Carnegie-Rochester Conference Series on Public Policy 41 (1994) 157-219
North-Holland

The post-war U.S. Phillips curve: a revisionist econometric history

Robert G. King*

University of Virginia, Charlottesville, VA 22901

Federal Reserve Bank of Richmond

and

Mark W. Watson†

Northwestern University, Evanston, Illinois 60208

Federal Reserve Bank of Chicago

Abstract

In 1958, A.W. Phillips discovered a strong negative correlation between inflation and unemployment in United Kingdom data. Continuing controversy surrounds the long-run trade-off suggested by a curve he drew through these observations.

We conduct a wide-ranging investigation of the post-war U.S. Phillips correlations and Phillips curve. Many economists view the Phillips correlations as chimerical, given the rise in both inflation and unemployment during the 1970s, and the Phillips curve as plagued by subtle identification difficulties raised by Lucas and Sargent. Yet, a strikingly stable negative correlation exists over the business cycle, and recent theory indicates the Lucas-Sargent critique may not be empirically relevant. When we estimate the long-run trade-off as Gordon and Solow did, we find it is roughly one-for-one. This traditional Keynesian identification also makes business cycles entirely due to demand shocks. However, the Gordon-Solow model is not the only one that fits the data well. Alternative identifications lead to much more modest effects of demand on business cycles and essentially negligible long-run trade-offs.

*We have received many constructive comments on this paper: we particularly thank Charles Evans, Robert J. Gordon, Bennett McCallum, and Charles Plosser. Support was provided by the National Science Foundation via grant NSF-91-22463.

† *Correspondence to:* Mark W. Watson, Northwestern University, Department of Economics, 2003 Sheridan Road, Evanston, Illinois 60208.

SSDI 0167-2231(94)00018 H

1 Introduction

The relationship between inflation and unemployment has been among the most controversial macroeconometric topics of the post-war period. Early work by Phillips (1958) for the United Kingdom documented a pronounced negative correlation between these series over 1861–1957. This *Phillips correlation* was subsequently investigated for other countries and other periods. Notably, Samuelson and Solow (1960) suggested that a similar negative relationship held for the United States over roughly the same time period that Phillips had studied.

Phillips and Samuelson-Solow drew curves through the inflation and unemployment data; they used these as structural relations to discuss a long-run *Phillips curve trade-off* between inflation and unemployment. Subsequent macroeconometric research built structural models of inflation and unemployment designed to add more detailed theoretical underpinnings to the Phillips curve; it also uncovered quantitatively important dynamic interactions between these variables (e.g., in Gordon [1970]) so that short-run as well as long-run trade-offs could be explored.

During the 1960s and 1970s, economists of widely varying perspectives focused their research on the Phillips correlation and the Phillips trade-off. There was also substantial exchange of views on the origins and significance of these relationships. For example, the classic Eckstein [1972] volume on the econometrics of price determination includes contributions on the testing of the natural rate hypothesis by a neoclassical economist (Lucas), on the effect of money on prices and output by monetarist economists from the St. Louis Federal Reserve Bank (Anderson and Carlson), on models of price-setting by a theoretically oriented Keynesian (Nordhaus), on the wage-price sectors of large econometric models by academics (Hymans and Klein) and Federal Reserve researchers (de Menil and Enzler), on the construction of wage and price statistics by an NBER business-cycle researcher (Moore), and on the typical spectral shape of price series by a time series econometrician (Nerlove). Through roughly 1980, the nature and stability of the linkages between inflation and unemployment was the central topic for macroeconometric research.

1.1 *The great divide*

In the last decade, however, research on this topic has fallen into a period of quiescence. Curiously, this relative lack of research activity does not reflect the emergence of a general consensus, but rather the division of macroeconomists into two groups with widely different perspectives on the structure and stability of the linkages between inflation and unemployment.

Neoclassical and monetarist economists. On one side of the street, the Phillips curve and the Phillips correlation essentially disappeared as research topics as a result of three related factors.

First, new empirical evidence suggested a striking change in the behavior of inflation and unemployment after 1970. Notably, the pronounced negative correlation of Phillips apparently disappeared from the U.S. data after 1970: since 1970, there have been lengthy time periods in which inflation and unemployment were positively rather than negatively associated. Thus, it became possible to view the Phillips correlation as effectively dead in the post-1970 period, relegating it to the list of facts that held for some periods but not for others. To some, this indicated that the Phillips correlation was an empirical feature of secondary interest for business-cycle theory and forecasting.

The second factor was new theory: more than two decades ago, Lucas [1972] and Sargent [1971] argued that studies of the links between inflation and unemployment were subject to a subtle identification problem. Adopting the then-prevailing view that the real effects of nominal disturbances depended on whether these were anticipated or unanticipated, Lucas and Sargent showed that it could be impossible to estimate the long-run Phillips trade-off using then-standard econometric methods. In particular, if there was no permanent variation in inflation over the sample period, then the effect of a change in trend inflation could not be determined without a fully articulated behavioral model. Even more strikingly, Lucas and Sargent constructed examples of "natural rate models"—settings without any effect of sustained inflation on unemployment—that displayed an apparent long-run trade-off. To many economists, the increase in U.S. inflation during the 1970s corresponded to the "grand experiment" of permanently increasing the growth rate of nominal aggregate demand as envisioned earlier by Samuelson and Solow [1960]. Thus, the simultaneous rise in the average levels of inflation and unemployment during the 1970s provided striking evidence against any long-run trade-off, consistent with predictions by Friedman [1968] and Phelps [1967], and suggested that earlier apparent trade-offs were due to the identification problem stressed by Lucas and Sargent. Overall, in the eyes of neoclassical and monetarist economists the result was that the Keynesian macroeconometric models displayed "econometric failure on a grand scale," in the phrase of Lucas and Sargent [1979], due mainly to the structural specifications of wage and price adjustment underlying the Phillips curve trade-off in these models.

The third factor leading to disappearance of research on the Phillips curve was the necessity of developing new methods to execute the program advocated by Lucas and Sargent. On the one hand, to construct the fully articulated dynamic models that Lucas and Sargent advocated, many neoclassical economists spent the bulk of their energy working on real theories of aggregate fluctuations. While these models are arguably a natural starting point for macroeconomic analysis, they are ones in which Phillips trade-offs are essentially absent, and the implications of these models for the Phillips correlation have typically not been explored. On the other hand, the solution of the identification problem uncovered by Lucas and Sargent involved a new style of econometrics and associated technical problems.

Keynesian economists. On the other side of the street, many Keynesian macroeconomists were surprised by the strong reaction of neoclassical and monetarist economists to the new theory and new evidence. That is, particularly for those Keynesian macroeconomists engaged in forecasting, the remarkable feature of the Phililps curve in the post war period was its *stability*. To be sure, there was the embarrassing failure to predict the post-1970 increase in average levels of unemployment and inflation; that difficulty could be fixed, however, by requiring that there be no long-run trade-off implied by the appropriate structural specifications of Keynesian macroeconometric models and by incorporating exogenous variables to capture the effects of "supply shocks." But, with these modifications, the standard structural identification of Phillips [1958], Solow [1969], and Gordon [1970] continued to be a powerful tool for organizing the dynamics of unemployment and inflation. For the short term, it provided a reasonable forecasting tool, even after 1970. In the longer term, the world was a more difficult place to forecast after 1970 and the conventional equations did as well as anything else. The Keynesian macroeconometricians argued that one could continue to use the structural Phillips curve for considering the effects of monetary policy acceleration or deceleration: they computed the dynamic effects of these policy shifts and evaluated their benefits and costs using methods that were essentially unchanged by the arguments of Lucas and Sargent.[1]

Thus, during the 1980s, business-cycle research was basically conducted by two groups. The first did not study the Phillips correlation or the Phillips curve because the former was viewed as chimerical and the latter as subject to deep identification problems. The second viewed the Phillips curve as an essentially intact structural relation: most research activity sought to add variables to represent supply shocks and to build in a zero long-run trade-off. Even without these modifications, conventional Phillips curves continued to be a much-used tool for medium term forecasting and policy analysis.

1.2 *Our revisionist history*

In this paper, we reexamine the post war U.S. unemployment and inflation experience using some new theoretical results and two complementary time

[1] For example, see Gordon and King [1982].

series econometric methodologies. Our history of the Phillips curve and the Phillips correlation is revisionist because it challenges important aspects of the views of each of the prevailing schools.

The neoclassical/monetarist position. We provide challenges that are theoretical and empirical. On the theoretical side, we find that the Lucas-Sargent identification problem may not be as devastating for the long-run estimates of Solow-Gordon as originally supposed. (In this regard, we detail arguments made earlier in Fisher and Seater [1993] and King and Watson [1992].) Notably, we show that if inflation contains important low-frequency variation, as captured by a "unit root" stochastic process, then one can estimate long-run trade-offs using procedures like those of Gordon and Solow. Further, these types of estimates can be large even after 1970, as shown in King and Watson [1992].

On the empirical side, when we reexamine the nature of the "stylized facts," we find evidence that the Phillips correlation is not dead. To do this, we decompose the post war inflation and unemployment time series into three parts, which may usefully be labelled as "trend," "business cycle," and "irregular components." We find that there is a pronounced negative relationship between the business-cycle components of inflation and unemployment. That is, for the post war U.S. data, the Phillips correlation is alive and well, once one recognizes that it lives at the business-cycle frequencies. After 1970, simultaneous increases and decreases in trend inflation and unemployment obscure these business-cycle comovements. We also document that there is a changing cyclical pattern of inflation after 1970: as inflation became more volatile there was a corresponding increase in its covariance with unemployment, so that the overall correlation remained roughly over the entire postwar period.

