

ME
WEEK 1
7/16 1:00 P.M.

RECEIVED
JUN 12 1999
FEDERAL RESERVE BOARD

What's Happened to the Phillips Curve?

Flint Brayton, John Roberts, and John Williams*

Division of Research and Statistics

Federal Reserve Board

Washington, D.C. 20551

June 1999

*The authors acknowledge the comments and assistance of Peter Tulip. Views presented are those of the authors and do not necessarily represent those of the Federal Reserve Board or its staff.

1 Introduction and Summary

The rise and fall of inflation in the United States during the 1970s and 1980s provided a testing ground for the Phillips curve model of inflation dynamics. And, according to some observers (Fuhrer 1995, Gordon 1997), such models performed admirably well in tracking actual inflation, both within and out of sample. Based on this achievement, many became convinced of the usefulness of such models as tools in predicting inflation. Unfortunately, the main feature of empirical Phillips curve models, that is, that inflation rises when markets tighten, appears to be turned on its head during the economic expansion of the 1990s, when the unemployment rate fell below its long-run average of around 6 percent and then slid under 5 percent, while inflation fell.

One objective of this paper is to ascertain whether the recent performance of inflation is statistically surprising in the context of a canonical Phillips curve in which price inflation depends on the unemployment rate, past price inflation, and standard measures of price supply shocks. To provide a graphical preview of our conclusion that a significant shift likely has taken place in such a baseline model, figure 1 shows forecast errors for our preferred CPI equation.¹ The upper panel of the figure contains one-quarter-ahead forecast errors based on successive reestimates of the equation over samples whose starting point is held fixed at 1955:Q1 and whose ending point advances one quarter at a time. The forecasts make use of the actual values of explanatory variables.² While no individual error since 1987 lies outside of the 95 percent confidence range, inflation has consistently fallen short of the model's predictions over the past five years.

Given the run of negative prediction errors, multi-period forecast errors and their confidence bands provide a better graphical depiction of the probability of the equation's recent performance. The lower panel of the figure plots the sequence of four-quarter-ahead forecast errors. These errors are derived in a manner analogous to that used for the one-step-ahead errors, except that within each four-quarter forecast interval, the forecast is a

¹The inflation series is the the all-items CPI index, modified from 1967 to 1983 so that homeowner-ship costs are on a rental-equivalent basis, and adjusted since 1995 to eliminate effects of methodological changes. The explanatory variables consist of 24 quarterly lags of past inflation, contemporaneous values of a demographically weighted unemployment rate and a composite measure of relative food and energy price movements, an intercept, and a variable for the 1970s wage-price controls. The structure of this equation is discussed in more detail below.

²The use of conditional forecasts is appropriate because the issues being examined concern the structure of the Phillips curve itself. Errors in forecasting the unemployment rate, for example, are not germane because conceptually they are associated with the performance of other equations in a forecasting system.




Figure 1

dynamic simulation in which inflation lags are set equal to simulated values. The most recent four-quarter forecasts have overpredicted inflation by as much as 1-1/2 percentage points, well outside of the 95 percent confidence range, which encompasses forecast errors of up to 1 percentage point.

A tendency to over-predict inflation starting in the middle of this decade characterizes baseline equations for all six measures of inflation we study. For each of four consumption prices, the magnitude of the prediction errors is sufficient to reject the hypothesis of stability of the equation's intercept at traditional significance levels: for two product prices, the rejection of stability is not as strong, although errors are almost as big. In the framework of these equations, instability of the intercept is equivalent to instability of the NAIRU.

We find a close association between the number of lags of inflation included in the baseline equations and whether or not evidence of significant structural change is found. Equations such as ours with long inflation lags (15-24 quarters) are unstable while those with short lags are not. In essence, estimates of the NAIRU are much more precise in long-lag equations than in short-lag equations, and consequently the level of the unemployment rate over the past few years is well below the 95 percent NAIRU confidence range of long-lag equations but not that of short-lag equations.

The relationship between inflation lag length and NAIRU precision has been documented by Staiger, Stock, and Watson (1997)—hereafter, SSW—but they emphasize the wide confidence intervals for the NAIRU found in short-lag equations on grounds that such lag lengths are optimal according to either standard hypothesis tests or the use of information criteria to select lag lengths. Our long-lag equations are the outcome of using an information criteria to choose simultaneously lag lengths and coefficient smoothing restrictions. Permitting smoothing restrictions leads to equations that have much longer lags on past inflation than is the case when lag coefficients are unrestricted, but with only a few independently estimated parameters because the lag coefficients are restricted to lie on low-order polynomials. We show that these long-lag Phillips curves, in addition to yielding much more precise estimates of the NAIRU, have better out-of-sample forecasting properties than do the short-lag ones. Monte Carlo evidence indicates that the procedure of jointly searching over lag lengths and smoothing restrictions is reasonably accurate at finding the “correct” model whether or not it contains smoothed inflation coefficients.

Our baseline analysis provides support for the type of Phillips curve long favored by

Robert Gordon who for many years has reported a specification with 24 quarterly lags of past inflation.³ The tendency of our baseline equations to significantly over-predict inflation since the mid-1990s, however, is an indication of structural change—perhaps a decline of the NAIRU—or misspecification. Because the baseline equations contain the relative rates of import price inflation and food and energy price inflation, the statistical evidence also permits the conclusion that recent declines in the relative prices of these supply variables are not large enough to account fully for the low rate of inflation.

We investigate two possible explanations for the breakdown or near-breakdown of the price-price Phillips curve. One is the possibility that the capacity utilization rate—which has been near its historical average during the past few years—is a better measure of macroeconomic slack than the unemployment rate, in which case the recent behavior of inflation would not be that extraordinary. From a conceptual perspective, the unemployment rate might be preferred because of its broader sectoral coverage. But as it turns out, movements of the utilization and unemployment rates are sufficiently similar over the past 40 or so years that the goodness of fit of equations based on one is quite similar to the fit of equations based on the other. Nonetheless, we find that Phillips curve equations that include the unemployment rate fit somewhat better over the postwar years than do those using capacity utilization. The outcome is closer to a toss up in out-of-sample forecasts, however. All told, the evidence does not provide strong support for the use of capacity utilization but neither does it suggest that the utilization rate is dominated by the unemployment rate.

