specification are the minimum wage, employers' contributions for social insurance, and a term capturing the impact of the wage-price controls during the Nixon presidency.

¹⁶ We estimate trend inflation and trend labor productivity using a Hodrick-Prescott filter on the deflator and on output per hour in the non-farm business sector, respectively.

¹⁷ The 10-year inflation expectations are taken from the Hoey/Federal Reserve Bank of Philadelphia survey.

the results from dynamic simulations over the period 1996:Q1 to 2007:Q2 for wage and price inflation. In the simulations, we feed back the lags of forecasted values for wage and price inflation into the Phillips curve specification, so that the only actual values for wage and price inflation entering the simulations are the pre-1996 values.¹⁸ The chosen period for the simulations is relevant because in the mid-to-late 1990s and since 2005, there have been divergences between the unemployment rate gap and the summary measure of labor market pressure. Therefore, it is interesting to discover the impact these differences may have for predicting wage and price inflation.

Figure 6 shows actual wage inflation and predicted wage inflation from the Phillips curve specification using either the unemployment rate gap or the summary measure of labor market pressure to represent y . In order to smooth the noise in quarterly movements in inflation, note that the figure depicts the 4-quarter moving average of actual and predicted wage inflation.¹⁹ The fraction of inflation variability captured by the Phillips curve specification is relatively small. As a result, the predicted values appear to capture, at most, the broad contours of actual wage movements. This imprecision notwithstanding, it is interesting to note that, over the simulation period, the predicted value for wage inflation using the summary measure of labor market pressure has a smaller root mean squared error than the predicted value using the unemployment rate gap. The difference in root mean squared errors, however, is less than 10 percent. In the late 1990s, the summary indicator of labor market pressure signalled more labor market tightness than did the unemployment rate gap, and hence indicated higher wage inflation. The opposite holds true since 2005.

Figure 7 shows actual price (personal consumer expenditures) inflation and predicted inflation from the Phillips curve specification using either the unemployment rate gap or the summary measure of labor market pressure to represent y . Again, the figure depicts 4-quarter moving averages. The difference between the predictions from

¹⁸ We still use historical values for expected inflation in the simulations. The weight that inflation expectations receive in the estimated Phillips curve specifications is no greater than 25 percent.

¹⁹ For this reason, Figures 6 and 7 start in 1997:Q1 rather than at the beginning of the simulation period.

the specification using the principal components summary measure of labor market pressure versus the specification using the unemployment rate gap is very small. The reason is that the estimates for the coefficient β , which measures the sensitivity of inflation to labor market conditions, are much smaller than the estimates obtained from the wage Phillips curve specification. This is true irrespective of the chosen measure of labor market slack. As with wage inflation, the root mean squared error from the specification with the summary measure of labor market activity is marginally smaller than the root mean squared error from the specification with the unemployment rate. The figure suggests, though, that this difference is very small.

Overall, the simulations in this section do not offer clear-cut evidence whether the summary measure of labor market pressure or the unemployment rate gap is the more accurate predictor of wage and price inflation. There is some indication that the summary measure of labor market pressure is more consistent with the recent evolution of wage inflation than is the unemployment rate gap. But there is little statistical evidence that this finding would lead to preferring one measure over the other. In this respect, it should also be noted that the Phillips curve estimates over the sample we consider, for both wage and price inflation, yield essentially the same R^2 when using the two different measures of labor market pressure.

V. Conclusions

We have shown that a single factor—the first principal component—can summarize a large portion of the variance among a wide array of labor market indicators. This summary indicator is highly correlated with the unemployment rate gap. Given this finding, the emphasis usually placed on the unemployment rate as measuring the current state of the labor market is justified. For aggregate analysis, the unemployment rate gap is a good summary statistic for a wider set of labor market information.

Periodically there are occasional deviations between the unemployment rate gap and the summary indicator of labor market pressure. Under these circumstances, the consideration then becomes which variable the policymaker should follow. As for the

implications for inflation drawn from different estimates of labor market pressure, the answer is not clear-cut. Phillips curve specifications for wage and price inflation yield essentially the same fit when activity is measured either by the summary indicator of labor market pressure or by the unemployment rate gap. In other words, deviations between the two measures have not been statistically sizable enough, either in magnitude or duration, to favor one measure over the other.

While in this brief we have explored some of the implications of labor market pressure for predicting wage and price inflation, the summary indicator of labor market pressure can also be useful for inferring the size of the output gap. Okun's law relates the unemployment rate gap to the output gap, but it is possible to estimate a similar relationship by using the summary labor market indicator in place of the unemployment rate gap. In the present situation, in which the summary indicator suggests that the labor market is weaker than the unemployment rate gap suggests, the output gap measure derived from the summary indicator would be essentially zero as of 2007:Q2. This result compares to an estimated positive output gap of about 0.9 percentage points from a standard Okun's law analysis, where the trend value of the unemployment rate is estimated at 5 percent.

