

The Phillips Curve Is Alive and Well

The Federal Reserve Bank of Boston's 1978 economic conference was entitled "After the Phillips Curve." Many of the papers in the conference volume sounded the death knell of the Phillips curve, citing its dismal performance in the face of oil shocks in the 1970s, and its inappropriateness as a policy guide because of its presumed sensitivity to shifts in the underlying macroeconomic structure.

Seventeen years after the publication of the conference volume, rumors of the death of the Phillips curve appear to have been greatly exaggerated. In fact, the Phillips curve is alive and well, and living in a good number of (although certainly not all) widely used macroeconomic models. This paper takes the view that the primary reason for its longevity is that, in contrast to the common perception at the time of the 1978 conference, the Phillips curve has been an extremely robust empirical relationship, showing little or no sign of instability over the past 35 years. To outward appearances, at least, the Phillips curve is as structural a relationship as macroeconomists have ever had at their disposal.

The major criticism levied against the Phillips curve and many macroeconomic models of the 1960s and 1970s was that they were not truly structural. That is, they captured empirical regularities between aggregate variables like the unemployment rate and the rate of inflation, but they did not take account of all the interactions in the underlying or structural behavior of consumers and firms in the economy. The risk to this approach, as articulated by Lucas and Sargent (1978), was that even if the underlying structural behavior of economic agents remained stable, the measured relationships among aggregate variables could easily shift as conditions changed in the linkages not incorporated into the aggregate relationships. A fuller description of this criticism (the "Lucas critique") and a simple example are described in Section III below.¹

Ultimately, however, as Lucas and Sargent (1978) emphasized, "the question of whether a particular model is structural is an empirical, not a theoretical, one." Thus, while the theoretical force of Lucas and Sargent's

Jeffrey C. Fuhrer

Assistant Vice President and Economist, Federal Reserve Bank of Boston. The author thanks Lynn Browne, Mark Hooker, Richard Kopcke, Stephen McNees, and Geoffrey Tootell for helpful comments, and Alicia Sasser for excellent research assistance.

arguments still holds—macroeconomic models not based on underlying structure are subject to shifts—it is the empirical force of their point that is of interest for macroeconomists and policymakers. This paper examines an array of empirical evidence bearing on the stability of the Phillips curve, and calls into question the empirical force of the Lucas critique as it applies to the Phillips curve.

I. A Brief History of the Phillips Curve

The essence of the modern Phillips curve is that the rate of change of nominal wages depends upon the expected rate of change of the overall price level—workers want their wages to keep pace with inflation, all else equal—and on the level of the unemployment rate relative to its natural rate. The dependence of

The Phillips curve is alive and well, and living in a good number of widely used macroeconometric models.

wage inflation on expected inflation (as well as unemployment) is the difference between the “expectations-augmented” Phillips curve, first implemented empirically in the 1960s, and the original Phillips curve.² The natural rate of unemployment, or NAIRU (non-accelerating-inflation rate of unemployment), is defined by the behavior of inflation: It is the rate of unemployment that exerts neither downward nor upward pressure on wage inflation, given expectations of price inflation.³

Many current versions of the Phillips curve cast it as the link between the rate of inflation in overall prices and the level of unemployment. In fact, this portrayal is a simplification of an underlying relationship between the rate of change of wages and unemployment, and among the levels of wages, prices, and productivity. The first link in this chain is the original Phillips curve; the second link relates prices to unit labor costs, defined as the difference between wages and productivity.⁴

The strong empirical bond between the rate of change of wages and the level of unemployment was

first documented in Phillips (1958) for United Kingdom data from 1861 to 1957.⁵ Although subsequent authors (notably Robert J. Gordon) built upon the work of Phillips, many of their significant alterations to the original Phillips curve were anticipated by Phillips.

For example, Phillips postulated that the rate of change in wages could depend on the (expected) rate of change of retail prices, “operating through cost of living adjustments” (p. 283). This notion is the grandfather of the “expectations-augmented” Phillips curves of today, as outlined above. In addition, Phillips argued that the rate of change of unemployment, as well as its level, could be an important determinant of wages. These “speed limit” effects are motivated by the possibility that inflation responds more when unemployment is changing rapidly than when it changes gradually, holding constant the level effect. Phillips also recognized the importance of the effect of prices of imported goods on the overall determination of wages (p. 284). Finally, the original Phillips curve recognized the possibility of requiring larger and larger increments of unemployment to reduce inflation as inflation approached zero.⁶ Thus, while his statistical work, undertaken before the advent of modern computing technology, focused solely on the correlation between wage inflation and unemployment,

¹ The essence of Lucas’s critique was articulated earlier by Haavelmo (1944), in the context of interpreting multi-equation econometric models, and Duesenberry (1948), in the discussion of his consumption model. Lucas (1976) is the first development of his version of the critique.

² Perry (1966) includes lags of the consumer price index as adjustments for the cost of living in his wage-price Phillips curve. In his view, wages adjust to *past* changes in the cost of living. Gordon (1970) explicitly includes expected changes in the overall price level as a determinant of wage changes. He uses a distributed lag of past price changes as a proxy for expected price changes.

³ Equivalently, the NAIRU is the rate of unemployment at which the rate of change of nominal wages equals the expected change in the overall price level. For more detail on specification and interpretation of the Phillips curve and the NAIRU, see Tootell (1994).

⁴ If the unemployment rate proxies well for the tightness in markets for all factors of production, then the Phillips correlation may reflect the underlying correlation between general price pressures and inflation.

⁵ Klein and Goldberger’s model (1955) includes an equation that makes the change in wages a function of the unemployment rate and the lagged change in the price level. However, they are not commonly cited as the discoverers of the inflation/unemployment relationship.

⁶ Phillips’ original curve is estimated as $\log(a + \Delta w_t) = \log(b) + c \log(U_t) + \epsilon_t$, where the constant a is included in order to ensure non-negativity of the left-hand side of the equation. This form implies that as wage inflation falls below some level (typically as it approaches zero), the proportionate effect of unemployment decreases.

his theoretical work suggested many of the improvements made to the empirical Phillips curves in the last 40 years.

II. Strengths and Weaknesses of the Phillips Curve

The great strength of the Phillips curve is that it captures an economically important and statistically reliable empirical relationship between inflation and unemployment. Figure 1 portrays the essence of this relationship. The figure plots the change in the inflation rate (three standard measures are displayed in the solid lines) against the difference between the civilian unemployment rate and an assumed NAIRU of 6 percent. When unemployment falls below about 6 percent, inflation in wages and prices tends to rise;

Perhaps the greatest weakness of the Phillips curve is its lack of theoretical underpinnings: No one has derived a Phillips curve from first principles, beginning with the fundamental concerns and constraints of consumers and firms.

when unemployment rises above 6 percent, inflation tends to fall. After publication of Phillips's seminal paper, the presence of this correlation was documented by many researchers for many countries.

The figure focuses on the post-1979 period for two reasons. First, adjusting for the effects of oil price shocks that mask the underlying correlation is not as important in this period as in the 1970s.⁷ Second, the academic literature of the 1970s proclaimed the death of the Phillips curve; finding that the underlying correlation appears robust from that time to today sheds some light on the empirical relevance of the theoretical objections to the Phillips curve. Section IV examines in more detail the stability of the Phillips curve.

Perhaps the greatest weakness of the Phillips curve is its lack of theoretical underpinnings: No one

has derived a Phillips curve from first principles, beginning with the fundamental concerns and constraints of consumers and firms. Few models that articulate the supply of and demand for labor imply a simple aggregate relationship that looks at all like a Phillips curve. This is not to say that the empirical relationship makes no sense. Labor costs account for about two-thirds of the total cost of producing output, so that pressures in the labor markets should strongly influence changes in wages and prices. Still, some feel that this lack of rigorous theoretical foundations is a fatal flaw; many find this deficiency less life-threatening. The next section discusses a specific implication of this asserted deficiency of structure, widely referred to now as "the Lucas Critique."