The Keynesian position. Our documentation of the post war Phillips correlation at business-cycle frequencies would seem to be good news for Keynesian macroeconomics. But challenges also appear for the traditional Keynesian position when we estimate structural models of the Phillips curve. The traditional Keynesian identification of the structural Phillips curve is not the only one which fits the post war data well and, on many dimensions, it appears to be an extreme one.

Structural modeling. When we begin the process of structural modeling, we uncover an additional key feature of the dynamic interactions of inflation and unemployment: there a near-causal ordering, in the sense of Granger [1969], of inflation and unemployment. More specifically, in a reduced form vector autoregression, past unemployment is important for predicting current inflation but past inflation contains little information about current unemployment.

This finding has important implications for structural modeling of inflation and unemployment. As in any simultaneous equations setting, we must make some identifying assumptions to build a structural Phillips curve and to explore Phillips trade-offs. When we consider alternative identifying assumptions about the short-run effect of inflation on unemployment, we find that these have dramatic effects on the dynamic structure as well. For example, there is a ready real business-cycle interpretation of the data: if we assume that there is little short-run effect of inflation on unemployment, then there is also little effect at any horizon. This outcome, which results from the near Granger-causal structure, then implies that most of the negative Phillips correlation at business-cycle frequencies arises from the negative effect of unemployment on inflation.

Traditional Keynesian estimates of the Phillips curve by Solow [1969], Gordon [1970] and others were based on the assumption that unemployment is dominated by aggregate demand disturbances, sufficiently so that it may be used as a regressor in wage and price equations. While this identification is one that many macroeconomists would now question, it was standard practice in Keynesian macroeconomics before 1970: we find that it leads to a very traditional Keynesian interpretation of business cycles. That is, when we assume that there is a large short-run effect of inflation on unemployment, then the near Granger-causal structure has the reverse implication: post war fluctuations in macroeconomic activity are entirely driven by shocks to demand which manifest themselves mainly in unemployment in the short-run and with little short-run effect on inflation. Further, there are also major costs of disinflation. Notably, under the Gordon-Solow identification, we find that there is a large long-run Phillips curve slope in both the pre-1970 and post-1970 data. We also find large "sacrifice ratios," defined as the unemployment cost of disinflation over a five-year period as in Okun [1978], which are in line with the traditional Keynesian estimates surveyed by Okun.

An intermediate interpretation of the time series results from an identification that has its origins in the rational expectations monetarist (REM) studies of Sargent [1976] and Barro and Rush [1980]. Under this identification, there is a Phillips curve with a small long-run slope. Business cycles turn out to be about half a result of disturbances to demand and those to aggregate supply. There is also a much smaller sacrifice ratio, which is broadly in line with recent estimates by Ball [1993] and Mankiw [1990].

All of these structural models fit the data equally well. Yet, they have substantially different implications for the sources of business cycles, for the trade-off between inflation and unemployment, and for the interpretation of particular historical episodes.

1.3 *Implications for the long-run trade-off*

To preview our results, it is useful to consider the problem of estimating the long-run trade-off between inflation and unemployment. Looking across various subsamples of the post war period in King and Watson [1992], we found clear evidence of structural change in the behavior of post war U.S. inflation and unemployment that occurs around 1970. Further, as predicted by Lucas and Sargent, higher average values of inflation and increased persistence of inflation are associated with a decline in the estimate of the long-run trade-off between unemployment and inflation constructed from our analogues of the methods of Gordon [1970] and Solow[1969].

This decline in the long-run trade-off is quantitatively important, as illustrated in the first column of Table 1. (These estimates are based on King and Watson [1992] and explained in greater detail in Sections 3 and 4 below.) In the sample period that corresponds to Gordon's [1970] study of the Phillips curve, 1954–1969, we find that our variant of his methods suggests that a permanent one-percent increase in inflation — arising from an increased growth rate of aggregate demand — is associated with a 1.3 percentage-point decline in the unemployment rate. Strikingly, this estimate is close to the unit value which Solow [1970] extracts from Gordon's [1970] study; the unit value is also suggested by a graph in the 1969 Annual Report of the President representing the relevant trade-off for the United States. In the latter sample period, 1970–1992, the estimate is cut by roughly one-half to -0.57.[2] Thus, there is clear evidence of the importance of the Lucas-Sargent critique for the estimation of structural Phillips curves, at least as it bears on the estimated extent of long-run trade-offs between inflation and unemployment. However, the extent of the long-run inflation and unemployment trade-off remains quantitatively important for the full sample period. One view is that this is because inflation is an I(0) process, as in the models of Lucas and Sargent, but results of unit root tests suggest considerable persistence, consistent with an I(1) inflation process. Indeed, as we will show below unit root tests are consistent with the presence of permanent components in inflation even in the 1954-69 sample period.

The traditional Keynesian (TK) estimates of the Phillips curve by Solow [1969], Gordon [1970], and others were based on the assumption that unemployment is dominated by aggregate demand disturbances, sufficiently so that it may be used as a regressor in wage and price equations. When we use an alternative identification derived from the "rational expectations mone-

[2]Throughout this paper we measure the Phillips curve tradeoff by the ratio of the change in the unemployment rate to the change in the inflation rate, i.e., $\partial u/\partial \pi$ in the notation used below. This is the inverse of the traditional Phillips curve slope. From a testing standpoint, it is useful that for the measure $\partial u/\partial \pi$, neutrality corresponds to a value of zero, instead of $-\infty$ in the usual Phillips curve slope.

Table 1:
Estimated Long-Run Phillips Curve Trade-Offs

Sample Period	Keynesian	Monetarist
1954–69	-1.30	-0.47
1970–92	-0.57	-0.23
1954-92	-0.71	-0.29

Notes: The long-run trade-off is defined as: $lim_{k \to \infty}[\partial u_{t+k}/\partial \epsilon_{dt}][\partial \pi_{t+k}/\partial \epsilon_{dt}]^{-1}$, where ϵ_{dt} denotes a "demand" shock, defined formally in Sections 3 and 4.

tarist" (REM) studies of Sargent [1976] and Barro and Rush [1980], we find a very different estimate of the long-run trade-off between inflation and unemployment. Our estimate is given in column 2 of Table 1: it is -.47 for the pre-1970 sample period and -.23 for the later sample period. Further, when we implement a real business-cycle (RBC) identification, which is that there is no short-run link between inflation and unemployment, we find that there is an essentially zero trade-off in both periods.

By looking down the two columns of Table 1, one gets a sense of how the sample period affects the long-run trade-off between inflation and unemployment. By looking across the rows, one gets a sense of how short-run identifications affect the long-run trade-off between inflation and unemployment. One reaction to this table is that the short-run identifications are quantitatively at least as important as the sample period for the long-run trade-off between inflation and unemployment. In the remainder of the paper, we will see that this reaction is consistently appropriate as we consider a range of evidence on the dynamic interaction of inflation and unemployment.

1.4 *Plan of the paper*

The organization of the remainder of the paper is as follows. We begin in Section 2 by reporting the data that we study in the paper, both in their original form and using moving averages designed to highlight "trend" and "cyclical" components of the data. While the basic time series display little evident correlation, this section documents a remarkably stable negative correlation between unemployment and inflation over business-cycle frequencies, which we refer to as the "Phillips correlation" throughout the paper. Section 3 begins our discussion of the "Phillips trade-off," defined as the relative effects of aggregate demand on unemployment and inflation within a particular structural model. It begins with the theoretical background to our study, considering the types of models that Keynesian and Monetarist economists have typically used to give structure to the Phillips correlation. This section then considers the identification problems raised by Lucas [1972] and Sargent [1971], and discusses how these problems are affected by unit roots in the inflation process. In Section 4, we develop a bivariate dynamic structural model of inflation and unemployment. Working with the three alternative short-run identifications discussed above, we investigate the model's implications for (i) the dynamic response of unemployment and inflation to demand disturbances; (ii) the contribution of "demand" and "supply" shocks to post war U.S. economic fluctuations suggested by this model; (iii) the contribution of demand disturbances to specific historical episodes; and (iv) the long-run trade-off between unemployment and inflation. In Section 5, we consider issues of the econometric stability of the post war inflation and unemployment processes, seeking to assess the importance of changing structure for both forecasting and structural estimates of the Phillips trade-off. In Section 6, we conduct a sensitivity analysis of our main empirical results to various assumptions made elsewhere in the paper, including evaluating the importance of unit roots, data frequency, measures of inflation, and inclusion of exogenous supply shocks. Section 7 is a summary and conclusion.

2 The Phillips correlation in post war U.S. data

We begin with a brief review of the post war U.S. inflation and unemployment experience; this also serves to introduce the data that we will study in the remainder of the paper. Panel A of Figure 1 plots the monthly inflation rate series, π_t, which is the annualized percentage rate of change in the consumer price index.[3] Because the monthly inflation series is very choppy, we also graph the annual average inflation rate, which is the bold solid line in the figure.[4] The vertical lines in the panels of Figure 1 indicate the sample period that we use to estimate our "early" Phillips curve: it is chosen to match the effective sample period of Gordon [1971], who excluded the earlier observations to eliminate the inflation of the Korean War and some large outliers immediately following the lifting of World War II price controls. Panel B of Figure 1 plots the unemployment rate.[5] Some summary features

[3]We use the seasonally adjusted index for all urban consumers (Citibase series PUNEW). The annual inflation rate is $\pi_t = 1200 * log(cpi_t/cpi_{t-1})$, where cpi_t is the value of the index at time t.

[4]The choppiness in the raw data reflects rounding error in the CPI, which is reported to one decimal place (with a base of 100 in 1982–84). The annual average inflation rate shown in the figure is the centered moving average $\pi_t^a = (\sum_{i=-6}^{6} \pi_{t-i})/13$.