Our preferred explanation is based on an augmented Phillips curve that includes the level of the markup of price relative to trend unit labor costs as an error correction term. We find that the level of the markup in the nonfarm business sector is highly significant in equations for all measures of inflation examined, with a high markup estimated to restrain inflation and a low markup putting upward pressure on inflation. Equations that include the markup display much weaker evidence of instability. In particular, a high level of the markup in the mid-1990s is estimated to have depressed price inflation.

The markup story is not without caveats, however. First, the markup has recently fallen back to a “normal” level and thus has not been much of a restraint on predicted inflation over the past year or so—a period in which the markup equations overpredict inflation by a substantial amount. Because the overprediction episode is as yet fairly short, it is too soon

³The most recent exposition is Gordon (1998).

choose the best number of lags.

Our method of choosing lag specification—on inflation, unemployment, and the two supply variables—draws on both of these approaches. Similar to SSW, we use an information criterion to choose optimal lag lengths. And in the spirit of Gordon, we permit smoothing of the lag coefficients, but by constraining the coefficients to lie on a polynomial function of the lag index—PDL restrictions. The Schwartz criterion is used to choose lag lengths and the optimal degree of PDL smoothing.

We constrain coefficients on lagged inflation to sum to one, a restriction consistent with the (near) unit root behavior of postwar inflation. A maximum of 25 lags of inflation is permitted. The unemployment rate, which may enter contemporaneously and with up to three lags, is a fixed-weight aggregate of unemployment rates of six age and gender categories using 1993 labor force shares.⁴

Relative import price inflation, which may enter contemporaneously and with up to three lags, is constructed by subtracting from a measure of import price inflation the lagged value of the aggregate inflation series being explained. For the four consumption price equations, a core measure of import prices is used that excludes oil, computers, and semiconductors. This choice reflects the small share of “high-tech” goods in personal consumption expenditures. The two product prices have bigger high-tech shares than does consumption, but the shares are still well short of the contribution of such goods to imports. Equations for the product prices use an import price that has an intermediate weight on computers and semiconductors.⁵

⁴Baseline results are unaffected if the contemporaneous unemployment rate is excluded.

⁵Roughly speaking, GDP and NFB have about one-third the high-tech intensity of imports. Consequently, the import price series we use in the product price equations includes about one-third of the contribution of computers and semiconductors to import prices.

The measure of relative food and energy price inflation entering each equation also varies with the particular measure of aggregate inflation being explained. For CPI equations, the difference between overall and core CPI inflation is used, while PCE equations contain the difference between overall and core PCE inflation. This PCE-based measure is also used in GDP equations, while an adequate account of the short-run dynamics of NFB inflation requires that separate food and energy price variables be included.⁶ Relative food and energy price inflation may enter each equation lagged once up to a maximum of four lags, with smoothing restrictions permitted on the associated coefficients. In addition, contemporaneous relative food and energy price inflation is assumed to enter equations for all measures of inflation aside from CPIX and PCEX and is assumed to be excluded from the latter.

Each inflation equation also includes a constant and a dummy variable for wage and price controls in the 1970s. The latter is constructed so that it has a mean of zero and equals 1.0 from 1971:Q3 to 1974:Q1 and -3.67 from 1974:Q2 to 1974:Q4.

2.2 Equations with structure chosen through 1987

The explanatory variables, lag lengths, and coefficient smoothing restrictions in our initial set of baseline equations are selected through 1987 and represent specifications that a researcher might have developed prior to observing the low inflation of the 1990s.

The Baseline-87 CPI specification, which is shown in the first column of table 1 and was used for the out-of-sample forecasts reported in figure 1, contains 24 lagged values of past CPI inflation whose coefficients are constrained to lie on a second degree polynomial, and the contemporaneous values the demographically weighted unemployment rate and the relative price of food and energy. Import prices are excluded. The lag length on inflation is identical to that in Gordon's equations, but here the inertial element of inflation is based on only three estimated parameters compared with Gordon's six. In terms of the demographically weighted unemployment rate, the NAIRU is precisely estimated to be 6.0

⁶While lags of the PCE-based measure appear to work well in NFB equations, the contemporaneous observation does not. Rather, the contemporaneous value of a separate food price variable is required to capture the tendency of NFB prices to move initially in the opposite direction of a change in the price of the farm sector, because the cost of farm sector inputs is not immediately and fully passed through to the price of finished food products. In recognition of this, our NFB equations contain the contemporaneous values of relative farm sector price inflation and relative PCE energy price inflation, and lags (if warranted by the selection process) of the PCE-based food-energy aggregate.

Table 1: Baseline-87 Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Specification Selection Period (-1987.4)						
Equation Specification: ^a						
inflation lags ^{pdl degree}	24 ²	24 ²	24 ²	23 ⁴	24 ²	18 ¹
unemployment rate	x	x	x	x	x	x
food and energy price	x	x	x	x	x	x
import price		x			x	x
NAIRU	6.02	6.01	6.25	6.37	5.96	5.79
(std. err. ^b)	(.13)	(.15)	(.22)	(.25)	(.16)	(.15)
Extended Estimation Period (-1998.4)						
Mean residual (97.1-98.4)	-.74	-.79	-.54	-.61	-.60	-.64
Stability tests: ^c						
all coefficients	**	***	-	**	-	-
intercept only	***	***	**	***	*	-
NAIRU shift: ^d						
most likely date	94.4	95.1	93.3	95.3	95.3	97.3
pre-shift NAIRU	6.10	6.06	6.22	6.18	6.03	5.89
post-shift NAIRU	4.76	4.06	4.81	3.74	4.63	4.07

^aCoefficient estimates and regression statistics are presented in the appendix.

^bOne-half of the 70 percent confidence range.