Appendix


Series Name	Used In:
Civilian Unemployment Rate: 16 yr +	PC, S&W
Civilian Participation Rate: 16 yr +	PC
Civilians Unemployed For Less Than 5 Weeks (Thou.)	PC, S&W
Civilians Unemployed For 27 Weeks or More (Thou.)	PC
Work Part Time: For Economic Reasons, All Ind.	PC
All Employees: Total Nonfarm	PC, S&W
Average Weekly Hours: Total Private Industries	PC
Aggregate Hours: Nonfarm Payrolls, Total	PC
Initial Claims for Unemployment Insurance	PC
Job-finding rate, R. Shimer (2005)	PC
Index of Help-Wanted Advertising in Newspapers	PC, S&W
Conference Board: Jobs Hard to Get	PC
Ratio of Help-Wanted Ads to Number of Unemployed	S&W
Civilian Labor Force: Total Employed	S&W
Civilian Labor Force: Nonag. Ind. Empl.	S&W
Duration of Unemployment: Average (in weeks)	S&W
Civilians Unemployed For 5-14 Weeks (Thou.)	S&W
Civilians Unemployed For 15 Weeks or More (Thou.)	S&W
Civilians Unemployed For 15-26 Weeks (Thou.)	S&W
Employees: NonAg. Payrolls: Total Private	S&W
Employees: NonAg. Payrolls: Goods-Producing	S&W
Employees: NonAg. Payrolls: Mining	S&W
Employees: NonAg. Payrolls: Contract Construction	S&W
Employees: NonAg. Payrolls: Manufacturing	S&W
Employees: NonAg. Payrolls: Durable Goods	S&W
Employees: NonAg. Payrolls: Nondurable Goods	S&W
Employees: NonAg. Payrolls: Service-Producing	S&W
Employees: NonAg. Payrolls: Trans. & Public Util.	S&W
Employees: NonAg. Payrolls: FIRE	S&W
Employees: NonAg. Payrolls: Government	S&W
Employees: NonAg. Payrolls: Wholesale & Retail	S&W
Avg. Wkly. Hrs., Prod. Workers: Manufacturing	S&W
Avg. Wkly. Hrs., Prod. Workers: Manufacturing O/T	S&W
NAPM Employment Index	S&W

Note: Total Employed in Civilian Labor Force and Nonagricultural Industries Employed in Civilian Labor Force are both adjusted for population growth. "Used In" indicates which principal components measure(s) the series were used in. "PC" denotes the 12 series used in our principal components-based labor market activity measure, and "S&W" indicates that the series was included in Stock and Watson (2002).