[5]This is the unemployment rate for all workers 16 years and over, in percent and seasonally adjusted (Citibase series LHUR).

of these data are presented in Table 2, which presents means and standard deviations for various subintervals of the post war (1950–1992) period.

Table 2:
Summary Statistics

	Unemployment Rate		Inflation	
Sample Period	$\bar{X}$	S	$\bar{X}$	S
1950–53	3.61	1.08	3.27	5.27
1954–92	5.98	1.54	4.27	3.98
1954–59	5.11	1.03	1.48	3.28
1960–69	4.77	1.07	2.48	2.66
1970–79	6.22	1.16	7.12	3.81
1980–92	7.11	1.38	4.73	3.71

Notes: $\bar{X}$ denotes the sample mean and S denotes the sample standard deviation.

There are three distinct features of the time series plotted in Figure 1. First, there is large high-frequency variation in inflation. Second, both series show significant variation over periodicities associated with U.S. business cycles. Finally, both series show slowly varying average levels or trend behavior; for example, the average inflation rate increased from 2.5% in the 1960s to 7.1% in the 1970s and then fell to 4.7% in the 1980s.

These features of the time series are highlighted in Figure 2, which shows the results from passing the data through symmetric two-sided filters designed to highlight the contribution of various periodicities. Panel A displays low-frequency variation (components with periodicities longer than eight years), panel B displays business-cycle variation (components with periodicities between eighteen months and eight years), and panel C displays high-frequency variation (components with periodicities less than eighteen months).[6]

[6]Baxter and King [1993] discuss the construction of optimal approximate band-pass filters: some approximation is necessary because exact band-pass filters are infinite two-sided moving averages. Their definition of optimal is a standard one in the frequency domain literature, i.e., the filters are designed to minimize the integrated square losses over all frequencies subject to constraints on specific points. The low-pass filter that produces Figure 2, panel A, is constrained to place unit weight at the zero frequency and the band-pass filter that produces Figure 2, panels B and C, is constrained to place zero weight at the zero frequency. The quality of approximation depends on the length of the moving average: here we choose relatively long moving averages that use data from $t-K$ to $t+K$, with $K=60$ months.

FIGURE 1

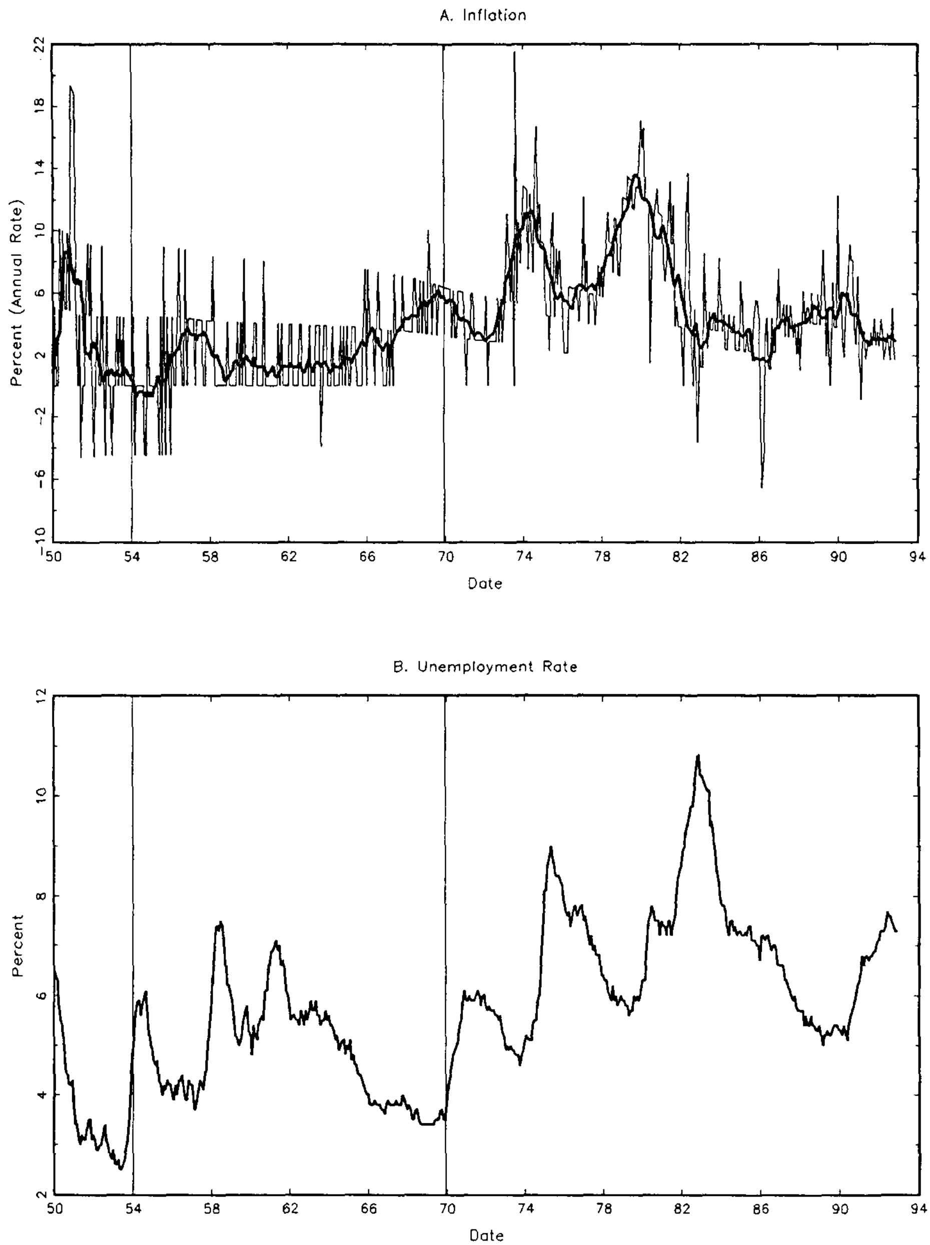

Notes: Panel A: Raw Data (thin line), Centered 13-Month Moving Average (thick line).

FIGURE 2

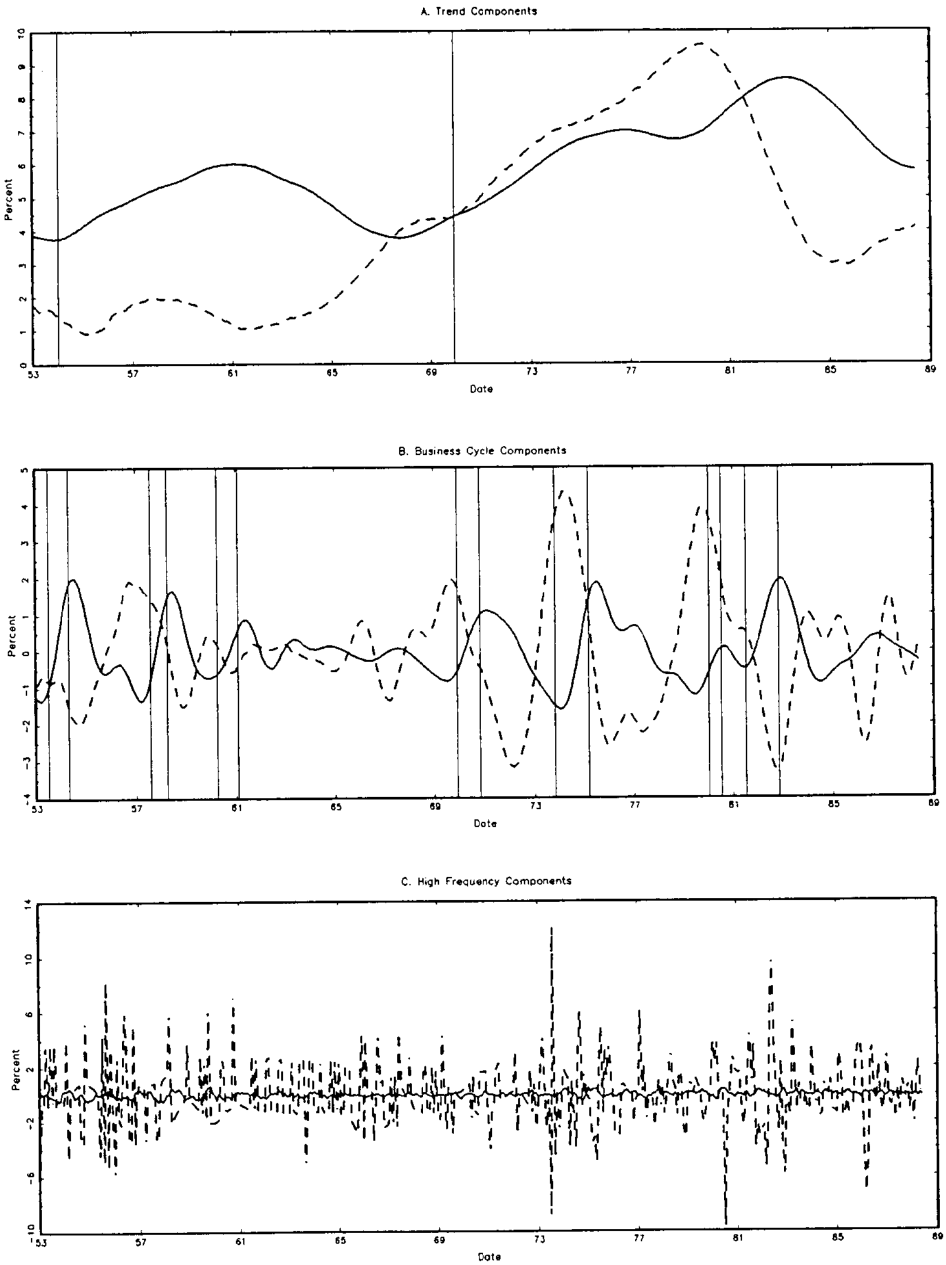

Notes: Unemployment (solid line), Inflation (dashed line).