^cSignificance level: > .10(-), .10(*), .05(**), or .01(***).

^dFrom the equation whose dating of an intercept shift yields the Chow statistic with the lowest p-value.

percent with a standard error of 0.13, where the latter is measured as one-half of the width of the 70 percent Gaussian confidence interval.

Common to the set of Baseline-87 equations are NAIRU estimates clustered around six percent with small standard errors, lag lengths on inflation that are long (ranging from 18 to 24 quarters), coefficients on lagged inflation that are highly smoothed, and the inclusion of relative food and energy prices. The notable difference across equations is the appearance of import prices in the two product price equations but in only one of the consumption price equations.

Before addressing the question of how well the baseline equations hold up in the 1990s, we need to ascertain if its coefficients are stable within its estimation sample. A finding of instability in the current decade would be less noteworthy if the equation were unstable earlier on. We test for two aspects of coefficient stability using the $Exp - F_\infty$ test of Andrews, Lee, and Ploberger (1996) for structural change at an unknown date. One is a joint test that all coefficients are stable; the other is for stability of the intercept, which in most instances is equivalent examining the stability of the NAIRU. In the absence of a prior about the type of structural change that might have occurred, testing the stability of all coefficients simultaneously is appropriate. If we believe that any instability is confined to the constant term, however, the joint test has low power and the test of intercept stability is preferred. Shifts in the intercept seem more likely than shifts in other coefficients, but we do not want to rule out the latter. Thus, we report both tests but focus on the stability of the intercept.

No evidence of coefficient instability is found in the CPI equation over a sample period of 1955:Q1 to 1987:Q4. Similar results obtain in equations for the five other measures of inflation, although for the two product prices, the stability outcome depends on advancing the start of the estimation interval from 1955:Q1 to 1960:Q1. Even later starting dates for the two core inflation equations are dictated by data availability.⁷

As shown earlier in figure 1, out-of-sample forecast errors for the baseline CPI equation indicate a persistent overprediction of the rate of inflation starting in 1994. We now formally test the stability of this equation over a time span that includes the period of surprisingly low inflation. Although many have come to believe that factors such as increased international competition or lower wage demands arising from heightened job insecurity

⁷Estimation periods begin in 1963:Q3 for CPIX and in 1965:Q3 for PCEX.

have restrained inflation, our view is that these hypotheses have been motivated primarily by the *ex post* observation that inflation has been low relative to predictions. Rather, in the absence of a specific prior, we assess the stability of the CPI equation at all dates in the expanded estimation period, which runs from 1955:Q1 to 1998:Q4.

In the tests for stability of all coefficients jointly, the test rejects stability at the five percent level, as shown in the bottom panel of table 1. An even stronger case for a shift in the Phillips curve comes from the test sequence when only the intercept is allowed to change. Here, the test indicates a probability of no shift that is less than 1 percent. Based on the shift date that provides the best fit (and for which the Chow statistic attains its lowest p-value), the shift in the intercept most likely occurred at 1994:Q4, with the NAIRU estimated to fall from 6.1 percent up until 1994 to 4.8 percent thereafter. Clearly, a limitation of this approach is that it assumes that the shift occurred in a single quarter whereas it may in fact have taken place over an extended period. Still, it gives us some idea of the timing and magnitude of the shift.

Based on the stability tests, all four consumption price equations have unstable constants at a significance level of five percent or lower for the estimation period that extends through the end of 1998. Equations allowing an intercept shift at the date that yields the best fit indicate that the NAIRU declined from just above 6 percent prior to the shift date (ranging from 1993 to 1995) to between 3-3/4 and 4-3/4 percent thereafter.

The evidence for a structural break is weaker in the two product price equations, with only GDP having a p-value for an intercept shift of less than ten percent. A likely explanation for the less significant estimates of intercept shifts is that product price inflation is predicted less precisely than consumption price inflation. The noisier behavior of product prices makes it more difficult to detect a structural change of given size. This explanation is consistent with the finding that the mean residuals for the product price equations over the past two years are not markedly different than the mean residuals for the consumption price equations.

2.3 Equations with structure chosen through 1998

We next examine whether our conclusions about equation stability hold up when the period over which the specifications are determined is extended through 1998 (Baseline-98). As shown in table 2, use of the longer sample has only minor effects. As before, all specifica-

Table 2: Baseline-98 Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Specification Selection Period (-1998.4)						
Equation Specification: ^a						
inflation lags ^{pdl degree}	24 ²	24 ²	23 ²	24 ²	15 ²	17 ²
unemployment rate	x	x	x	x	x	x
food and energy price	x	x	x	x	x	x
import price		x		x	x	x
NAIRU	5.97	5.88	5.99	5.86	5.88	5.80
(std. err. ^b)	(.11)	(.15)	(.18)	(.26)	(.18)	(.16)
Mean residual (97.1-98.4)	-.74	-.77	-.52	-.44	-.50	-.51
Stability tests: ^c						
all coefficients	**	***	-	-	-	-
intercept only	***	***	**	**	-	-
NAIRU shift: ^d						
most likely date	94.4	94.4	93.3	95.1	97.3	97.3
pre-shift NAIRU	6.10	6.06	6.21	6.17	5.97	5.88
post-shift NAIRU	4.76	4.12	4.83	4.00	3.79	4.11

^aCoefficient estimates and regression statistics are presented in the appendix.

^bOne-half of the 70 percent confidence range.

^cSignificance level: > .10(-), .10(*), .05(**), or .01(***).

^dFrom the equation whose dating of an intercept shift yields the Chow statistic with the lowest p-value.

tions contain long lags on past inflation and yield estimates of the NAIRU that are close to six percent. Import prices appear in four of six equations, rather than three.

With regard to coefficient stability, the most likely date for a shift in the NAIRU again ranges from late 1993 to early 1995 in the consumption price equations, but estimates of the post-shift NAIRU—between 4.1 and 4.8 percent—are a tad higher than before. All four consumption price equations continue to have unstable constants at a significance level of five percent or lower, but for both product prices the significance level of an intercept shift now exceeds ten percent.