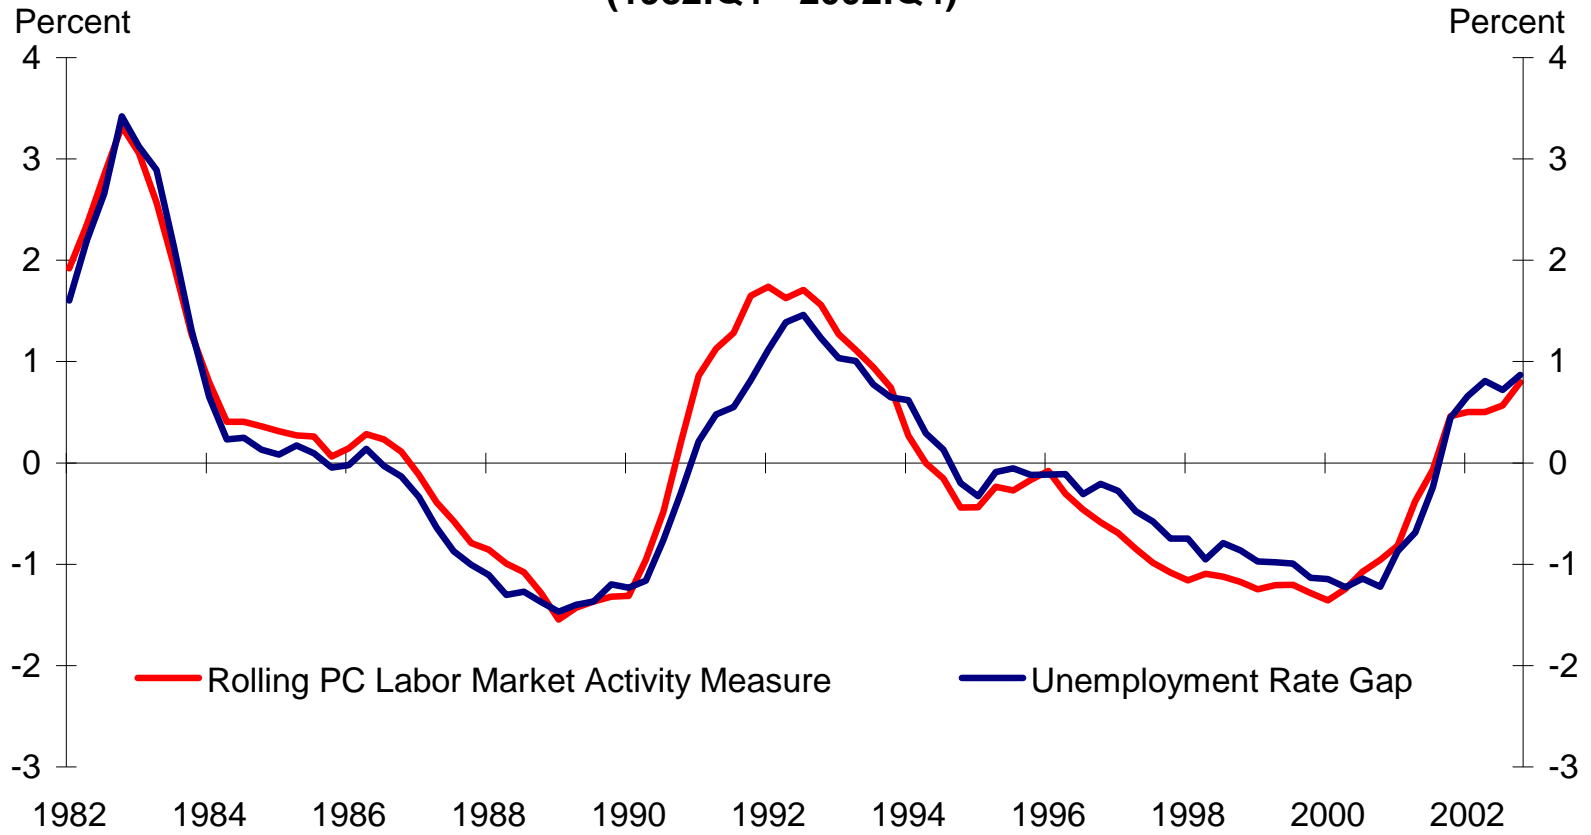
References

- Aguiar, Mark, and Hurst, Erik, "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *Quarterly Journal of Economics* (122, no. 3), pp. 969–1006.
- Bernanke, Ben S., Boivin, Jean, and Elias, Piotr, "Measuring the Effects of Monetary Policy: A Factor-Augmented Vector Autoregressive (FAVAR) Approach," *Quarterly Journal of Economics* (120, no. 1), pp. 387–422.
- Hodrick, Robert J., and Prescott, Edward C., "Post-War U.S. Business-Cycles: An Empirical Investigation," *Journal of Money, Credit, and Banking* (29, no. 1), pp. 1–16.
- Mise, Emi, Kim, Tae-Hwan, and Newbold, Paul, "On Suboptimality of the Hodrick-Prescott Filter at Time Series Endpoints," *Journal of Macroeconomics* (27, no. 1), pp. 53–67.
- Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *The American Economic Review* (95, no. 1), pp. 25–49.
- Stock, James, and Watson, Mark, "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business Economics and Statistics* (20, no. 2), pp. 147–162.


**Figure 1: Unemployment Rate and PC Labor Market Activity
Gap Measures
(1973:Q1 - 2002:Q4)**



**Figure 2: Unemployment Rate and Rolling PC Labor Market Activity
Gap Measures
(1982:Q1 - 2002:Q4)**



**Figure 3: PC Labor Market Activity Gap Measures With and Without
Current Population Survey Data
(1973:Q1 - 2002:Q4)**



**Figure 4: PC and S&W Labor Market Activity Gap Measures
(1973:Q1 - 2002:Q4)**




Figure 5a: Unemployment Rate and PC Labor Market Activity Gap Measures (using HP filter trend) (2003:Q1 - 2007:Q2)