III. The Lucas Critique

Robert E. Lucas, Jr. (1976) criticized the use of econometric models, such as the Phillips curves that were estimated at that time, for policy evaluation. In principle his critique applies to the Phillips curves presented in this paper. The best way to understand his critique is through an example, and for our purposes, the best example is the Phillips curve.

Recall that the Phillips curve assumes that workers set wage changes so as not to lose ground relative to expected changes in the cost of living. Most econometric models during the late 1970s modeled expected changes in the cost of living, π_t^e , as a distributed lag of past changes in the cost of living, or

$$\pi_t^e = \sum_{i=1}^k \alpha_i \pi_{t-i}. \quad (1)$$

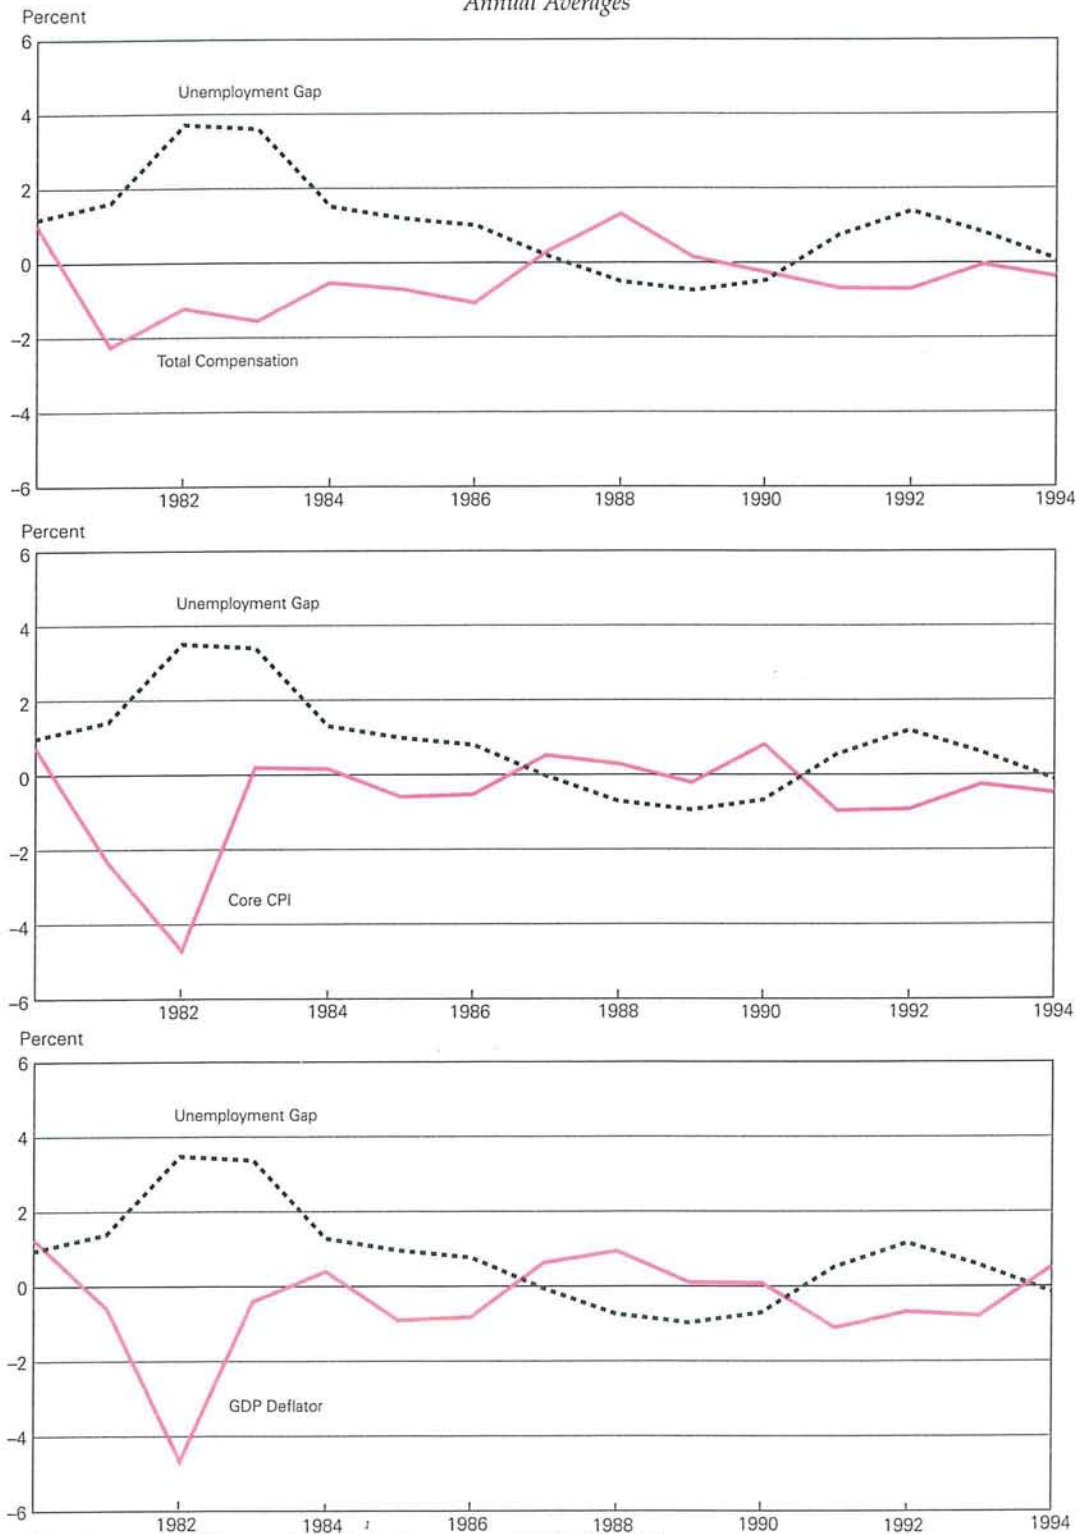
This characterization of expected inflation is implemented empirically by estimating the coefficients α_i on lagged inflation from historical data.⁸ Such a simple

⁷ The significant drop in oil prices in the mid-1980s appears to be of relatively small importance for overall inflation. However, during the two major oil price shocks of the 1970s, inflation increased at the same time that unemployment increased. This does not negate the underlying Phillips curve, it simply makes clear that changes in important relative prices can adversely affect both inflation and unemployment. Thus, a careful empirical implementation of the Phillips curve must at the very least control for the effects of rapid oil price increases. The regressions presented in section IV capture the underlying inflation-unemployment correlation, controlling for changes in oil prices.

⁸ The sum of the coefficients on lagged inflation is normally constrained to one, so that in the long run, inflation equals expected inflation.

Figure 1

Phillips Curve Correlation in the 1980s *Change in Inflation vs. Unemployment Gap* *Annual Averages*



Unemployment gap is civilian unemployment rate minus an assumed NAIRU of 6 percent.
The three measures of inflation are shown as the change in the annual average of the four-quarter rate of change.

Sources: Civilian unemployment rate: U.S. Bureau of Labor Statistics.
Total compensation: U.S. Bureau of Labor Statistics.
Before 1982: Nonfarm business sector, compensation per hour.
1982-present: Employment cost index, compensation, civilian workers.
Core CPI: consumer price index, all items excluding food and energy: U.S. Bureau Labor Statistics.
Gross domestic product deflator: U.S. Bureau of Economic Analysis.

description of inflation expectations is assumed to provide a reasonable forecast of inflation in the short run, particularly if monetary policy and other important influences on inflation have remained relatively stable.

There are two theoretical objections to using this description of expectations formation for policy questions. First, equation (1) incorporates no knowledge of the effect of unemployment on the rate of inflation, even though the premise of the Phillips curve is that unemployment is critical in determining the direction of the inflation rate. Second, the estimated equation implies a particular path for inflation, depending on where it has been in the past, that does not depend on how vigorously or sluggishly monetary policy is pursuing its inflation target (if any).