2.4 Import prices

In the four Baseline-98 equations which contain the relative rate of import price increase (all but total and core CPI), prices of imported goods are estimated to have restrained do-

mestic inflation lately. Over the past five years—since the unemployment rate fell below 6 percent—the measures of import price inflation we use have increased between 1-1/2 and 2 percentage points more slowly at an annual rate, on average, than have the consumption and product prices under examination. Based on dynamic simulations of the Baseline-98 equations starting in 1994, the relative weakness of import prices is estimated to have restrained inflation last year 1/4 percentage point for the two PCE measures, 1/2 percentage point for GDP prices, and nearly 3/4 percentage point for NFB prices. These effects of import prices are, nonetheless, small relative to dynamic simulation errors over the same period, which build to overpredictions of 1-1/2 to 2-1/4 percentage points in 1998. For import prices to fully explain the low level of inflation last year, their coefficients would have to be from four (NFB) to ten (PCE) times larger than estimated.⁸ Tests of the stability over time of the import price coefficients reveal no significant shifts, as long as equation intercepts are also permitted to change. Thus, recent inflation can be better explained by a permanent but constant adjustment to predicted inflation rather than by a permanent increase in the coefficient on import prices, with effects that vary with the actual pattern of these prices.

Import prices are always included in the product price equations but in only some of the consumption price equations. A quick investigation suggests that this is a result of the price of investment goods having a much higher sensitivity to import prices than does the price of consumer products.⁹ The pattern in which import prices appear in product price equations but less often in consumption price equations is not sensitive to the particular measure of import price, whether it be one of the two import prices we use or the price of nonoil imports without any “high-tech” adjustments, and it also obtains when import prices are weighted by the ratio of imports to GDP.

⁸Estimated coefficient sums are .04 in equations for overall and core PCE, .06 for GDP, and .12 for NFB.

⁹An equation for price inflation of gross domestic purchases contains an import price coefficient that is similar to the one in the GDP price equation. (This implies that export prices are positively correlated with import prices, a result which could reflect responses to exchange rate changes.) Equations for the prices of the three main components of purchases—consumption, gross private domestic investment, and government spending—reveal that the largest and most significant effect of import prices is estimated to occur for the price of investment.

Table 3: No Smoothing Restrictions

	CPI	PCE	CPIX	PCEX	GDP	NFB
Specification Selection Sample Period (-1998.4)						
Equation Specification: ^a						
inflation lags ^{pdl degree}	4 ³	4 ³	5 ¹	5 ¹	4 ³	7 ⁶
unemployment rate	x	x	x	x	x	x
food and energy price	x	x		x	x	
import price	x	x	x	x	x	x
NAIRU	6.02	5.95	5.96	5.79	5.88	5.74
(std. err. ^b)	(.31)	(.38)	(.28)	(.48)	(.38)	(.31)
Mean residual (97.1-98.4)	-.16	-.15	-.14	-.13	-.19	-.20
Stability tests: ^c						
all coefficients	-	-	-	-	-	***
intercept only	-	-	-	-	-	-

^aCoefficient estimates and regression statistics are presented in the appendix.

^bOne-half of the 70 percent confidence range.

^cSignificance level: > .10(-), .10(*), .05(**), or .01(***).

3 Sensitivity of results to inflation lag length and smoothing restrictions

A natural question to ask is whether our conclusions about the stability of baseline inflation equations are sensitive to the use of long inflation lags and coefficient smoothing restrictions. Evidence about NAIRU uncertainty suggests that this may be the case. Standard errors of NAIRU estimates in the Baseline-98 equations, which range from .11 to .26 percentage points, are much smaller than those others have reported using equations with shorter lags on inflation. SSW (1997, p. 196) state that “a typical estimate of the NAIRU in 1990 is 6.2 percent, with a 95 percent confidence interval for the NAIRU in 1990 being 5.1 percent to 7.7 percent.” Their result implies a standard error for the NAIRU that is four times larger than our average estimate.

When our specification selection procedure is modified so that coefficient smoothing restrictions are not permitted (table 3), we estimate equations that are similar to those of SSW in that they have short inflation lag lengths—from four to seven quarters. Point estimates of the NAIRU are little changed, but their standard errors are about twice as large as

Table 4: Out-of-Sample Forecast RMSEs

	CPI	PCE	CPIX	PCEX	GDP	NFB
Four-Quarters-Ahead						
max=25	.61	.72	1.27	.88	.71	1.38
max=12	.75	.73	.85	.74	.74	1.21
no-pdl	.85	.82	1.03	.71	.93	1.65
Eight-Quarters-Ahead						
max=25	1.26	1.48	2.65	2.23	1.36	2.96
max=12	1.80	1.53	1.70	1.58	1.47	2.47
no-pdl	2.12	1.84	2.72	1.46	2.04	3.68

those in the Baseline-98 set of equations. The less precise NAIRU estimates are associated with the result that stability of intercepts cannot be rejected at the ten percent significance level. Clearly, permitting coefficient smoothing restrictions has a substantial effect on key characteristics of inflation equations.

3.1 Out-of-sample forecasts

To verify the robustness of the long-lag equations relative to that of the short-lag specifications, we compute rolling out-of-sample forecasts for equations based on three variants of the selection procedure. The first (max=25), which allows PDL smoothing restrictions and permits up to 25 inflation lags, is the approach used to derive the Baseline-87 and Baseline-98 sets of equations, except that now the selection procedure is rerun over a sequence of estimation periods whose end advances one quarter at a time. The second alternative (max=12) shortens the feasible inflation lag length to 12 quarters while still permitting coefficient smoothing, while the last case does not permit smoothing (no-pdl).