The series plotted in panel A show the long-run movements in unemployment and inflation over the post war period obtained from the low-pass filter (isolating periodicities greater than 8 years).[7] Inflation was low in the late 1950s and early 1960s with an average level of approximately 1.5%; this was followed by a rising trend in inflation that peaked at nearly 10% in 1980, followed by a decrease to 4% by 1990. Similarly, the trend unemployment rate drifted up from approximately 4% in the late 1960s to 8% in the early 1980s, before falling to 5% in the late 1980s. The figure also shows an apparent change in the correlation between the long-run movements in unemployment and inflation. During 1954–1969 (shown by vertical lines on the graph) there is a strong negative correlation (-.62); from 1970–1987 there is no consistent relation (the sample correlation is .03), and over the entire period there is a positive correlation (.50).[8]

Panel B shows the analogous "business-cycle" components of the series, that is, the results of applying a band-pass filter that isolates periodic components of between 18 months and 8 years in duration. Before and after 1970, these series vary as Phillips would have expected, with the unemployment rate rising and the inflation rate falling during NBER-dated recessions: in this sense, there is a remarkable stability of the business-cycle regularities that is masked by the trend and irregular components in Figure 1.[9] But our plot also provides evidence of a change in the inflation process: the inflation process is more volatile post-1970 than it is before 1970, with the standard deviation increasing from 0.93 to 2.04. Despite this changing variability, there is a strong and stable Phillips correlation: the sample correlation of the filtered series is -.69 for 1954–69, -.67 from 1970–1987, and -0.66 for the full sample.[10]

[7]The trend components of inflation and unemployment are very close to the series that results from the familiar procedure of applying a five-year centered moving average — with equal weights — to inflation and unemployment.

[8]We use the sample correlation as a summary measure of association over the two sample periods. Of course, if the series are I(1), as we assume below, these sample correlations are poor estimates of the correlation between the stochastic trends in the series.

[9]One concern about exploration of filtered data such as these is that one is uncovering "spurious relations" that arise from the filtering rather than from the original series. For example, it is well-known that spurious periodic characteristics can be induced by filtering (as discussed in Sargent [1979, chapter XI], for example). It is also well-understood that shifts in timing can be introduced by filtering. However, the features that we stress are unlikely to be spurious. We are not concerned with the periodic nature of the univariate series, which is an artifact of the filtering. Rather, we are interested in the comovements of the two series and their behavior relative to the NBER peak and trough dates. These comovements will be summarized by estimated correlations and associated standard errors that account for serial correlation in the filtered data. Note also that we have applied the same symmetric linear filter to each series, so that no phase shifts have been induced.

[10]Standard errors (computed using an AR(12) spectral estimator) for the estimated correlations are .17, .19, and .18, respectively.

Finally, panel C shows that inflation has much larger high-frequency variation than unemployment. These components have a small negative correlation over the sample period.

Overall, Figures 1 and 2 display two important features of the post war U.S. data. First, low frequency components of inflation and unemployment became more important after 1970 and these did not display the original, negative Phillips [1958] correlation. Second, and by contrast, the Phillips correlation has been remarkably stable over business-cycle frequencies.

3 Theoretical background

In this section, we review the theoretical background to our research. The presentation is designed to highlight the identification problems that are crucial to the study of unemployment and inflation.

3.1 *The traditional Keynesian macroeconometric model*

The IS-LM model of Hicks [1937] and Modigliani [1944] provided the main point of departure for the Keynesian macroeconometric model-building program of the 1950s and 1960s. Given the behavior of nominal wages and prices, the model could determine employment, national income and its components, and the rate of interest. However, in this form, the IS-LM model was incomplete for macroeconometric purposes: a specification of the dynamic adjustment of wages and prices had to be added.

The impressive evidence of Phillips [1958] regarding the relationship between unemployment and wage inflation offered macroeconometric model-builders a way to complete their system. There were two steps in this process. First, the Phillips curve was treated as a structural relation for wage determination. Second, the Phillips curve was combined with a "price equation," as in Eckstein and Fromm [1968], that determined prices as a function of wages and other economic variables.

For expositional convenience, we will conduct our initial discussion without being explicit about dynamics and focus on the contemporaneous interaction of inflation (π_t) and unemployment (u_t). (Thus, for now we make no distinction between π_t and the logarithm of the price level.) Later, we will use these results to discuss short-run identification restrictions in dynamic models, treating the variables that we consider now as the unpredictable components of inflation and unemployment.

Our Keynesian system is given by the two structural equations (1) and (2). The first of these two equations is a "price equation." This specification describes how the inflation rate responds to the unemployment rate, which is an indicator of aggregate demand conditions in the model. The parameter "a" indicates the extent of price adjustment to demand: rapid adjustment of inflation to demand conditions arises for large absolute values of "a"; correspondingly, prices are essentially unresponsive to demand with "a" close to zero. In addition, there is a price shock term, p_t, that is the residual in the price equation.

$$\pi_t = au_t + p_t \tag{1}$$

More generally, the price equation was part of the wage-price block in Keynesian macroeconometric models. Since many Keynesian modelers viewed prices as a relatively fixed markup over wages, however, we use the simpler specification (1).

The second equation in our Keynesian model determines the unemployment rate as a function of inflation π_t and a demand shift variable d_t. This equation captures the Hicksian (IS-LM) determination of real variables as functions of wages and prices.

$$u_t = h\pi_t + d_t \tag{2}$$

The parameter h governs the extent of the short-run Hicksian influence of inflation on demand. Dating from the work of Klein [1950], the conventional Keynesian macroeconometric view was that the short-run dependence of real variables on the price level was minor, suggesting small values of h in equation (2), and that demand variations were dominated by exogenous shocks (d_t). Econometrically, this last observation suggested that the extent of short-run price adjustment in (1) could be estimated via ordinary least squares procedures, as in Solow [1969] and Gordon [1970].

3.2 *Monetarist models of inflation and unemployment*

Monetarist models of inflation and unemployment typically specified an alternative behavioral structure (as, for example, in Anderson and Carlson [1972]). Working in the same static terms as above, they posited an aggregate supply curve, equation (3), and an aggregate demand curve, equation (4), frequently deriving the latter from a simple quantity equation. The aggregate supply specification took the form,

$$u_t = f\pi_t + s_t \tag{3}$$

where large absolute values of f imply large effects of inflation on real activity and s_t is a shock to aggregate supply. Comparably, the demand side took the form:

$$\pi_t = qu_t + m_t \tag{4}$$

where q indicates the sensitivity of inflation to real activity (perhaps combining the income elasticity of money demand with an Okun's law relation between real income growth and unemployment) and m_t is a shock to the inflation rate, typically viewed as originating in changes in the money growth rate.

3.3 *Observational equivalence of the two frameworks*

Generally, the two models described above are observationally equivalent: there is a simple translation from the shocks and parameters in one model into the shocks and parameters of the other. For example, the price equation of the Keynesian model, $\pi_t = au_t + p_t$, can be rewritten as the aggregate supply equation of the monetarist model, $u_t = f\pi_t + s_t$, with the change of variables $f = (\frac{1}{a})$ and $s_t = -(\frac{1}{a})p_t$. That is, in Figure 3, the price shock is the vertical displacement in the supply schedule/price equation and the supply shock is the horizontal displacement. Similarly, the reduced form of the IS-LM model, $u_t = h\pi_t + d_t$, can be rewritten as $\pi_t = qu_t + m_t$ with a change of variables $q = (\frac{1}{h})$ and $m_t = -(\frac{1}{h})d_t$.

While these systems are observationally equivalent, choice between them likely affects the identification strategies used for estimation and inference. Indeed, in our empirical sections below, we will see precisely how important these choices can before one's views of the long-run and short-run trade-off between inflation and unemployment as well as for one's views about the dominant sources of macroeconomic fluctuations.

3.4 *Dynamic Phillips curve specifications*

The dynamic generalization of the Phillips curve specification (1) is:

$$\pi_t = au_t + \beta_{\pi u}(L)u_{t-1} + \beta_{\pi\pi}(L)\pi_{t-1} + p_t \tag{5}$$

where $\beta_{\pi u}(L)u_{t-1}$ permits aggregate demand to have an effect on inflation that is distributed over time and $\beta_{\pi\pi}(L)\pi_{t-1}$ permits price shocks to have rich dynamic effects on the inflation rate. Keynesian macroeconometricians like Eckstein and Fromm [1968] and Gordon [1970, 1971] found that specifications like (5) were necessary to fit the post war data well.

In this context, there are two "slopes" to the Phillips curve that are of interest. First, the short-run slope of the Phillips curve is given by

$$\partial\pi_t/\partial u_t = a. \tag{6}$$

The long-run slope is obtained by contemplating sustained changes in inflation and unemployment – i.e., by setting $u_t = u$ and $\pi_t = \pi$ for all t before taking the partial derivative – so that:

$$\partial\pi/\partial u = [a + \beta_{\pi u}(1)]/[1 - \beta_{\pi\pi}(1)], \tag{7}$$

where $\beta_{\pi u}(1)$ denotes the sum of lag coefficients in $\beta_{\pi u}(L)$ and $\beta_{u\pi}(1)$ is defined analogously. The derivatives in (6) and (7) are the reciprocals of our measure of the unemployment-inflation trade-off.