Lucas argues that these omissions in modeling inflation expectations make econometric equations, such as the Phillips curve, not useful for evaluating alternative policies. The reasoning is essentially as follows: Suppose that in period I monetary policy has an inflation target of 5 percent, and that it moves only very gradually to achieve that target (perhaps because policymakers do not wish to bear the cost of high unemployment entailed in a more rapid move toward the inflation target). A distributed lag model of inflation estimated on period I data will build in the slow return of inflation to its target, reflected in the pattern of the coefficients on lagged inflation.⁹

Now consider period II, in which monetary policy has the same inflation target, but in which it moves quickly and vigorously to achieve the target. A distributed lag model of inflation will build in the vigor of (presumably effective) monetary policy, shortening the lag responses to reflect the quicker return of inflation to its target. Thus, the appropriate lag pattern in forming expected inflation will differ for different monetary policy regimes.¹⁰

Suppose that monetary policymakers who live in period I entertain changing monetary policy to the more vigorous policy of (as yet unexperienced) period II. The policymakers want to know how inflation will respond under the new regime, so they pull out the Phillips curve (estimated under period I data), alter the forecast of the unemployment rate based on the more vigorous policy assumption, and produce a forecast of inflation.

Will the inflation prediction so obtained be useful for policymakers? This is the crux of the Lucas critique. If wage and price setters know about the new monetary policy behavior and understand its implications for expected cost-of-living changes (and its im-

plications for the unemployment rate), then presumably their expectations of cost-of-living changes will change under the new regime, and the old Phillips curve model will not be a good predictor of inflation. Lucas's insight in the context of the Phillips curve is that *price setters' expectations of future events may depend upon the behavior of other agents in the economy, particularly upon those who set monetary policy*. Thus, as monetary policymakers consider changing policy, they need to use econometric models that take into account the possible change in expectations that this will induce in other actors in the economy, in this case wage and price setters.

This insight certainly holds water theoretically, as long as the assumptions hold about how expectations are formed, but it need not hold water empirically. We cannot know *a priori* how agents form their expectations. In particular, we do not know whether they adjust their expectations to changes in monetary policy behavior, as Lucas's theoretical objection suggests they might. Even if agents' expectations react to changes in monetary policy, we cannot know *a priori* whether the historical changes in monetary policy have been large enough to cause empirically significant shifts in expectations in the Phillips curve. Evidence bearing on these empirical issues is presented in the following sections.

IV. An Empirical Assessment of the Lucas Critique for the Phillips Curve

Lucas's critique suggests that the Phillips curve will not be stable over long periods of time, particularly as the behavior of the monetary authority changes. This hypothesis is testable, and this section provides some straightforward tests for instability in the Phillips curve, with concern for both the statistical and the economic significance of any changes in the relationship over time. We begin with the simplest and least stringent, and progress to somewhat more complex and stringent tests.

The first test is to see whether a Phillips curve estimated on data over the last 30 years shows signs of going off track. This is an important first step in looking for signs of instability, but it is the least

⁹ The expectations component of an estimated Phillips curve is examined in more detail in section IV.

¹⁰ A fully internally consistent model of expectations would adjust the lag coefficients on past inflation *and* build in the expected effect of swings in the unemployment rate, as suggested by the Phillips curve.

stringent test because it uses the information that was used to estimate the equation to gauge its success within the estimation sample. More stringent tests that gauge the performance of the equation outside its estimation sample are presented below.

The basic specification used here is a "price-price" Phillips curve that subsumes the "wage-price" Phillips curve and the unit labor cost equations into one. Current inflation in the core CPI—the CPI excluding its food and energy components— Δp_t , depends on lagged inflation (with coefficients α_i constrained to sum to 1), two lags of the unemployment rate, U_t , and the rate of change of oil prices, $\Delta p_{o,t}$:¹¹

$$\Delta p_t = \sum_{i=1}^{12} \alpha_i \Delta p_{t-i} + \sum_{j=1}^2 \beta_j U_{t-j} + \gamma \Delta p_{o,t} + \epsilon_t. \quad (2)$$

While results are reported for the "price-price" version of the Phillips curve using the core CPI, parallel results have been computed for the "wage-price" version of the Phillips curve, and for "price-price" versions using the overall CPI, and for the GDP deflator.¹² All of these versions give qualitatively similar results (see Tables 1–1c for the estimated equations). The results for the core CPI are presented because it is perhaps the most widely monitored measure of the "core" or "trend" rate of inflation. Important quantitative differences for different price measures will be noted throughout.

We expect the sum of the coefficients on unemployment to be negative and the coefficient on the change in oil prices to be positive.¹³ The estimated

equation is presented in Table 1. The estimated overall response to the unemployment rate, at -0.28 , is quite similar to those estimated in similar studies, for example Gordon (1994). The estimated NAIRU, at 6.1 percent, is completely in line with conventional wisdom, and the uncertainty surrounding the estimate (its standard error) easily admits NAIRUs of 5.5 to 6.5 percent.¹⁴

Note that the coefficients on the two lags of the unemployment rate may be interpreted as a "level of unemployment" effect and a "change in unemployment" or "speed limit" effect.¹⁵ If the unemployment contribution to the Phillips curve is written as $aU_{t-1} + bU_{t-2}$, it is straightforward to show that this is equivalent to a level and a change effect, $(a + b)U_{t-1} - b\Delta U_{t-1}$. Thus, the significance of the second lag coefficient determines the significance of the "speed limit" effect. In the estimates presented in Table 1, the level effect is displayed as the sum of the unemployment rate coefficients. The implied change effect is -1.8 . For a given level of unemployment, the speed limit coefficient says that every one-tenth percentage point increase in the unemployment rate drops the inflation rate by almost two-tenths percentage point. Note that, as shown in Table 1c, the compensation Phillips curve exhibits no rate-of-change effect; the estimated coefficient on a second lag of unemployment is small and not significantly different from zero, and is therefore omitted from the specification.

from zero (standard errors are corrected for heteroskedasticity and serial correlation). For the purpose of this paper, the exclusion of the instrumented value of contemporaneous unemployment is unlikely to change the fit or behavior of the estimated equation.

¹⁴ The NAIRU standard error is the standard error of a function of several estimated parameters: the negative of the constant (not reported) divided by the sum of the coefficients on lagged unemployment. Its standard error is computed using an asymptotic approximation:

$$SE(\text{NAIRU}) = \sqrt{\frac{\partial f(\beta)}{\partial \beta} \Omega \frac{\partial f(\beta)'}{\partial \beta}}$$

where $f(\beta)$ is the function that translates the underlying estimated parameters (constant and unemployment rate coefficients) into the NAIRU, and Ω is the estimated variance covariance matrix of those parameters. The NAIRU for the wage-price Phillips curve is not reported, as it varies over time, depending on the average level of productivity growth in the current and preceding seven quarters.

¹⁵ The speed limit effect gets its name from the suggestion that more rapid changes in the unemployment rate may cause larger changes in the inflation rate for a given level of the unemployment rate. The intuition is most clear for rapid decreases in the unemployment rate, when it is argued that rapidly growing demand for labor might put greater pressure on wages (due to bottlenecks, for example) than a gradual increase in labor demand and decline in the unemployment rate.

¹¹ An exchange rate term, meant to reflect the influence of the exchange rate on imported goods prices on the overall domestic price level, did not enter significantly. Separate terms that measure the effect of changes in the trade-weighted average of foreign CPIs or in the real exchange rate, which combines the ratio of domestic to foreign CPIs with the nominal exchange rate, were also insignificant. Similarly, terms that might measure the effect of productivity growth through the unit labor cost relationship were found to be insignificant.

¹² The wage measure used in these unreported regressions splices total compensation for nonfarm business with the employment compensation index (ECI) measure of compensation beginning in 1982.