The first out-of-sample forecast interval starts in 1975:Q1 based on models selected and estimated through 1975:Q4, except for PCEX whose initial forecast starts at 1979:Q1, given the shorter historical time span over which data on this measure of inflation is available. For all inflation measures, the last forecast period ends at 1998:Q4. The results in table 4, which reports root mean squared percentage forecast errors for the price level at horizons of four and eight quarters, confirm the conclusion based on in-sample results that long-lag inflation equations tend to be superior to short-lag ones. For all inflation measures aside from PCEX, our model selection procedure (max=25) yields price forecasts that have smaller root mean squared errors than does the procedure not permitting coefficient smoothing (no-pdl). In

some cases, however, the intermediate procedure ($\max=12$) in which smoothing restrictions are permissible but inflation lag lengths are constrained to be no more than 12 quarters yields more accurate forecasts than does the longer-lag variant.

3.2 Monte Carlo analysis

In this section, we address the question of whether model-selection criteria can discriminate between long-lag and short-lag models of the Phillips curve in samples such as the one we are examining. This issue is of interest because some prominent previous work using model-selection criteria to choose lag lengths in the Phillips curve chose short lags (SSW, 1997). However, this work did not entertain the possibility that the coefficients on the lagged dependent variable lie on a low-order polynomial, as we consider here.

We run four main experiments with models of CPI inflation. In the first two, we repeatedly generate artificial data assuming that our long-lag, low-order model for own-lags in the Phillips curve is the correct model (table 2, column 1) and then run two different model-selection procedures, one of which allows for the possibility of low-order polynomial-distributed lags, and the other of which does not. We report the lag lengths chosen by each procedure. In the next pair of experiments, we first run the model-selection procedure ruling out low-order PDLs—that is, with the coefficients on the lags unconstrained. Consistent with the results of SSW, the chosen model has relatively short lags (table 3, column 1). We then repeatedly generate data using this model and again run the two different model-selection exercises on each set of artificial data, and tabulate the length of the lags chosen on inflation.

We use a six-lag VAR to generate simulated data for the other endogenous variables in the model, namely, the unemployment rate; the relative food-and-energy price; and relative import prices. We assume that the unemployment rate is unaffected in the current period by any of the other variables in the model; that relative food-and-energy prices are affected by the unemployment rate but not by any of the other variables; and that relative import prices are affected by unemployment and food-and-energy prices but not aggregate inflation.

Based on one thousand draws, we find that when the model used to generate the artificial data is the long-lags version of the Phillips curve, the specification-search procedure that allows for reduced-order PDLs generally chooses long lags, and that the search procedure which does not allow for reduced-order PDLs selects a short-lag model. In particular, the

median lag length when reduced-order PDLs were allowed is 23, compared to the true lag length of 24 (with the coefficients constrained to lie on a second-order polynomial). In 75 percent of the draws, the chosen lag length is 20 or more. In only five percent of the draws is the chosen lag length four or less. By contrast, when only unconstrained lags are allowed, the median lag length chosen is three, and in 99 percent of cases, the chosen lag length is ten or less. Hence, when the true model involves low order and long lags, a model-specification procedure that does not allow for this possibility will choose a specification with far too few lags.

As noted previously, when we run the model-specification search for a model for the CPI allowing only for unconstrained own-lag specifications, the procedure chooses four lags. Using this specification to generate artificial data, the median number of lags chosen when the model search procedure is limited to consider only unconstrained lags is three, and in 95 percent of the cases, the number of lags is five or fewer. When reduced-order polynomials are also allowed, the median number of lags selected is six while the ninety-fifth percentile of the distribution is 14 lags.

On the whole, the more-general model-selection procedure that allows low-order polynomials out-performs the procedure that permits only unrestricted lags. The median lag estimated by the first method is close to the inflation lag length used to generate the artificial data in all experiments, while the median lag estimated by the second method is severely biased down in one instance.

3.3 Summary of baseline equations

Viewing the complete set of baseline equations, out-of-sample forecasts, and Monte Carlo analysis, the evidence favors the long-lag inflation equations in which the NAIRU is estimated precisely but appears to have been unstable recently over the short-lag equations with wide confidence bands and no evidence of instability. We next turn to the question of whether there are economic explanations other than a decline in the NAIRU for the recent low rates of inflation.

Table 5: Capacity Utilization Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Specification Selection Period (-1998.4)						
Equation Specification: ^a						
inflation lags ^{pdl degree}	24 ²	10 ⁴	9 ⁴	24 ²	15 ²	17 ²
capacity utilization	x	x	x	x	x	x
food and energy price	x	x	x	x	x	
import price	x	x	x	x	x	x
NAICU	81.3	81.6	81.7	81.9	81.7	81.9
(std. err. ^b)	(.43)	(.63)	(.65)	(.77)	(.61)	(.50)
Mean residual (97.1-98.4)	-.08	-.18	.08	-.18	-.03	.11
Stability tests: ^c						
all coefficients	***	-	-	-	-	-
intercept only	*	-	-	-	-	-

^aCoefficient estimates, and regression statistics are presented in the appendix.

^bOne-half of the 70 percent confidence range.

^cSignificance level: > .10(-), .10(*), .05(**), or .01(***).

4 Capacity Utilization

Much has been made of late of the divergent signals about resource utilization coming from the labor and production markets. Although utilization rates in the two typically move together fairly closely, labor utilization has been higher relative to its historical average than has capacity usage. Thus, the rate of capacity utilization seems more consistent with the recent inflation picture than does the unemployment rate. We use the Schwartz criterion to develop a set of equations that include capacity utilization and exclude the unemployment rate. The specification selection period runs through 1998 and thus the capacity utilization equations are comparable to those containing the unemployment rate in table 2. As reported in table 5, only the CPI equation shows significant evidence of coefficient instability when capacity utilization is the measure of demand pressure. And, the instability in this instance is associated primarily with the early part of the estimation period, not the end.