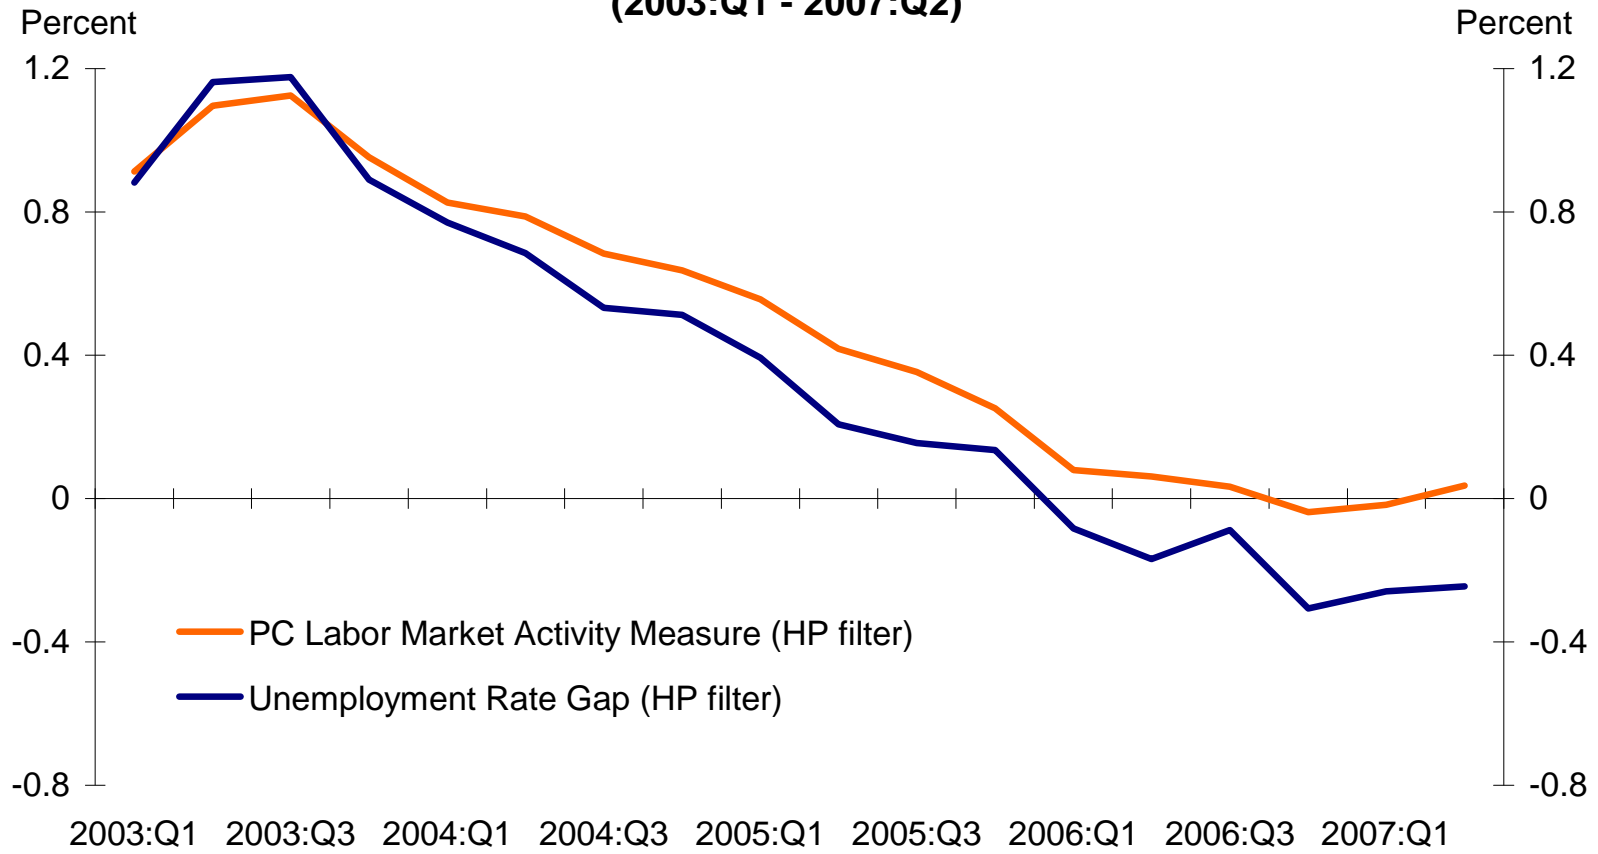


Figure 5b: Unemployment Rate and PC Labor Market Activity Gap Measures (with trends kept flat) (2003:Q1 - 2007:Q2)




Figure 6: Actual and Predicted Wage Inflation from Phillips Curve Specification Using the Unemployment Rate and PC Labor Market Activity Gap Measures (4-quarter moving average, 1996:Q1 - 2007:Q2)




Figure 7: Actual and Predicted Price Inflation from Phillips Curve Specification Using the Unemployment Rate and PC Labor Market Activity Gap Measures (4-quarter moving average, 1996:Q1 - 2007:Q2)