Figure 3
The Price Equation/Supply Curve

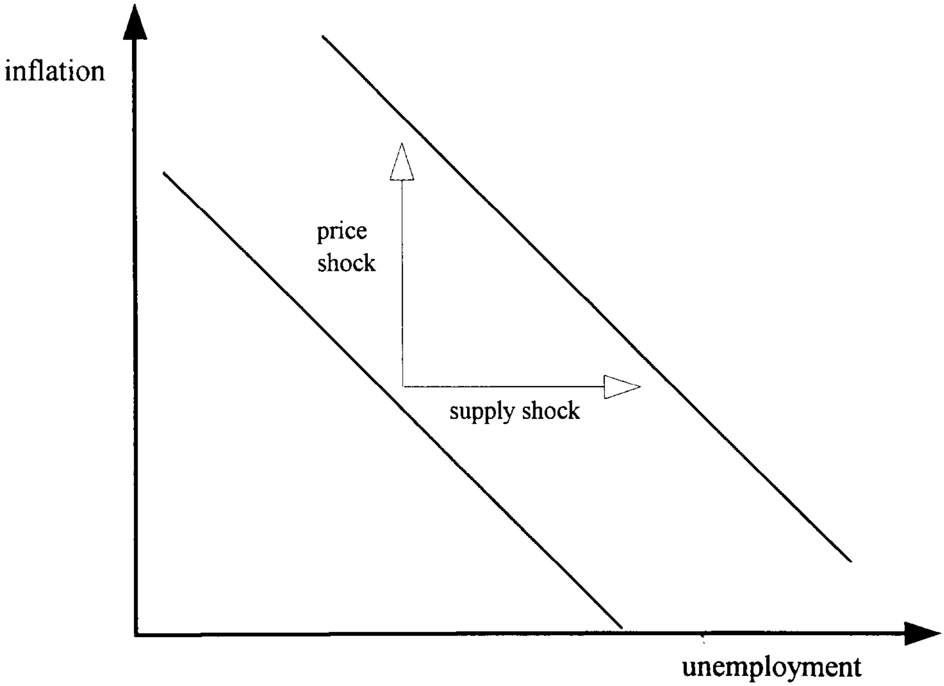

Note: The price equation interpretation of this figure is that $\pi = au + p$, so that the slope is ``$a$'' and the vertical shift is the price shock ``p.'' The supply curve interpretation is that $u = f\pi + s$, so the slope is $(a.f)$ and the horizontal displacement is ``s.'' These models are observationally equivalent with $f = (1/a)$ and $s = -(1/a)p$.

Computation of short-run and long-run trade-offs as in (6) and (7) were standard practice in traditional Keynesian macroeconometrics, including studies that simulated complete models and those that concerned the properties of the wage-price block, i.e., of structural wage and price specifications. Typically, these studies also traced out a full set of dynamic multipliers. Our structural Phillips curve investigation in Section 4 below will follow this path.

3.5 *Testing the natural rate hypothesis*

The natural rate hypothesis of Friedman [1968] and Phelps [1967] suggested no long-run trade-off between inflation and unemployment, i.e., a very large value of $\partial\pi/\partial u$ in (7). Gordon [1970] and Solow [1969] sought to test this hypothesis by making two modifications to the Phillips curve. They began with a structural equation of the general form of (1) augmented to introduce the effects of expected inflation, $\pi_t^* : \pi_t = au_t + b\pi_t^* + p_t$. They then proxied expected inflation using a distributed lag method. In the simplest version, π_t^* was treated as the "adaptive" expectation $\pi_t^* - \pi_{t-1}^* = \theta(\pi_{t-1} - \pi_{t-1}^*)$: this specification had the implication that if there were a permanent increase in inflation, expectations would ultimately capture it in the long run, i.e., $\partial\pi^*/\partial\pi = 1$ for such changes. More complicated schemes allowed for richer dynamic patterns of expectation adjustment, $\pi_t^* = v(L)\pi_t = \sum_{i=0}^{\infty} \nu_i\pi_{t-1-i}$, but continued to impose the requirement that $\partial\pi^*/\partial\pi = 1$ or equivalently the requirement that $\sum_{i=0}^{\infty} \nu_i = 1$. Hence, the Gordon [1970] and Solow [1969] procedures fit naturally into the dynamic Phillips curve specification (5) with $\alpha_{\pi\pi}(L) = b\nu(L)$. If the natural rate hypothesis was invalid, then it was also possible to estimate the extent of the long-run trade-off in an "expectations adjusted" Phillips curve. Indeed, Gordon and Solow found substantial long-run trade-offs in U.S. and U.K. data. The unit value of the $\partial\pi/\partial u$ trade-off estimates that Solow [1970] extracts from Gordon [1970], for example, is not very different from the reciprocal of the value of $\partial u/\partial\pi$ in Table 1: our estimate of $\partial\pi/\partial u$ is .8 for the same sample period.

3.6 *The Lucas-Sargent critique*

The Gordon-Solow tests were criticized by Lucas [1972] and Sargent [1971] in a pair of papers that set the stage for a revolution in macroeconometrics. Our presentation will be a blend of the Lucas and Sargent examples, dealing explicitly with the "monetarist" specification of the Phillips curve as suggested by Lucas' analysis and using a general autoregressive specification of the inflation process as in Sargent [1971].

The structural framework contains two equations. First, the aggregate supply specification makes u_t simply a function of unexpected inflation: it is an expectations-augmented version of (3).

$$u_t = f\pi_t - g\pi_t^* + s_t. \tag{8}$$

The natural rate hypothesis is then that $f = g$. Inflation is generated by the autoregressive process,

$$\pi_t = \rho_1\pi_{t-1} + ...\rho_n\pi_{t-n} + m_t., \tag{9}$$

where m_t is an unpredictable shock. The rational expectations assumption is:

$$\pi_t^* = E_{t-1}\pi_t = \rho_1\pi_{t-1} + ...\rho_n\pi_{t-n}. \tag{10}$$

Then, the reduced form unemployment and inflation relation is:

$$u_t = f\pi_t - \sum_{i=1}^{n} g\rho_i\pi_{t-i} + s_t. \tag{11}$$

Hence, if a macroeconometrician computed our long-run trade-off on the data from this economy, he would find that $\partial u/\partial\pi = (f - g\sum_{i=1}^{n}\rho_i)$. Even if the long-run system is neutral ($f = g$), the assumption that inflation is a stationary stochastic process implies that $\partial u/\partial\pi = f(1 - \sum_{i=1}^{n}\rho_i)$ is not zero. That is, Lucas and Sargent pointed out that there would be an apparent long-run trade-off when none was implied by the structural specification (8). Lucas [1972], Sargent [1971], and Lucas and Sargent [1979] elaborated on the macroeconometric implications of this example: proper tests of models with rational expectations would necessarily involve cross-equation tests. In particular, they argued that it was necessary to provide a detailed structural description of the inflation process in order to test the natural rate hypothesis.

Economists were persuaded by the Lucas-Sargent argument for three reasons. First, it was so clearly correct in its analytics. Second, rational expectations offered a way to avoid treating expectations as a source of "free parameters" in empirical work. Third, their argument provided a coherent explanation of the breakdown of empirical Phillips curves. Notably, it suggested that as inflation became more persistent, in the sense that $\sum_{i=1}^{n}\rho_i$ became closer to unity, then there should be smaller estimates of the long-run trade-off.

3.7 *The Lucas-Sargent example once again*

The structural model (8) with $f = g$ has a very strong natural rate property: inflation affects date t real activity only if it is unexpected as of date $t - 1$. Natural rate models of the class developed by Fischer [1977], Gray [1976], and Phelps and Taylor [1977] permit inflation forecasting errors of a longer duration to affect real activity, suggesting the value of studying specifications like (12).

$$u_t = f\pi_t - \sum_{i=1}^{q} g_i E_{t-i}\pi_t + s_t. \tag{12}$$

In this model, the natural rate property is that simultaneous changes in π_t and $E_{t-i}\pi_t$ by the same amount have no affect on unemployment, i.e., that $f - \sum_{i=1}^{q} g_i = 0$.

As above, we assume that inflation is generated by an autoregressive process, $\pi_t = \rho_1\pi_{t-1} + ...\rho_n\pi_{t-n} + m_t$, which we write as

$$\rho(L)\pi_t = m_t, \tag{13}$$

with $\rho(L) = 1 - \rho_1 L - ... - \rho_n L^n$, where L is the lag operator. Throughout, we require that the change in inflation, $\Delta\pi_t = \pi_t - \pi_{t-1}$, is a stationary random variable with moving average representation $\Delta\pi_t = \mu(L)m_t$, where $\mu(z) = 1 - z/\rho(z)$. We explore two cases: (i) the Lucas-Sargent case in which inflation is stationary, so that $\rho(L)$ has all of its roots outside the unit circle, $\rho(1) \neq 0$ and $\mu(1) = 0$; and (ii) an alternative in which there is a single unit root in $\rho(L)$, so that $\rho(L) = (1 - L)\phi(L)$ with $\phi(L)$ having all of its roots outside the unit circle.

With a unit root in the inflation process, we can derive two results (the details are given in appendix A). First, we find that a unit root in the inflation process implies that the sum of coefficients restriction tested by Solow and Gordon is equivalent to the long-run neutrality restriction. Second, we find that the presence or absence of a unit root is critical to whether the time series data are informative about the consequences of sustained inflation, as variously argued by Sargent [1971] and Fisher and Seater [1993].

The first result is that the reduced form of (12) and (13) is:

$$u_t = \beta_{u\pi}(L)\pi_t + s_t \tag{14}$$

where $\beta_{u\pi}(L)$ is a $(q+n)^{th}$ order polynomial in L. Further, under the unit root assumption, it follows that:

$$\beta_{u\pi}(1) = f - \sum_{i=1}^{q} g_i. \tag{15}$$

That is, if there is a unit root in the inflation process, then the sum of coefficients is informative about the slope of the long-run Phillips curve as suggested by Solow and Gordon.