Although capacity utilization might have a slight edge in explaining recent inflation, that is not the case for the whole estimation period. We reestimated the unemployment rate equations with capacity utilization added as an explanatory variable while also reestimating the capacity utilization equations with the unemployment rate added. In terms of statisti-

Table 6: Unemployment Rate (U) vs Capacity Utilization (CU): In Sample

	t-statistic			
	Equations based on U		Equations based on CU	
	U	CU	U	CU
CPI	5.0	0.4	4.2	0.8
PCE	4.2	0.1	1.8	0.9
CPIXFE	2.9	0.6	1.9	1.1
PCEXFE	1.7	0.8	1.7	0.8
GDP	3.2	0.4	3.2	0.4
NFB	3.2	0.6	2.8	1.0

Table 7: Unemployment Rate (U) vs Capacity Utilization (CU): Out of Sample


	CPI	PCE	CPIX	PCEX	GDP	NFB
Four-Quarters-Ahead						
U (max=25)	.61	.72	1.27	.88	.71	1.38
CU (max=25)	.65	.77	.89	.84	.81	1.20
Eight-Quarters-Ahead						
U (max=25)	1.26	1.48	2.65	2.23	1.36	2.96
CU (max=25)	1.37	1.68	1.67	2.06	1.80	2.53

cal significance, the unemployment rate dominates capacity utilization in all instances—irrespective of whether the specification was initially tuned to contain the unemployment rate or capacity utilization (table 6).

A final comparison of unemployment and capacity utilization makes use of out-of-sample forecast errors (table 7). Based on forecasts of equations chosen by our preferred model selection procedure (max=25), the equations with the unemployment rate are more accurate at forecasting some measures of inflation, while equations with capacity utilization are more accurate at forecasting other inflation series. The outcome is a draw.

All told, our evidence neither provides strong support for the use of capacity utilization nor does it suggest that the utilization rate is dominated by the unemployment rate.

Figure 2



5 The Markup

Thus far, the Phillips curves we have examined have clearly been reduced-form relationships. A more structural perspective might view prices as a markup over costs. To pursue this approach, we look at the behavior of the markup of prices in the nonfarm business sector over trend unit labor costs, where the latter is compensation per hour from the Productivity and Cost release, relative to a measure of the trend in productivity that allows a single break in its slope in 1973. As shown by the solid line in figure 2, the markup over trend unit labor costs has no discernable trend historically and has been relatively high over much of the past five years—the period of “unusually” low inflation.¹⁰

Our analysis starts with the specification selection procedure used earlier, but with the set of potential explanatory variables augmented to include the the log of the markup over

¹⁰If the ECI measure of hourly compensation, which begins in 1980, is spliced with the Productivity and Cost series, a similar pattern is found. In addition, the statistical conclusions presented below are largely unaffected if the spliced series is used.

Table 8: Markup (CPH) Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Specification Selection Period (-1998.4)						
Equation Specification: ^a						
inflation lags ^{pdl degree}	24 ³	24 ²	23 ²	24 ²	15 ²	17 ²
unemployment rate	x	x	x	x	x	x
food and energy price	x	x	x	x	x	x
import price		x			x	x
trend markup	x	x	x	x	x	x
NAIRU ^b	5.97	5.88	6.05	6.01	5.92	5.85
(std. err. ^c)	(.09)	(.12)	(.16)	(.18)	(.16)	(.13)
Mean residual (97.1-98.4)	-.54	-.47	-.32	-.44	-.38	-.46
Stability tests: ^d						
all coefficients	-	-	-	-	-	-
intercept only	-	-	-	*	-	-

^aCoefficient estimates and regression statistics are presented in the appendix.

^bMeasured at the mean value of the markup.


^cOne-half of the 70 percent confidence range.

^dSignificance level: > .10(-), .10(*), .05(**), or .01(***).

trend unit labor costs—which we will refer to as the “trend markup.” In an earlier stage of our research, the lagged level as well as several lagged growth rates of the trend markup were part of the set of potential regressors, but the growth rates were never selected for inclusion in the final equations. To economize on computations, the estimates presented here only include the lagged level as part of the candidate set. This specific variable, which is a broad measure of labor cost pressure on the price of domestic production, is used in equations for all six measures of inflation. The unemployment rate is the measure of slack included in the set of explanatory variables.

According to the selection criterion, the lagged level of the trend markup appears in all inflation equations, as shown in table 8. This variable has the implication that, when the price level is high relative to trend unit labor costs, downward pressure is exerted on all the measures of inflation investigated. Given the relatively high level of the markup over much of this decade, residuals of all equations containing the labor cost measure are small until very recently, and only the PCEX equation shows significant evidence of coefficient instability. The instability in this instance is not associated with the end of the estimation

Figure 3



period, however.

The estimated role of the markup in inflation dynamics is not dependent on including the decade of the 1990s in the sample. We find that the markup variable is statistically significant in most of our inflation equations over sample periods that end in 1987.

The markup story is not without caveats, however. First, as shown in figure 2, the markup has recently fallen back to a “normal” level and thus has not been much of a restraint on predicted inflation over the past year or so—a period in which the markup equations overpredict inflation by a substantial amount (figure 3). Because the overprediction episode is as yet fairly short, it is too soon to say whether the errors are transitory or evidence of some other factor restraining inflation. Another concern is the sensitivity of outcomes to the particular measure of productivity entering the markup calculation. Although our main results go through when the trend is estimated with the HP filter, the markup variable is generally insignificant when the trend is estimated with a one-sided Kalman filter or when actual productivity is used. Finally, we do not have a clear story for the rise of the markup rose in the early 1990s. The residuals of the markup equation for NFB inflation are generally positive during this period, suggesting that some of the rise in

the markup may have been a result of unexplained price increases. Low rates of increase of hourly compensation in the first half of this decade may also have contributed, but an analysis of compensation developments is beyond the scope of this paper.

6 Conclusions

The baseline version of the Phillips curve cannot adequately explain why inflation has been so low recently. While import prices, which are part of the baseline specification, are estimated in most equations to have restrained inflation over the past few years, a full attribution of inflation developments to these prices would require a many-fold increase in their coefficients. Looking beyond the baseline framework, capacity utilization predicts inflation in the recent period more accurately than does the unemployment rate, but the reverse holds for the estimation sample as a whole on an in-sample basis. Out-of-sample forecast accuracy provides no basis for discriminating between the two measures of utilization.