¹³ Including the contemporaneous unemployment rate with the two lagged rates yields an estimate of the contemporaneous coefficient that is insignificantly different from zero. The possibility of simultaneity bias for the contemporaneous coefficient suggests that the coefficient should be instrumented. However, likely candidates for instruments would include lags of unemployment and inflation that are already in the regression model. An instrumental variables estimate that uses lags of unemployment, inflation, the federal funds rate, and federal government expenditures, yields an estimate of the contemporaneous coefficient that is insignificantly different

Table 1
*Phillips Curve Estimates: Quarterly
 Inflation Rate, CPI excl. Food and Energy*
 Seasonally Adjusted (SA)

Variable	Estimated Coefficient	Standard Error	T-Statistic
Δp_{t-1}	.31	.090	3.4
Δp_{t-2}	.27	.095	2.9
Δp_{t-3}	.34	.097	3.5
Δp_{t-4}	-.05	.096	-5
Δp_{t-5}	.03	.095	.3
Δp_{t-6}	.14	.095	1.5
Δp_{t-7}	-.01	.095	-1
Δp_{t-8}	-.17	.095	-1.8
Δp_{t-9}	.12	.095	1.3
Δp_{t-10}	.04	.094	.4
Δp_{t-11}	.04	.092	.4
Δp_{t-12}	-.07	.084	-.8
U_{t-1}	-2.08	.499	-4.2
U_{t-2}	1.81	.518	3.5
Sum of U Coefficients	-.28	.121	-2.3
$\Delta \rho_t$.01	.006	1.9
NAIRU	6.09	.557	10.9

Standard Error of Regression: 1.762
 Estimation Range: 1960:II to 1993:IV

Table 1a
*Phillips Curve Estimates: Quarterly
 Inflation Rate, CPI All Items (SA)*

Variable	Estimated Coefficient	Standard Error	T-Statistic
Δp_{t-1}	.17	.088	2.0
Δp_{t-2}	.16	.083	1.9
Δp_{t-3}	.39	.080	4.9
Δp_{t-4}	.02	.086	.2
Δp_{t-5}	.24	.086	2.8
Δp_{t-6}	.03	.083	.4
Δp_{t-7}	-.03	.081	-.3
Δp_{t-8}	-.19	.079	-2.4
Δp_{t-9}	.07	.079	.9
Δp_{t-10}	-.04	.076	-.5
Δp_{t-11}	.12	.074	1.6
Δp_{t-12}	.06	.070	.8
U_{t-1}	-3.93	.533	-7.4
U_{t-2}	5.91	1.056	5.6
U_{t-3}	-3.63	1.135	-3.2
U_{t-4}	1.30	.608	2.1
Sum of U Coefficients	-.34	.119	-2.9
$\Delta \rho_t$	-.02	.011	-2.3
Δe_{t-1}	.03	.006	5.7
NAIRU	5.81	.418	13.9

Standard Error of Regression: 1.56
 Estimation Range: 1960:II to 1993:IV

Table 1b
*Phillips Curve Estimates: Quarterly
 Inflation Rate, GDP Deflator (SA)*

Variable	Estimated Coefficient	Standard Error	T-Statistic
Δp_{t-1}	.30	.091	3.3
Δp_{t-2}	.26	.093	2.7
Δp_{t-3}	.19	.095	2.0
Δp_{t-4}	.16	.097	1.7
Δp_{t-5}	-.09	.097	-1.0
Δp_{t-6}	-.12	.097	-1.2
Δp_{t-7}	.08	.097	.8
Δp_{t-8}	-.02	.096	-.2
Δp_{t-9}	.03	.095	.4
Δp_{t-10}	.16	.093	1.7
Δp_{t-11}	.09	.088	1.0
Δp_{t-12}	-.03	.083	-.3
U_{t-1}	-.38	.102	-3.7
NAIRU	6.12	.329	18.6

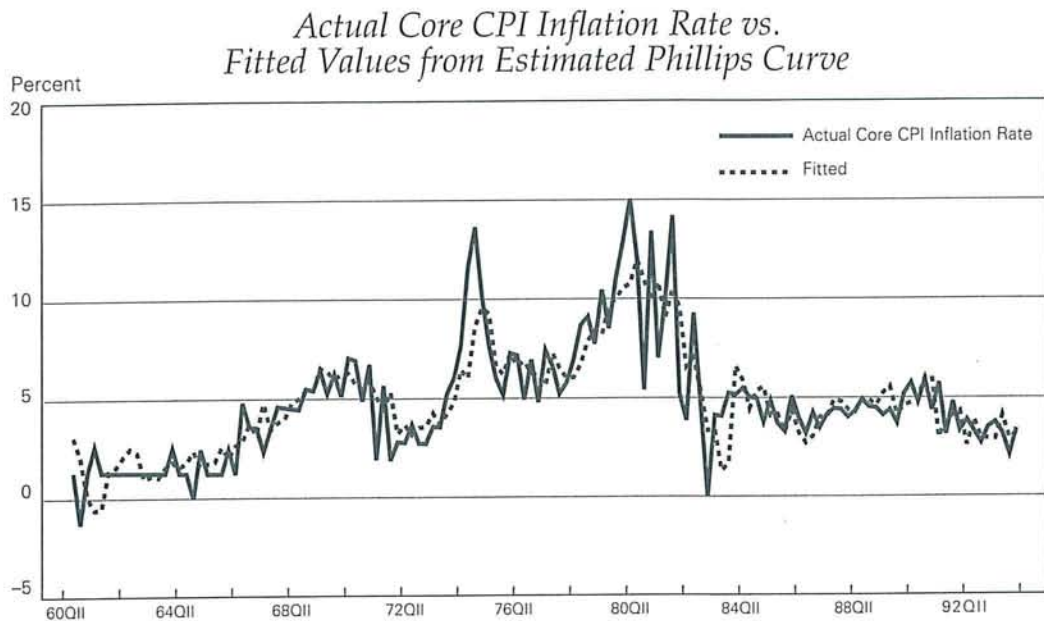
Standard Error of Regression: 1.44
 Estimation Range: 1960:II to 1993:IV

Table 1c
Wage-Price Phillips Curve Estimates

Variable	Estimated Coefficient	Standard Error	T-Statistic
Wage Inflation Equation			
Dependent variable: Quarterly Inflation Rate, ECI (SA)			
Δp_{t-1}	.17	.077	2.3
Δp_{t-2}	.17	.082	2.1
Δp_{t-3}	.09	.083	1.1
Δp_{t-4}	.15	.087	1.8
Δp_{t-5}	-.09	.084	-1.1
Δp_{t-6}	.12	.083	1.5
Δp_{t-7}	-.08	.083	-.9
Δp_{t-8}	.06	.083	.8
Δp_{t-9}	.07	.084	.9
Δp_{t-10}	.08	.079	1.0
Δp_{t-11}	.12	.078	1.6
Δp_{t-12}	.12	.071	1.7
U_{t-1}	-.86	.116	-7.3
Nixon/off	-5.05	1.422	-3.6
Constant	6.47	.736	8.8
Unit Labor Cost Equation			
Dependent variable: Quarterly Inflation Rate, Total CPI			
Δw_t	.46	.131	3.5
Δw_{t-1}	.40	.133	3.0
Δw_{t-2}	.14	.138	1.0
Δw_{t-3}	.11	.138	.8
Δw_{t-4}	.03	.137	.3
Δw_{t-5}	.00	.137	.0
Δw_{t-6}	-.05	.131	-.4
Δw_{t-7}	-.11	.129	-.8
Δ Productivity	-1.00	.000	.0
Constant	.34	.200	1.7

Standard Error of Wage Inflation Equation: 1.60
 Standard Error of Unit Labor Cost Regression: 2.23
 Estimation Range: 1963:I to 1993:IV

Figure 2



Source: Consumer price index, all items excluding food and energy, seasonally adjusted, U.S. Bureau of Labor Statistics. Fitted values are author's calculations.

Figure 2 displays the actual data for the core CPI inflation rate and the fitted values from the estimated regression equation. As the figure indicates, there is no sign within the estimation sample that the Phillips curve has run amok. In fact, its within-sample prediction errors have decreased in the last ten years.