To develop the second result, let $\bar{\pi}_t$ denote the Beveridge and Nelson [1981] measure of trend inflation, and let $M_t = M_{t-1} + m_t = \sum_{j=1}^{t} m_j + M_o$

denote the sum of the (demand) shocks in the inflation equation (13). Then, $\bar{\pi}_i = \mu(1)M_t$, and (14) can be rearranged as:

$$u_t = [f - \sum_{i=1}^{q} g_i]\mu(1)M_t + \psi(L)m_t + s_t, \tag{16}$$

where $\psi(L)m_t$ is a stationary component of unemployment arising from the demand shifter m_t. In (16) the long-run parameter, $[f - \sum_{i=1}^{q} g_i]$, appears as the coefficient on trend inflation, $\mu(1)M_t = \bar{\pi}_t$. When inflation is stationary, $\mu(1) = 0$, trend inflation is identically zero, and the neutrality parameter $[f - \sum_{i=1}^{q} g_i]$ is not identified in (16). Thus, as stressed by Lucas and Sargent, the relevant experiment-permanent changes in the rate of inflation-are absent from the inflation data and so the long-run Phillips trade-off could not be estimated from (16). In the unit root case, by contrast, it follows that $\mu(1) = 1/\phi(1) \neq 0$, so that the relevant experiments are present in the data: variation in $\bar{\pi}_t$ allows the long-run slope to be determined.

4 Estimating a structural Phillips curve

In this section we investigate the structural Phillips curve and implied trade-offs between inflation and unemployment. For this purpose, we use a bivariate VAR of the form:

$$\Delta u_t = \lambda \Delta \pi_t + \sum_{i=1}^{p} \phi_{u\pi,i} \Delta \pi_{t-i} + \sum_{i=1}^{p} \phi_{uu,i} \Delta u_{t-i} + \epsilon_{st} \tag{17}$$

$$\Delta \pi_t = \delta \Delta u_t + \sum_{i=1}^{p} \phi_{\pi\pi,i} \Delta \pi_{t-i} + \sum_{i=1}^{p} \phi_{\pi u,i} \Delta u_{t-i} + \epsilon_{dt} \tag{18}$$

We will interpret equation (17) as the Phillips Curve. In terms of the Keynesian model discussed in Section 2, equation (17) is the "price" or "mark-up" equation (1) rearranged so that u_t appears on the right-hand side. Thus, the parameter λ in (17) corresponds to $1/a$ in (1), and ϵ_{st} is proportional to the shock in the "price" equation. In terms of the monetarist model, equation (17) is the "supply" equation (3), so that λ in (17) corresponds to f in (3) and ϵ_{st} is the "supply" shock. In either model, ϵ_{dt} in (18) corresponds to the demand shock, and our interest here focuses on the dynamic effects of this shock on both u_t and π_t. Equations (17) and (18) are written in first difference form, so that both u_t and π_t are assumed to be $I(1)$ and not cointegrated. The specification (17)-(18) thus permits us to estimate the "long-run" effects of the disturbances ϵ_{dt} and ϵ_{st} in ways similar to those used by Solow [1969] and Gordon [1970]. However, rather than computing $\partial u/\partial \pi$, we will be concerned with $lim_{k\to\infty}\{[\partial u_{t+k}/\partial \epsilon_{dt}]/[\partial \pi_{t+k}/\partial \epsilon_{dt}]\}$ and thus focus on the relative effects of demand shocks on unemployment and inflation.

The first difference form of this specification is consistent with the stochastic trends evident in the series as displayed in Section 2. More formal statistical tests do not reject the unit root restriction built into (17) and (18): for example, as reported in Table 3, the 95% confidence intervals for the largest autoregressive root are (0.968, 1.007) for the unemployment rate and (0.960, 1.006) for inflation over the full sample period. Table 3 provides these largest root estimates and also comparable information on the estimates for the subperiods 1954–1969 and 1970–1992. In Section 6 we consider the robustness of our primary results to the unit root specification and present results for the model estimated in levels.

Table 3:
Unit Root Statistics

Series	Sample Period	Larg. AR Root	$\hat{\tau}_\mu$	95% Conf. Interval
Unemp.	1954–92	0.97	-2.05	(.97 1.01)
Unemp.	1954–69	0.97	-1.15	(.97 1.02)
Unemp.	1970-92	0.94	-3.12	(.89 1.00)
Infl.	1954–92	0.98	-2.34	(.96 1.01)
Infl.	1954–69	0.97	-1.04	(.97 1.02)
Infl.	1970-92	0.96	-2.09	(.94 1.01)

Note: These results are based on a univariate VAR(12) including a constant term. $\hat{\tau}_\mu$ denotes the t-statistic testing that the sum of AR coefficients is equal to 1. The 95% confidence intervals are constructed from $\hat{\tau}_\mu$ using the procedure developed in Stock (1991).

Evidently, equations (17) and (18) are a set of two dynamic simultaneous equations. Standard results imply that two *a priori* assumptions are required to econometrically identify the parameters and shocks in the equations. Here we consider three different sets of identifying assumptions that are in turn suggested by three interpretations of (17)-(18): (i) the traditional Keynesian interpretation; (ii) the rational expectations monetarist interpretation; and (iii) a real business-cycle interpretation. Each interpretation leads to different identifying assumptions, which in turn lead to different estimates of the equations and shocks. These differences are quantitatively very large and imply very different estimates of the inflation-unemployment trade-off and the corresponding "costs" of disinflation. They also lead to very different historical interpretations of postwar U.S. business cycles.

In all of the interpretations of (17)-(18) we will assume that the disturbances ϵ_{st} and ϵ_{dt} are mutually uncorrelated. Conceptually, this allows us to think of each shock as arising from distinct and independent sources, and means that any contemporaneous correlation between u_t and π_t arises from nonzero values of the parameters λ and δ. Only one additional assumption is necessary to econometrically identify the parameters and shocks in (17)-(18).

4.1 *Three short-run identifications*

The three identifications that we consider differ from one another in their assumed correlation between ϵ_{dt} and u_t. The first identification is suggested by econometric implementations of the traditional Keynesian model (1)-(2) which allowed little contemporaneous feedback between the wage-price block (summarized by (1)) and the IS-LM block (summarized by (2)). This structure implies that one-step-ahead forecast errors in the unemployment rate are dominated by aggregate demand disturbances. At the other extreme, real business-cycle models postulate that movements in real variables such as u_t are perfectly correlated with aggregate supply shocks. There is a large middle ground between these two extreme views; we use an identification suggested by the rational expectation monetarist models of Sargent [1976] and Barro and Rush [1980] which, as we show below, yields results midway between the two extreme views.

The traditional Keynesian (TK) identification: As discussed above, in traditional Keynesian models, changes in prices arising from realizations of ϵ_{st} have little contemporaneous effect on u_t: in the extreme, ϵ_{st} and u_t are contemporaneously uncorrelated. This means that in addition to lagged variables, the contemporaneous value of u_t can be used as an instrument to estimate equation (17); implicitly this defines λ in equation (17) as $\lambda = var(\tilde{u})_t/cov(\tilde{u}_t, \tilde{\pi}_t)$, where $\tilde{}$ denotes an unforecastable component. Equivalently, equation (17) can be estimated by OLS using the reverse regression of π_t onto u_t and relevant lags. In this form, we recognize the "price equation" estimation strategy used by Gordon [1970] and other researchers in the Keynesian tradition.[11] Using our data, this leads to an estimate of $\lambda = -1.56$; throughout we will use this value of λ to represent the traditional Keynesian specification.[12]

[11]This is a stylized characterization of the empirical Phillips curve literature of the 1960s and early 1970s. These researchers often included additional shift variables to account for particular events and transformed the data in a variety of ways to model nonlinearities in the Phillips curve. But, in estimation procedures and in policy evaluations, they treated unemployment as exogenous in wage and price equations. More recent econometric analysis of the Phillips curve (notably, Gordon [1982, 1990b]) allows for correlation between $\tilde{u}_t$ and ϵ_{st} and uses identification procedures like those used in the rational expectations monetarist models.

[12]This estimate of λ is quite imprecise with an estimated standard error of 1.61. To see the source of this imprecision recall (from the usual IV formula) that the estimated standard error for $\hat{\lambda}$ is given by $var(\tilde{u}_t)/\{T[cov(\tilde{u}_t, \tilde{\pi}_t)]^2\}$. In the data, the one-step-ahead

The rational expectations monetarist (REM) identification: In rational expectations monetarist models, researchers also looked for an additional instrument that would allow them to estimate (17). In the "supply" equation interpretation of (17), this instrument is required to be correlated with the unforecastable component of inflation and uncorrelated with the supply disturbance. While empirical researchers in this tradition did not settle on a consensus instrument (or set of instruments), they obtained similar estimates of λ and the resulting Phillips curve trade-off. For example, Sargent [1976] estimated an equation like (17) as part of his larger classical macroeconometric model. In addition to lags, he used money, population and government spending variables as instruments. Barro and Rush [1980] estimated the effects of "unanticipated money growth" on unemployment and the price level, resulting in an implicit instrumental variables estimator of the parameter λ in (17).[13] The implicit estimate of λ from Sargent's analysis is -0.07, while the Barro and Rush estimates ranged from -0.17 to -0.07. In our analysis, we will use -0.07 as the value of λ for the rational expectations monetarist model. As the results will make clear, this seems to be a reasonable intermediate value between the traditional Keynesian and real business-cycle model extremes.[14]

The real business cycle (RBC) identification: In real business-cycle models, the unemployment rate is unaffected by nominal shocks. Thus, inflation does not enter (17), *i.e.*, $\lambda = \phi_{u\pi}(L) = 0$ in (17). We will use the assumption that $\lambda = 0$ as the additional identifying assumption in this interpretation of (17). While this is an arbitrary interpretation of the identifying assumptions in this model (since we could have set λ or any value of $\phi_{u\pi,i}$ equal to zero to achieve identification), this restriction leads to an empirical model with clear RBC characteristics.