We find, however, that evidence favors a view in which price developments are, in part, determined by an “error-correction” mechanism involving the markup of prices over trend unit labor costs. A high markup puts downward pressure on price inflation and the reverse obtains when the markup is low. With markups relatively high through much of the 1990s, the error-correction channel is estimated to have held down inflation over much of this period. Although the reason for the rise in the markup is somewhat ambiguous, markup levels are now close to their historical averages and thus should not restrain inflation in the near future, unless, for some reason, firms are under pressure to reduce margins below historical norms.

Appendix Tables

Table 9: Baseline-87 Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Estimated coefficients: ^a						
inflation lags	1.00 $\frac{2}{1.21}$ (∞)	1.00 $\frac{2}{1.24}$ (∞)	1.00 $\frac{2}{1.24}$ (∞)	1.00 $\frac{1}{1.23}$ (∞)	1.00 $\frac{2}{1.24}$ (∞)	1.00 $\frac{1}{1.18}$ (∞)
unemployment rate	-.53 (11.1)	-.45 $\frac{0}{0.1}$ (9.0)	-.37 (7.4)	-.34 $\frac{0}{0.3}$ (5.9)	-.53 $\frac{0}{0.3}$ (7.7)	-.64 $\frac{0}{0.1}$ (9.0)
food and energy price (contemporaneous)	.88 (20.5)	.95 (15.2)			.52 (6.6)	
food and energy price (lags)			.45 $\frac{0}{1.4}$ (7.4)	.46 $\frac{0}{1.3}$ (5.4)		.44 $\frac{1}{1.1}$ (2.7)
energy price						.05 (4.3)
farm price						-.02 (6.6)
import price		.03 (2.7)			.04 $\frac{1}{0.1}$ (2.4)	.13 $\frac{1}{0.3}$ (5.7)
Equation statistics:						
sample period	55.1 - 87.4	55.1 - 87.4	63.3 - 87.4	65.3 - 87.4	60.1 - 87.4	60.1 - 87.4
$\bar{R}^2$	.9428	.9379	.8950	.8832	.9108	.9131
sum of squared residuals	71.58	66.22	54.01	43.48	73.87	98.59
regression standard error	.7537	.7278	.7662	.7282	.8428	.9831
Durbin-Watson	1.78	1.53	1.77	1.62	1.90	2.01

^aFor each regressor, the estimated coefficient, or coefficient sum, and t-statistic are reported. Entries that are coefficient sums are indicated by subscripts indexing the minimum and maximum lags, and, if polynomial restrictions are imposed, by a superscript reporting the polynomial degree. Equations also include an intercept and wage-price dummy.

Table 10: Baseline-98 Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Estimated coefficients: ^a						
inflation lags	1.00 ² _{1,24} (∞)	1.00 ² _{1,24} (∞)	1.00 ² _{1,23} (∞)	1.00 ² _{1,24} (∞)	1.00 ² _{1,15} (∞)	1.00 ² _{1,17} (∞)
unemployment rate	-.49 (11.3)	-.40 ⁰ _{0,1} (8.3)	-.35 (8.0)	-.27 ⁰ _{0,3} (5.1)	-.38 (6.8)	-.50 ⁰ _{0,1} (7.8)
food and energy price (contemporaneous)	.90 (23.1)	.91 (15.1)			.38 (5.4)	
food and energy price (lags)			.45 ⁰ _{1,4} (8.2)	.36 ⁰ _{1,3} (3.8)		.35 ⁰ _{1,4} (2.4)
energy price						.04 (5.5)
farm price						-.02 (6.5)
import price		.04 ⁰ _{0,1} (3.8)		.04 ⁰ _{0,1} (3.3)	.06 ¹ _{0,1} (4.3)	.12 ¹ _{0,2} (6.5)
Equation statistics:						
sample period	55.1 - 98.4	55.1 - 98.4	63.3 - 98.4	65.3 - 98.4	60.1 - 98.4	60.1 - 98.4
$\bar{R}^2$	.9349	.9227	.8911	.8804	.9068	.9024
sum of squared residuals	89.67	94.22	68.87	71.25	93.03	129.09
regression standard error	.7263	.7467	.7116	.7490	.7928	.9403
Durbin-Watson	1.66	1.51	1.74	1.53	1.91	2.00

^aFor each regressor, the estimated coefficient, or coefficient sum, and t-statistic are reported. Entries that are coefficient sums are indicated by subscripts indexing the minimum and maximum lags, and, if polynomial restrictions are imposed, by a superscript reporting the polynomial degree. Equations also include an intercept and wage-price dummy.

Table 11: No Smoothing Restrictions

	CPI	PCE	CPIX	PCEX	GDP	NFB
Estimated coefficients: ^a						
inflation lags	1.00 $\frac{3}{1.4}$ (∞)	1.00 $\frac{3}{1.4}$ (∞)	1.00 $\frac{4}{1.5}$ (∞)	1.00 $\frac{4}{1.5}$ (∞)	1.00 $\frac{3}{1.4}$ (∞)	1.00 $\frac{6}{1.7}$ (∞)
unemployment rate	-.21 $\frac{1}{0.1}$ (4.2)	-.17 (3.5)	-.24 (4.8)	-.15 (3.0)	-.19 $\frac{1}{0.1}$ (3.5)	-.30 $\frac{2}{0.2}$ (4.0)
food and energy price (contemporaneous)	.89 (17.5)	.95 (13.2)			.31 (3.9)	
food and energy price (lags)	-.72 $\frac{3}{1.4}$ (10.6)	-.81 $\frac{3}{1.4}$ (8.2)		.17 $\frac{1}{1.1}$ (2.5)	-.20 $\frac{1}{1.1}$ (2.5)	
energy price						.04 (4.7)
farm price						-.02 (5.1)
import price	.04 (3.6)	.05 (4.7)	.05 (5.3)	.04 (4.2)	.08 $\frac{1}{0.1}$ (5.2)	.15 $\frac{2}{0.2}$ (7.5)
Equation statistics:						
sample period	55.1 - 98.4	55.1 - 98.4	63.3 - 98.4	65.3 - 98.4	60.1 - 98.4	60.1 - 98.4
$\bar{R}^2$	.9256	.9157	.8760	.8791	.8950	.8814
sum of squared residuals	98.24	99.70	77.29	70.87	102.78	150.43
regression standard error	.7763	.7797	.7594	.7530	.8419	1.0366
Durbin-Watson	2.09	2.05	2.02	2.01	2.14	1.81

^aFor each regressor, the estimated coefficient, or coefficient sum, and t-statistic are reported. Entries that are coefficient sums are indicated by subscripts indexing the minimum and maximum lags, and, if polynomial restrictions are imposed, by a superscript reporting the polynomial degree. Equations also include an intercept and wage-price dummy.