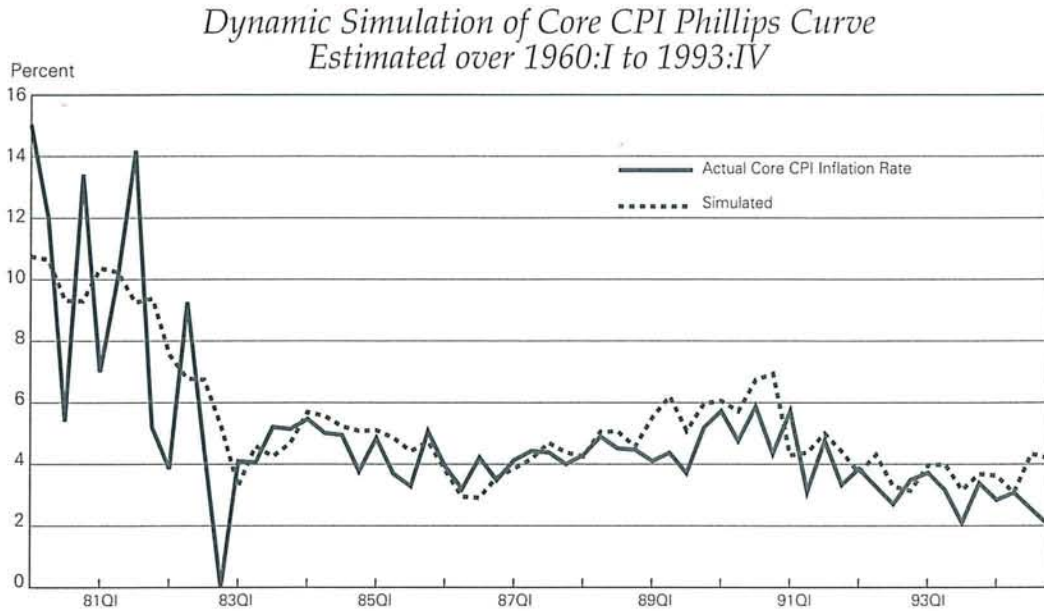
Another test of the specification involves simulating the Phillips curve over a fairly long period, feeding in the simulated values for the current period as lagged values in subsequent periods. This "dynamic" simulation is designed to reveal the multi-period forecast performance of the equation. The in-sample fit test is a one-period-ahead forecast test, and as such it allows the equation to "get back on track" by feeding in actual lagged observations for inflation that the equation did not predict. The dynamic simulation does not give the equation this information, so that one large error in predicting inflation can feed into all subsequent predictions, and the equation will wander significantly off track if it is not well specified.

Figure 3 shows the results of the dynamic simulation described above. Using the coefficients esti-

mated over the 1960:I to 1993:IV sample, the simulation begins with actual lagged values for all variables in 1980:I, and then proceeds for the next 15 years. As the figure indicates, the Phillips curve does not wander off even without referring to an actual inflation rate for 15 years, in large part because it relies on the robust correlation between inflation and unemployment.

Two considerably more stringent tests involve out-of-sample (outside the estimation sample) simulations. In the first, the Phillips curve is estimated from 1960:II to 1979:IV. The estimates are not presented here, but the differences between this estimation and the results shown in Table 1 are few: The estimated NAIRU is 5.3 (not statistically significantly different from the estimate in Table 1). The estimated sum of the unemployment coefficients is still -0.28 . These results in themselves indicate that the Phillips curve is quite stable across fairly long stretches of time. However, because the estimated NAIRU is different (likely due to changes in demographics) and the lag patterns in the expected inflation component of the curve differ

Figure 3



Source: Consumer price index, all items excluding food and energy, seasonally adjusted, U.S. Bureau of Labor Statistics. Fitted values are author's calculations.

somewhat between the two estimates, it is interesting to see if the 1960s and '70s Phillips curve can accurately predict the 1980s and '90s inflation outcomes.

Figure 4 shows the "fitted values"—the one-quarter-ahead forecasts of inflation, feeding actual lagged inflation into the equation—for 1980 to the present, using the coefficients from the Phillips curve estimated up through 1979. As the figure indicates, the inflation forecasts made using coefficients estimated from the earlier period are quite accurate, and show no significant sign of consistent under- or overprediction.

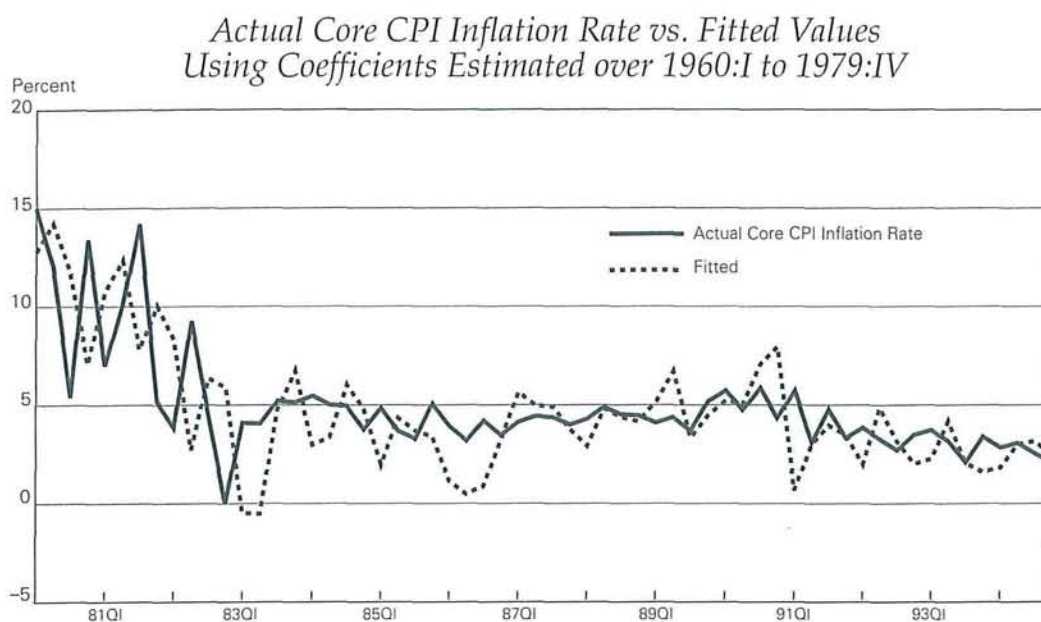
Figure 5 shows the results of a dynamic simulation of the same estimated Phillips curve outside its sample. Using coefficients estimated before 1980, and giving the equation no information about actual inflation during the period, the Phillips curve tracks the 1980s disinflation extremely well. In a sense, the simulated values from the Phillips curve reveal the underlying trend in inflation better than the measured inflation rate. The out-of-sample tests of Figures 4 and 5 would be difficult for any macroeconomic relation-

ship to pass, and overall, the Phillips curve does quite well.

A final out-of-sample simulation test pits the Phillips curve estimated from 1960:I to 1987:IV against the data in a dynamic simulation from 1988 to 1994. Again, the estimates for equation (2) themselves convey an overall impression of stability of the Phillips curve. The estimated sum of unemployment rate coefficients is -0.28 , and the estimated NAIRU is 5.9. Both are insignificantly different from the estimates displayed in Table 1. Figure 6 displays the dynamically simulated values of inflation versus the actuals. The equation does a very good job of capturing the general contours of inflation over this period. It overpredicts a bit in periods in which oil prices swung wildly, but this does not put the predictions permanently off track, and the equation completely regains its composure after 1990.

Overall, then, conventional tests of the stability of the Phillips curve indicate remarkable stability. There may be no other macroeconomic relationship that could perform as well by these criteria.

Figure 4



A More Direct Test of the Lucas Critique

The stability of the expectations component of the Phillips curve can also be tested directly. This, after all, was the main point of contention of the Lucas critique. To do so, a variety of statistical and graphical tests are performed to see if economically or statistically significant shifts have occurred in the expectations component of the Phillips curve. To be specific, we test whether there have been important changes in the coefficients on the lagged inflation terms in equation (2).