---

forecast errors for unemployment and inflation are very weakly correlated, leading to a small estimated value of $[cov(\tilde{u}_t, \tilde{\pi}_t)]$ and a corresponding large standard error for $\hat{\lambda}$. Interestingly, this is what would be expected from a model within which prices move very little on impact in response to a change in aggregate demand. In any event, Section 6 summarizes results for a large range of values of λ.

[13]The implicit instrumental variables estimator from Barro and Rush [1980] is constructed as follows: let b_{um} be the ordinary least squares estimate of the effect of (unanticipated) money on unemployment and let b_{pm} be the corresponding OLS estimate of the effect of (unanticipated) money on the price level (or inflation). Then, the implicit IV estimator is b_{um}/b_{pm}, *i.e.*, it is $cov(u, m)/cov(p, m)$, which is the conventional IV formula. See Appendix B for some additional discussion of how we derive estimates of λ from the studies of Sargent [1976] and Barro and Rush [1980].

[14]R.J. Gordon has suggested to us that "mainstream" is a better label for this identification, since it produces results in which both supply and demand disturbances play a major role in the business cycle, a result accepted by all mainstream economists irrespective of "Keynesian" or "Monetarist" perspectives.

4.2 *Implications of the identifications*

A summary of the results for each of these specifications is given in Figure 4 and Table 4. These results were obtained by estimating (17) using $\lambda = -1.56$ (the TK identification), $\lambda = -0.07$ (the REM identification), and $\lambda = 0.0$ (the RBC identification). The models were estimated using data from 1954:1-1992:12 and included twelve lags and a constant term. The figure shows the estimated impulse responses from the demand shock (panel A); the fraction of the k-step ahead variance of u and π attributed to the demand shock (panel B); and the inflation-unemployment trade-off at different horizons (panel C). This trade-off is $[\partial u_{t+k}/\partial\epsilon_{dt}]/[\partial\pi_{t+k}/\partial\epsilon_{dt}]$, and thus shows the relative effect of a demand shock on unemployment and inflation. (This slope formed as the ratio of the two impulse response functions in panel A.)

It is convenient to begin the discussion with the RBC identifying restriction ($\lambda = 0.0$); the results for this specification are summarized in the first column of panels in Figure 4. In this specification, ϵ_{st} corresponds to the one-step-ahead forecast error in u_t, and ϵ_{dt} is the portion of the forecast error in π_t that is orthogonal to the forecast error in u_t. The impulse response function and variance decomposition show that u_t is essentially Granger causally prior to π_t; that is, lagged values of π_t (and the associated lags of ϵ_{dt}) explain a tiny fraction of future values of u_t. (The F-statistic for the null hypothesis that π_t does not Granger-cause u_t is 1.25 with a p-value of .26.) At lag zero the unemployment-inflation trade-off is zero by assumption for this specification, since $[\partial u_{t+k}/\partial\epsilon_{dt}] = 0$ for $k = 0$. For other values of k, the trade-off is empirically determined, and panel C shows that it is positive but very small. In this sense, the identification yields a picture of the business cycle that is essentially real.

The next column of panels in Figure 3 shows the results for the REM identifying restriction $\lambda = -0.07$. Here, ϵ_{dt} explains roughly 40-50% of the variability of the unemployment rate at all horizons. This shock explains 52% of inflation on impact; 84% after 4 years, and over 95% of the long-run variance. The shock has a large impact effect on inflation and a moderate impact effect on the unemployment rate: hence, there is only a small value of $\partial u/\partial\pi$ on impact. After one year, there is a much larger effect on unemployment and inflation; the trade-off increases to $\partial u/\partial\pi = -.32$ and then falls to -.29 after 4 years. (The standard errors range from .07 at a year to .05 at four years.) The estimated long-run trade-off is -.29 (se = .05), which is not fully consistent with the natural rate hypothesis, but is a small value relative to the 1-for-1 trade-offs reported in the early empirical. Overall, the identification leads to a very monetarist picture of business cycles in that demand disturbances are of substantial importance for economic fluctuations–roughly

FIGURE 4

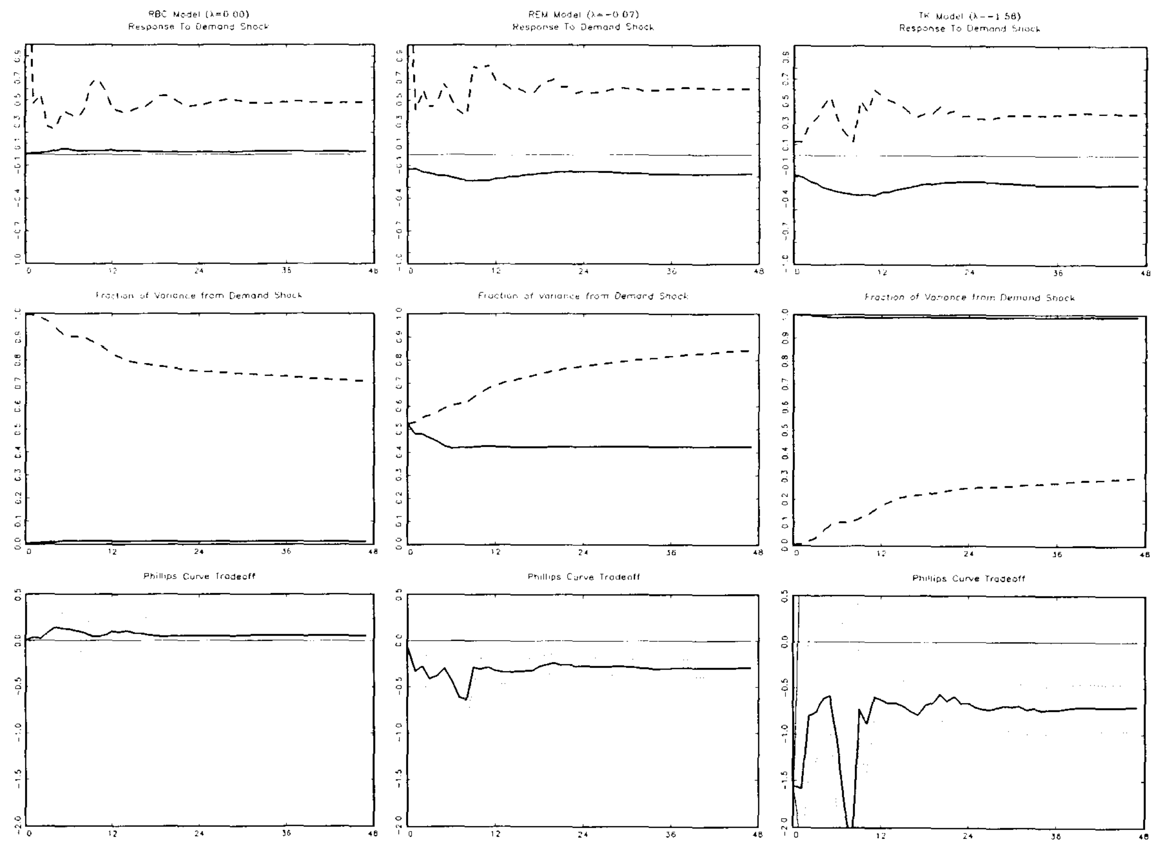

Notes: Panels A and B: Unemployment (solid line), Inflation (dashed line).
Panel C: Point Estimate (solid line), ± 2 Standard Errors (dotted line).

Table 4:
Summary of Results from Structural Models, 1954:1–1992:12

Lag	Demand IRF Unemp.	Demand IRF Infl.	Demand Var. Dec. Unemp.	Demand Var. Dec. Infl.	PC Trade-off
	A. Real Business Cycle Model ($\lambda = 0.00$)				
1	0.00	2.64	0.00	1.00	0.00
	(0.00)	(0.14)	(0.00)	(0.00)	(0.00)
12	0.03	0.58	0.01	0.85	0.09
	(0.04)	(0.13)	(0.02)	(0.04)	(0.11)
24	0.02	0.43	0.01	0.75	0.05
	(0.03)	(0.07)	(0.02)	(0.07)	(0.06)
36	0.03	0.48	0.01	0.73	0.06
	(0.03)	(0.07)	(0.02)	(0.08)	(0.07)
48	0.03	0.48	0.01	0.71	0.06
	(0.03)	(0.07)	(0.02)	(0.09)	(0.07)
∞	0.03	0.48	0.01	0.62	0.06
	(0.03)	(0.07)	(0.02)	(0.11)	(0.07)
	B. Rational Expectations Monetarist Model ($\lambda = -0.07$)				
1	-0.13	1.91	0.52	0.52	-0.07
	(0.01)	(0.07)	(0.04)	(0.04)	(0.00)
12	-0.24	0.84	0.43	0.68	-0.32
	(0.04)	(0.14)	(0.09)	(0.04)	(0.07)
24	-0.15	0.56	0.42	0.77	-0.27
	(0.04)	(0.08)	(0.10)	(0.04)	(0.05)
36	-0.18	0.60	0.42	0.81	-0.30
	(0.03)	(0.08)	(0.10)	(0.04)	(0.05)
48	-0.18	0.61	0.42	0.84	-0.29
	(0.03)	(0.08)	(0.11)	(0.04)	(0.05)
∞	-0.18	0.60	0.42	0.96	-0.29
	(0.03)	(0.08)	(0.11)	(0.03)	(0.05)
	C. Traditional Keynesian Model ($\lambda = -1.56$)				
1	-0.19	0.12	1.00	0.00	-1.56
	(0.37)	(0.23)	(0.01)	(0.00)	(0.00)
12	-0.37	0.61	0.99	0.15	-0.63
	(0.73)	(1.19)	(0.12)	(0.30)	(0.16)
24	-0.24	0.36	0.99	0.25	-0.66
	(0.48)	(0.71)	(0.12)	(0.70)	(0.11)
36	-0.27	0.37	0.99	0.27	-0.74
	(0.55)	(0.73)	(0.12)	(0.82)	(0.13)
48	-0.27	0.38	0.99	0.29	-0.70
	(0.54)	(0.75)	(0.11)	(0.93)	(0.13)
∞	-0.27	0.38	0.99	0.37	-0.72
	(0.53)	(0.74)	(0.11)	(1.36)	(0.12)

Note: Estimated standard errors are shown in parentheses.