Table 12: Capacity Utilization Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Estimated coefficients: ^a						
inflation lags	1.00 ² _{1,24} (∞)	1.00 ⁴ _{1,10} (∞)	1.00 ⁴ _{1,9} (∞)	1.00 ² _{1,24} (∞)	1.00 ² _{1,15} (∞)	1.00 ² _{1,17} (∞)
capacity utilization	.14 ⁰ _{0,1} (9.1)	.09 ⁰ _{0,3} (5.0)	.09 ⁰ _{0,3} (5.4)	.09 ⁰ _{0,3} (5.2)	.11 ⁰ _{0,1} (6.1)	.16 ⁰ _{0,3} (7.2)
food and energy price (contemporaneous)	.89 (19.4)	.87 (13.3)			.34 (4.6)	
food and energy price (lags)	-.24 ⁰ _{1,4} (3.2)	-.38 _{1,1} (4.8)	.28 ⁰ _{1,4} (4.0)	.37 ⁰ _{1,4} (3.7)		
energy price						.04 (5.6)
farm price						-.02 (6.7)
import price	.03 (2.6)	.04 (3.9)	.02 (2.4)	.04 (3.6)	.06 ¹ _{0,1} (4.1)	.14 ¹ _{0,2} (8.3)
Equation statistics:						
sample period	55.1 - 98.4	55.1 - 98.4	63.3 - 98.4	65.3 - 98.4	60.1 - 98.4	60.1 - 98.4
$\bar{R}^2$	.9306	.9249	.8957	.8792	.9020	.8958
sum of squared residuals	94.53	89.91	64.51	71.95	97.91	138.76
regression standard error	.7501	.7359	.6964	.7527	.8134	.9716
Durbin-Watson	1.60	2.02	2.09	1.54	1.86	1.95

^aFor each regressor, the estimated coefficient, or coefficient sum, and t-statistic are reported. Entries that are coefficient sums are indicated by subscripts indexing the minimum and maximum lags, and, if polynomial restrictions are imposed, by a superscript reporting the polynomial degree. Equations also include an intercept and wage-price dummy.

Table 13: Markup Equations

	CPI	PCE	CPIX	PCEX	GDP	NFB
Estimated coefficients: ^a						
inflation lags	1.00 ³ _{1,24} (∞)	1.00 ² _{1,24} (∞)	1.00 ² _{1,23} (∞)	1.00 ² _{1,24} (∞)	1.00 ² _{1,15} (∞)	1.00 ² _{1,17} (∞)
unemployment rate	-.55 ⁰ _{0,1} (11.9)	-.45 ⁰ _{0,3} (9.4)	-.36 (8.7)	-.37 ⁰ _{0,3} (7.3)	-.41 (7.3)	-.55 (9.0)
food and energy price (contemporaneous)	.88 (22.5)	.87 (15.8)			.36 (5.1)	
food and energy price (lags)			.41 ⁰ _{1,4} (7.7)	.49 ⁰ _{1,4} (5.9)		.43 ² _{1,4} (3.0)
energy price						.04 (5.2)
farm price						-.02 (6.7)
import price		.03 (3.3)			.05 ¹ _{0,1} (3.4)	.10 ¹ _{0,2} (5.3)
trend markup	-.22 ^{1,1} (4.9)	-.26 ^{1,1} (5.8)	-.17 ^{1,1} (3.7)	-.24 ^{1,1} (4.9)	-.15 ^{1,1} (2.6)	-.23 ^{1,1} (3.4)
Equation statistics:						
sample period	55.1 - 98.4	55.1 - 98.4	63.3 - 98.4	65.3 - 98.4	60.1 - 98.4	60.1 - 98.4
$\bar{R}^2$	.9429	.9344	.9005	.8903	.9104	.9134
sum of squared residuals	77.70	79.39	62.48	65.35	88.91	112.16
regression standard error	.6801	.6874	.6803	.7173	.7777	.8856
Durbin-Watson	1.82	1.69	1.88	1.63	1.95	2.05

^aFor each regressor, the estimated coefficient, or coefficient sum, and t-statistic are reported. Entries that are coefficient sums are indicated by subscripts indexing the minimum and maximum lags, and, if polynomial restrictions are imposed, by a superscript reporting the polynomial degree. Equations also include an intercept and wage-price dummy.

References

- Andrews, Donald, Inpyo Lee, and Werner Ploberger (1996) Optimal Changepoint Tests for Normal Linear Regression. *Journal of Econometrics*: 9-38.
- Fuhrer, Jeffrey C. (1995) The Phillips Curve is Alive and Well. *New England Economic Review*, March/April: 41-56.
- Gordon, Robert J. (1997) The Time-Varying NAIRU and its Implications for Economic Policy. *Journal of Economic Perspectives*, Winter: 11-32.
- Gordon, Robert J. (1998) Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU. *Brookings Papers on Economic Activity*, 2: 297-333.
- Staiger, Douglas, James H. Stock, and Mark W. Watson (1997) How Precise Are Estimates of the Natural Rate of Unemployment? *Reducing Inflation: Motivation and Strategy*, ed. Christina D. Romer and David H. Romer. University of Chicago Press.