Choosing a breakpoint at which the coefficients may have shifted is difficult. Here, we focus on October 1979, when the Fed changed its operating procedures and, according to many accounts, began in earnest its disinflation program. Formally, we test to see if we can reject the hypothesis that the estimated coefficients on lagged inflation—the α_i s—are the same before a breakpoint as they are after a breakpoint.¹⁶


Table 2 displays the results of several tests of coefficient stability for a variety of breakpoints centered around 1980. The table presents results both for

the core CPI inflation equation and for the total compensation equation, which more closely mirrors the spirit of the original Phillips curve. The values in the table (column pairs 1, 2, and 3) represent the probability that the estimated shift coefficients for inflation, unemployment, or both differ from zero due to chance. A low value in these columns indicates that we can be fairly sure that a shift occurred at the indicated breakpoint. A high value indicates that the estimated shift coefficients may differ from zero only

¹⁶ We accomplish this test by estimating equation (2), adding an interactive term that multiplies a dummy, which takes the value one after the breakpoint and zero before, with the inflation rate. The sum of the coefficients on the lagged interactive terms must be constrained to zero, in order to preserve the overall constraint that the sum of the lagged inflation terms equals 1. Note that the breakpoint is chosen in reference to a monetary policy breakpoint. It can be argued that important changes were made in fiscal policy as well during the 1980s, although those breakpoints would not necessarily correspond to the monetary policy breakpoints. As the results presented below show, however, there is little evidence to indicate a breakdown of the Phillips curve at any point in the sample.

Figure 5

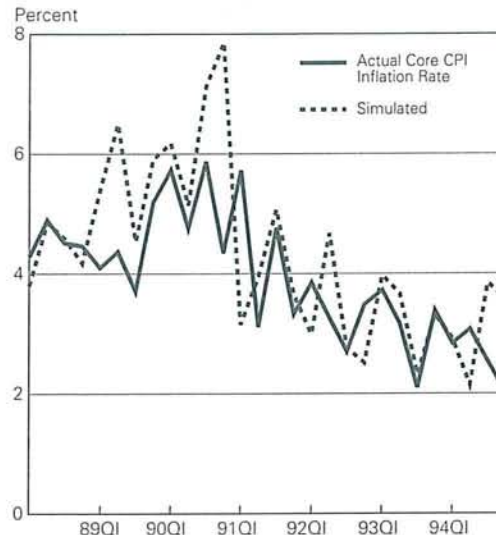
Dynamic Simulation of Core CPI Phillips Curve Estimated before 1980



Source: Consumer price index, all items excluding food and energy, seasonally adjusted, U.S. Bureau of Labor Statistics. Fitted values are author's calculations.

Figure 6

Dynamic Simulation of Core CPI Phillips Curve Estimated over 1960:I to 1987:IV



Source: Consumer price index, all items excluding food and energy, seasonally adjusted, U.S. Bureau of Labor Statistics. Fitted values are author's calculations.

due to chance. Thus, the values of 0.68 to 0.99 in the right-hand column under (1) show that we find no evidence of a shift in the expected inflation coefficients for the total compensation Phillips curve. The values of 0.01 and 0.03 in the left-hand column under (1) strongly suggest a shift in the expected inflation coefficients for the core CPI at those dates.¹⁷

As indicated in the pair of columns labeled (1) in the table, for most breakpoints around the 1980 shift in monetary policy, no compelling evidence (probabilities well below 0.1) can be found to indicate a shift in the lagged inflation (or expected cost of living increase) component of the Phillips curve. The exception arises for the core CPI when the breakpoint is chosen between 1980 and 1982. The pattern of test values suggests that if a shift occurred, it was isolated to a short period beginning around 1980 and probably ended by 1983. The final two rows of columns (1) present results of a test of the "transience" of the shift in lag coefficients. The test regression allows the first

four lagged inflation coefficients to shift in 1980:I, and for all the lagged coefficients to shift again in 1983:I. As the second-to-last row of columns (1) indicates, allowing the lags to shift in the 1980–82 period and thereafter strengthens the certainty with which we could accept the hypothesis of a shift in the lag coefficients for the core CPI. The final row in columns (1) tests for the equality of the lagged inflation coefficients in the pre-1980 and the post-1982 periods. We cannot reject the hypothesis that they were equal for either variable, strengthening the suspicion that if any shifting occurred, it was isolated to the 1980–82 period. Was the shift detected for the core CPI evidence of a three-year temporary change in expectations due to a temporary policy change, or is it simply evidence of overfitting a particular pattern of inflation outcomes?¹⁸

¹⁷ The values reported in the table are the fraction of the area to the right of the test statistic for the distribution constructed under the null hypothesis that the coefficients are equal before and after the breakpoint.

¹⁸ The results for the total CPI and for the GDP deflator provide, overall, much weaker evidence of shifts. For the total CPI, the obvious breakpoint for the lagged inflation coefficients also occurs in 1980:I, but it is not nearly as significant statistically as for the core CPI. The GDP deflator never develops a probability below 0.4 for the same test. Estimates of the baseline Phillips curves for the total CPI and the GDP deflator, as well as an estimated wage-price Phillips curve with a unit labor cost equation, are presented in Tables 1a–c.

Table 2
Tests for Shifts in Estimated Coefficients

Breakpoint(s)	Probability that estimated shift coefficients differ from zero due to chance							
	Joint test of all coefficients on:							
	(1)		(2)		(3)		(4)	
	Lagged inflation		Lagged unemployment		Both		Shift in sum of unemployment coefficients	
	Core CPI	Total Comp.	Core CPI	Total Comp.	Core CPI	Total Comp.	Core CPI	Total Comp.
1978:1	.08	.94	.64	.69	.15	.96	.04	-.01
1979:1	.09	.68	.68	.74	.17	.78	.02	-.02
1980:1	.00	.95	.93	.28	.00	.97	-.02	-.04
1981:1	.01	.99	.69	.49	.02	.99	.02	-.05
1982:1	.03	.99	.11	.77	.02	.99	.03	-.02
1983:1	.69	.98	.04	.90	.47	.99	.02	-.01
1984:1	.99	.95	.85	.78	.98	.94	-.02	-.02
1980:1, 1983:1	.00	.94						
Pre-1980 = Post-1982	.88	.79						

The dependent variable is the log change in the CPI excluding food and energy or the log change in total compensation, as defined in the text. The baseline equations are the estimated equations from Tables 1 and 1c, respectively.

To answer this question we examine the expected inflation components of the estimated Phillips curves during the potential shift periods. Figure 7 presents the Phillips curve's estimates of expected inflation, allowing for a possible shift in 1980 (the solid red line) and holding the coefficients constant for the entire sample (the dashed red line), for the core CPI and for total compensation. This figure shows graphically what the test results in Table 2 suggest. Core inflation from 1980 to 1982 gyrated wildly between 0 and 15 percent. Allowing the coefficients on lagged inflation to shift in 1980 allows the equation to better mimic this unusual pattern in inflation, and thus improves the in-sample fit of the Phillips curve. In addition, using this lag pattern does not worsen inflation predictions after 1982, when inflation moved more smoothly. Many would be reluctant to call this a true shift in the expectations component of the Phillips curve, especially when there is no evidence of any such shift in the compensation Phillips curve. Instead, the estimated shifts in the lagged inflation coefficients amount to an over-fitting of an unlikely-to-be-repeated pattern of inflation in the three-year span from 1980 to 1982.

The bottom panel of the figure shows the estimated change in the expectations component for the total compensation Phillips curve. Consistent with the results in Table 2, there is essentially no evidence of any shift in the expectations behavior at any year

around the proposed shift time of 1980. Because the underlying motivation of the Phillips curve makes most sense when applied to compensation data, this panel seems the most damning for the empirical significance of the Lucas critique for the Phillips curve.

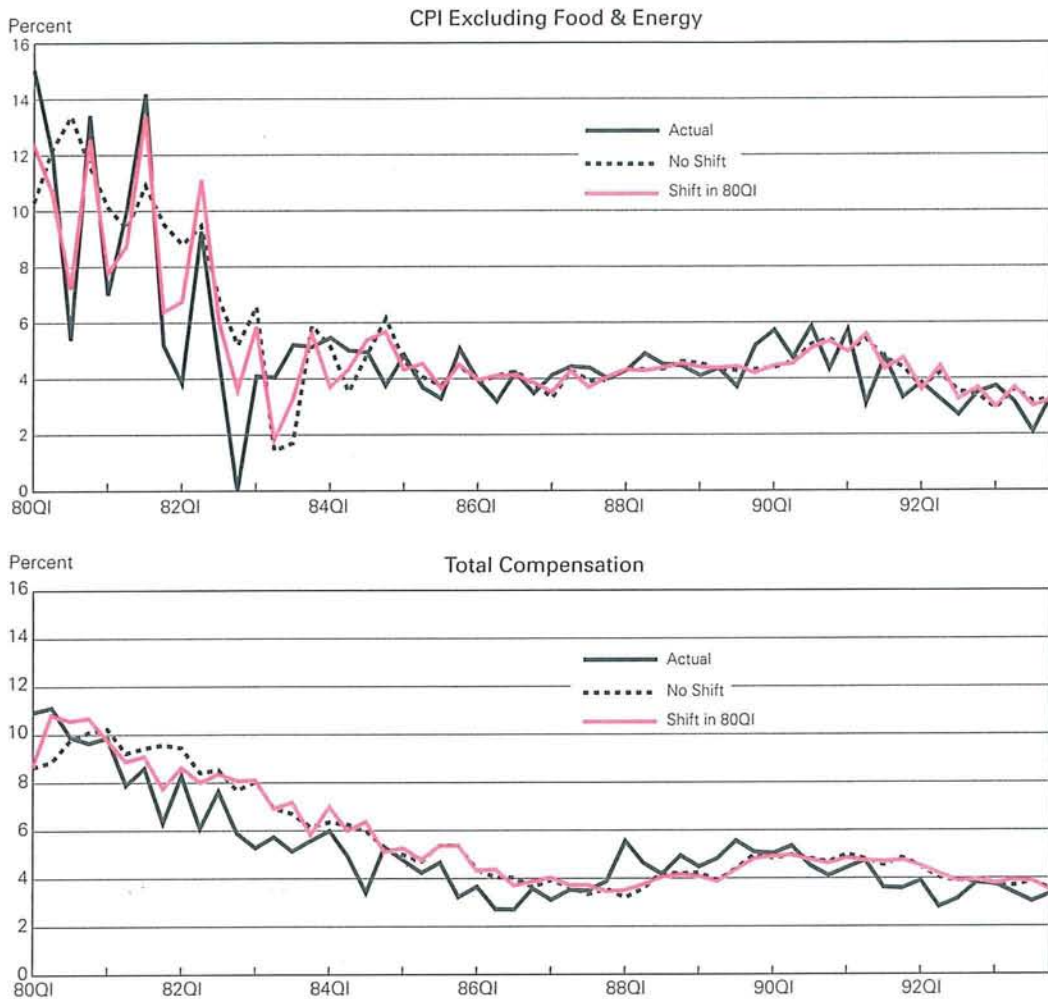
Shifts in the Unemployment Coefficients

What about the other coefficients in the Phillips curve? Were they stable as well? Tootell's article in the September/October 1994 issue of this *Review* addresses the stability of the NAIRU, finding no evidence for a shift in the NAIRU, and this research endorses that conclusion. But the NAIRU could remain stable even if the unemployment coefficients shift, as long as the sum of the unemployment coefficients remains the same or the constant in equation (2) shifts in the same direction as the sum to maintain the ratio that determines the NAIRU. Thus, it is of interest to see whether the unemployment effect in the Phillips curve has shifted over time.

In general, the possibility of a shift in unemployment coefficients can be viewed as another implication of the Lucas critique. Suppose the "underlying structure" behind the Phillips curve includes a monetary policy response that raises interest rates when inflation is currently higher than desired and raises the unemployment rate. Then part of the contemporaneous

Figure 7

Expected Inflation Component of Phillips Curve



Source: Consumer price index, all items excluding food and energy, seasonally adjusted, U.S. Bureau of Labor Statistics
 Total compensation, seasonally adjusted, U.S. Bureau of Labor Statistics.
 Before 1982 - Nonfarm business sector, compensation per hour.
 1982 to present - Employment cost index, compensation, civilian workers.
 Author's calculations.

ous inflation-unemployment correlation captured by the Phillips curve could include an imbedded monetary policy response to inflation. When the policy response changes, that part of the inflation-unemployment correlation will also change, and the Phillips curve will not be stable for that reason. This potential source of instability is less important when the Phillips curve is specified with a lagged unemployment rate, as it is in Tables 1 to 1c.

Columns (2) of Table 2 present tests of the signif-

icance of shifts in the lagged unemployment coefficients. For most of the breakpoints tested, no evidence is found of a shift in these coefficients. In 1983, however, there appears to be some benefit to the core CPI equation fit from allowing the unemployment coefficients to shift. Note, however, that while the coefficients may have shifted, the overall effect of unemployment on inflation (the sum of the coefficients) appears not to have shifted significantly. Columns (4) display the estimated shift in the sum of the

lagged unemployment coefficients. All of the shifts are 0.04 or less, not significantly different from the sum reported in Table 1. Thus, these results suggest that overall, the unemployment effect in the Phillips curve has also been quite stable. A slight shift in the lag pattern may have occurred around 1983, but it did not persist after 1983, and so it is as likely that this statistical result simply reflects overfitting of a few observations as was the case with the expected inflation component, discussed above. Once again, coefficients in the compensation Phillips curve show no sign of instability.¹⁹

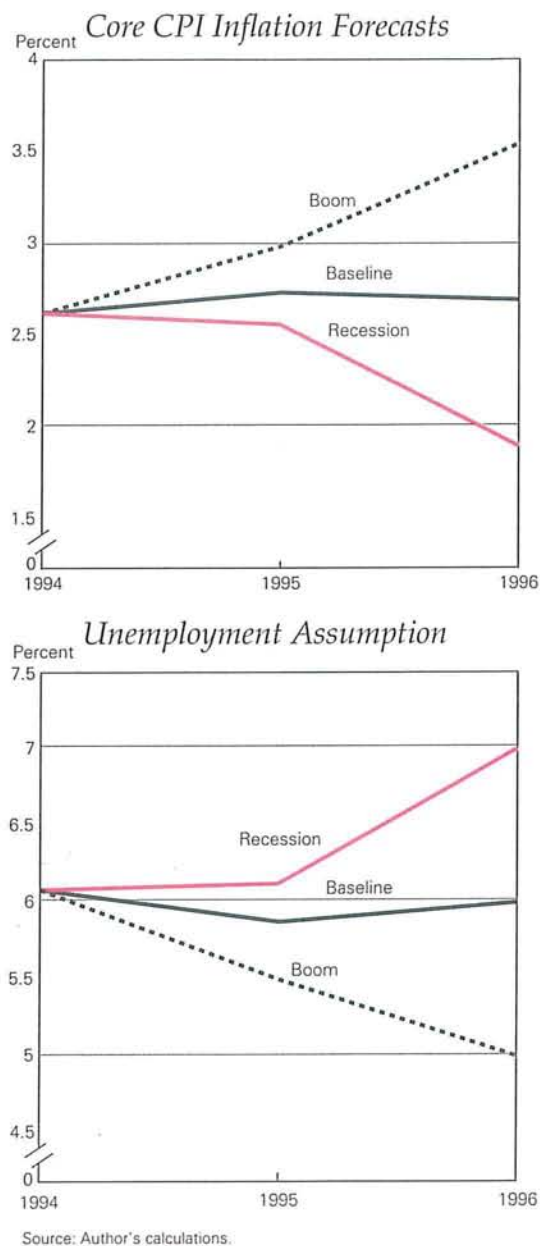
The Empirical Force of the Lucas Critique

Lucas and Sargent (1978) argued that "there is no reason, in our opinion, to believe that these models have isolated structures which will remain invariant across the class of interventions that figure in contemporary discussion of economic policy." Taking the results of Figures 2 to 7 and Tables 1 and 2 together, it would be difficult to argue that the Phillips curve is unstable across the observed changes in policy regimes. Whether using in-sample fit, out-of-sample simulation, coefficient stability tests, or a direct examination of the expected inflation component of the curve, the same conclusion emerges. The Phillips curve has exhibited remarkable stability over the past 35 years, even across data for what must be the most dramatic shift in monetary policy regime in the United States since World War II. As an empirical matter, the Lucas critique does not apply to the Phillips curve.

V. Phillips Curve Predictions

What does the estimated core CPI, price-price Phillips curve of Table 1 predict for inflation over the next two years? The answer depends, of course, on the assumed path of the unemployment rate.²⁰ Figure 8 depicts the annual average inflation rates predicted by the Phillips curve for three different unemployment

Figure 8



scenarios. In the baseline case (the solid black line), the unemployment rate rises back to the estimated NAIRU in 1995, remaining there through 1996. In the "recession" case (the red line), the unemployment rate rises to 1 percentage point above the NAIRU in 1996. In the "boom" case (the dashed black line), the unemployment rate falls to 1 percentage point below the NAIRU by 1996.

¹⁹ The results for total CPI are no more distinct than those for the core CPI. For total CPI, only in 1982 and 1983 does the probability dip below 0.05, and it then quickly reverts to 0.63 in 1984. The GDP deflator results are still weaker, as the probability gets only as low as 0.07 in 1981, and never dips below 0.30 for any other breakpoint date.

²⁰ A fully articulated model would include an equation that describes the behavior of the unemployment rate, as well as the price of oil. With such a model, we could jointly forecast the behavior of unemployment and inflation. Such an effort lies outside the scope of this paper.

In all cases, the changes in inflation predicted by the Phillips curve are modest. As a result of the inertia imparted by the expectations component of the estimated equation, inflation adjusts very gradually to the unemployment gaps. Within the first year of the forecast, a difference of about 0.5 percentage point either way in the unemployment rate alters the inflation forecast by no more than 0.3 percentage point. As the unemployment gaps widen to plus or minus 1 percentage point in the second year of the forecast, the differences in inflation forecasts become more noticeable, yielding changes in inflation of about 0.6 percentage points relative to the 1994 average rate.

Of course the strong persistence of inflation implies that a rising inflation rate at the end of 1996 will not immediately reverse its course in 1997. Thus, the inflation outlook beyond 1996 is also of interest. However, it would be unwise to use the Phillips curve in isolation for longer forecast horizons. The projected outcomes for inflation and unemployment shown in Figure 8 would likely elicit a significant monetary policy response that would alter the outcomes for inflation and unemployment. A model that incorporates the dynamic interactions among inflation, interest rates, and unemployment is required for such an exercise. Still, the basic structure of the Phillips curves explored in this paper suggests that unemployment outcomes like those displayed in Figure 8 should result in only modest changes in the core inflation rate in the short run.

VI. Conclusions

Do these empirical results suggest that the Phillips curve can and should be used for policy exercises, counter to the expectations of Lucas and Sargent at the 1978 Boston Fed conference? It would appear so, although some caveats are in order.

As Lucas and Sargent suggested, the ultimate test of whether equations are structural or not lies in their empirical performance. Taking the empirical results in section IV literally, there is no sign that the Phillips curve is not structural. The double negative in the preceding sentence, however, indicates the certainty

with which we can hold the conclusions of the empirical results. We cannot prove that the equation is structural; we simply cannot find evidence that it is not.

This last objection is not simply semantic. The economic content to the caveat is that we may not have observed the shifts in policy behavior that are important enough to derail the empirical performance of the Phillips curve. While the disinflation from the late 1970s levels of inflation to the present was dramatic, it was still relatively gradual (as compared to the ends of some hyperinflations, for example), and thus could be captured by a model that allows for a gradual influence of the unemployment rate on the rate of inflation. Perhaps more dramatic departures from historical monetary policy might not be captured as well by the Phillips curve.²¹

One such qualitatively different policy would be an effort by the Federal Reserve to target the level of prices, rather than its rate of change. Hall (1984, pp. 137–38) provides an entertaining narrative that highlights the difference between this strategy and the inflation-targeting strategy pursued over the past 30 years. Would the Phillips curve retain its remarkable stability across such a change in monetary policy? The experiment has not been attempted yet, and so we cannot know.²²

²¹ Thomas Sargent (1981) finds evidence in favor of the Lucas critique in economies that recover from hyperinflation. The causes of inflation rates that are measured in the 100s or 1000s are considerably different from the causes of more common single- and double-digit inflation rates (hyperinflations are most often due to extreme lack of fiscal restraint). Thus, policies that stop hyperinflation have fundamentally altered the process generating inflation, and expectations about inflation's behavior after the hyperinflation are much more likely to differ from expectations during the hyperinflation.

²² While Phillips' sample (1861–1957) included periods during which the inflation rate was approximately zero and the price level was fairly flat, the periods do not parallel the exercise contemplated in the text. First, monetary policy in the United Kingdom was not charged with actively leaning against the wind as it is currently in the United States, and in any event was almost certainly not directly targeting the price level. Second, Phillips' original curve did not include the expected inflation component at issue here, so we do not know from his evidence that an expectations-augmented Phillips curve would have performed well in zero-inflation or stable price-level environments.

References

- Duesenberry, James. 1948. "Income-Consumption Relations and Their Implications." In *Income, Employment, and Public Policy, Essays in Honor of Alvin Hansen*, pp. 54-81. New York: W.W. Norton.
- Gordon, Robert J. 1970. "The Recent Acceleration of Inflation and Its Lessons for the Future." *Brookings Papers on Economic Activity*, vol. 1, pp. 8-41.
- _____. 1977. "Can the Inflation of the 1970s Be Explained?" *Brookings Papers on Economic Activity*, vol. 8, pp. 253-77.
- _____. 1980. "A Consistent Characterization of a Near-Century of Price Behavior." *American Economic Review Papers and Proceedings*, vol. 70, pp. 243-49.
- _____. 1994. "Inflation and Unemployment: Where Is the NAIRU?" Paper presented to the Board of Governors of the Federal Reserve System, Meeting of Academic Consultants, December.
- Haavelmo, T. 1944. "The Probability Approach in Econometrics." Supplement to *Econometrica*, vol. 12, pp. 1-118.
- Hall, Robert E. 1984. "Monetary Strategy With an Elastic Price Standard." In *Price Stability and Public Policy*, Federal Reserve Bank of Kansas City conference volume, pp. 137-59.
- Klein, Lawrence R. and Arthur S. Goldberger. 1955. *An Econometric Model of the United States, 1929-1952*. Amsterdam: North-Holland.
- Lucas, Robert E., Jr. 1976. "Econometric Policy Evaluation: A Critique." In *The Phillips Curve and Labor Markets*, ed. K. Brunner and A. H. Meltzer, Carnegie-Rochester Conference Series on Public Policy, 1, pp. 19-46. Amsterdam: North-Holland.
- Lucas, Robert E., Jr. and Thomas J. Sargent. 1978. "After Keynesian Macroeconomics." In *After the Phillips Curve: Persistence of High Inflation and High Unemployment*, Federal Reserve Bank of Boston Conference Series No. 19, pp. 49-72.
- Perry, George L. 1966. *Unemployment, Money Wage Rates, and Inflation*. Cambridge: The M.I.T. Press.
- Phillips, A. W. 1958. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861-1957." *Economica*, November, pp. 283-99.
- Sargent, Thomas J. 1981. "The Ends of Four Big Inflation." Working paper, Federal Reserve Bank of Minneapolis.
- Tootell, Geoffrey M. B. 1994. "Restructuring, the NAIRU, and the Phillips Curve." *New England Economic Review*, Sept./Oct., pp. 31-44.