40%-and explain nearly all of the variation in inflation.[15]

The last panel shows the results for the TK identifying restriction $\lambda = -1.56$. Recall that this model identifies ϵ_{dt} as the one-step-ahead forecast error in u_t. Since π_t does not Granger-cause u_t, the one-step-ahead forecast error in u_t explains essentially 100% of the variance of u_t at all horizons. Thus, the identifying restriction imposed in the traditional Keynesian interpretation of (17) implies that the unemployment rate is essentially a perfect indicator of demand. In contrast, the demand shock explains little of the short-run variability in π_t (the point estimate at lag 0 is 0.00% and 15% at lag 12) and approximately 40% of the long-run variance of inflation. This interpretation of (17) leads to a trade-off of $\partial u/\partial \pi$ of -1.56 on impact, which falls to -.71 ($se = .12$).

4.3 *Why the short-run identification (λ) matters so much*

A key feature of the empirical results discussed in the prior section is that the assumed short-run effect of inflation on unemployment (λ) substantially affects all of the other features of the dynamic system. Notably, λ dictates the level of the long-run multiplier ($\partial u/\partial \pi$); the sources of business cycles as revealed by estimated decompositions of variance; the shape of the impulse responses, etc.

To gain some intuition for these results, consider the reduced form of the dynamic system (17) and (18):

$$\Delta u_t = a(L)\Delta u_{t-1} + b(L)\Delta \pi_{t-1} + e_{ut} \tag{19}$$

$$\Delta \pi_t = c(L)\Delta u_{t-1} + d(L)\Delta \pi_{t-1} + e_{\pi t}, \tag{20}$$

where all of the lag polynomials contain only positive powers of L (so that only lagged values appear on the left-hand side of (19) and (20)). In this reduced-form system, the forecasting errors e_{ut} and $e_{\pi t}$ are linear combinations of the structural disturbances ϵ_{st} and ϵ_{dt}; specifically since $e_{ut} = D(\lambda\varepsilon_{dt} + \varepsilon_{st})$ and $e_{\pi t} = D(\varepsilon_{dt} + \delta\varepsilon_{st})$, where $D = (1 - \lambda\delta)^{-1}$. Larger values of λ thus imply that there is a larger short-run effect of demand shocks on unemployment. Indeed, in the limiting case with $\lambda \rightarrow \infty$, it follows that shocks to demand and shocks to unemployment are identical.

Summary statistics for the estimated reduced form are given in Table 5. We now consider two aspects of the reduced form which are approximately,

[15] Another implementation of a monetarist identification is that long-run inflation is always and everywhere a monetary (demand) phenomenon, i.e., that inflation is unaffected by ε_{st} in the long run. These sorts of long-run identifications are variously explored in Fisher and Seater [1993], King and Watson [1992], and Roberts [1993]; the latter two papers explicitly consider the trade-off between inflation and unemployment under this identification. For our full sample period, this identification leads to an estimated short-run tradeoff of -0.05, and is thus intermediate to our REM and RBC identifications.

although not exactly, true for our estimates. These are: (i) that the sum of the b coefficients, which we denote $b(1)$, is essentially zero; and (ii) that the individual b_i coefficients are close to zero. Consideration of these two features of the reduced form helps us understand why the selection of λ is so critical to the empirical results discussed in the previous section.

First, when $b(1) = 0$, the long-run Phillips trade-off is monotonically increasing in λ. To see this, solve (19)-(20) for the long-run trends in unemployment and inflation: $\tau_{ut} = a(1)\tau_{ut} + b(1)\tau_{\pi t} + e_{ut}$ and $\tau_{\pi t} = c(1)\tau_{ut} + d(1)\tau_{\pi t} + e_{\pi t}$, and recognize that $lim_{k\to\infty}\partial u_{t+k}/\partial\epsilon_{dt} = \partial\tau_{ut}/\partial\epsilon_{dt}$ and $lim_{k\to\infty}\partial\pi_{t+k}/\partial\epsilon_{dt} = \partial\tau_{\pi t}/\partial\epsilon_{dt}$. Then, by direct calculation:

$$\lim_{k\to\infty}\frac{\partial u_{t+k}/\partial\epsilon_{dt}}{\partial\pi_{t+k}/\partial\epsilon_{dt}} = \frac{[(1-d(1))\lambda + b(1)]}{[(1-a(1)) + \lambda c(1)]}. \tag{21}$$

With the condition $b(1) = 0$ imposed and using the facts that $d(1) < 1$ and $a(1) < 1$, (21) provides a simple characterization of the relationship between λ and the long-run Phillips trade-off. First, at $\lambda = 0$, the long-run slope is zero (which confirms our finding with the RBC interpretation). Second, $\partial u/\partial\pi$ is increasing in λ: a larger short-run negative slope implies a larger long-run negative slope.[16] Hence, if one assumes that inflation has a major short-run effect on unemployment (which is our version of the traditional TK identification), then one also finds that there is a major long-run effect of inflation on unemployment. By contrast, if one assumes that inflation has a small short-run effect, then one also finds little long-run effect.

When all of the $b_i's = 0$, (19) becomes:

$$\Delta u_t = a(L)\Delta u_{t-1} + \{\frac{\lambda}{1-\lambda\delta}\epsilon_{dt} + \frac{1}{1-\lambda\delta}\epsilon_{st}\},$$

where we have written e_{ut} in terms of the structural errors. Two implications follow directly. First, when $\lambda = 0$ (as in the RBC identification), then the demand shock has no effect on unemployment at any horizon. This finding accords with the findings in panel C of Table 4, with minor discrepancies that are associated with the fact that the $b(L)$ is only approximately rather than exactly zero. Second, increases in λ lead to larger effects of demand shocks on unemployment at all horizons as we move to consideration of the impulse responses under the REM and TK identifications. Further, the fraction of k-step-ahead forecast error variance attributable to demand shocks is roughly constant across horizons and is increasing in λ. The large value of λ used in the TK identification, for example, essentially makes the demand shock equal to the unemployment forecast error (as occurs exactly when $\lambda \to \infty$ in the formula above).

[16] The point estimates reported in Table 5 suggest that $c(1) < 0$, which implies that $\partial u/\partial\pi$ does not have any discontinuities for values of $\lambda < 0$.

Table 5:
Summary of the Reduced Form VAR

$$\Delta u_t = a(L)\Delta u_{t-1} + b(L)\Delta \pi_t + e_{u,t}$$

$$\Delta \pi_t = c(L)\Delta u_{t-1} + d(L)\Delta \pi_t + e_{\pi,t}$$

Sums of Coefficients:

Parameter	Estimate	(SE)
a(1)	.36	(.10)
b(1)	.04	(.04)
c(1)	-6.83	(1.41)
d(1)	-4.16	(.59)

Residual Covariance Matrix:

$$sd(e_u) = .190$$
$$sd(e_\pi) = 2.717$$
$$cor(e_u, e_\pi) = -.05$$

Granger-Causality Test Statistics:

Hypothesis	F-Statistic	p-values
b(L) = 0	1.23	(0.26)
d(L) = 0	3.42	(0.00)

Notes: The estimates are constructed from a VAR(12), including a constant, estimated over 1954:1 – 1992:12.

Thus, the near Granger-causal relationship between π_t and u_t means that selection of the short-run identification dictates the relation of demand and supply shocks in economic fluctuations.

4.4 *Measuring the costs of disinflation*

These different models suggest dramatically different costs of disinflation. Table 6 shows the estimated responses of unemployment and inflation to a demand shock that eventually leads to a 1% permanent reduction in inflation. In addition, the table shows the "Sacrifice Ratio" defined as the cumulative annual percentage-point changes in unemployment required to produce this permanent reduction in inflation. Estimates predicated on the TK identification suggest that the unemployment rate rises by 0.9% after 1 year, is still 0.7% higher after 5 years, and that the five-year sacrifice ratio is 3.7. That is, over five years, the cost of a 1% permanent reduction in inflation is a cumulative 3.7% annual percentage-point increase in unemployment. By contrast, the REM identification yields a much smaller change in the unemployment rate and a correspondingly smaller value of the 60-month sacrifice ratio of 1.52. This value is similar to results found by Gordon and King [1982], Mankiw [1990], and Ball [1993] using different identifying assumptions.[17] In the RBC identification, unemployment is essentially exogenous, so that the reduction in inflation has no unemployment cost. All of these estimates can be contrasted to simulation results reported in Eckstein [1981] for the DRI model. The DRI model's cost of disinflation is very large (the sacrifice ratio is 8), more than twice as large as the sacrifice ratio of the Keynesian identification in the bivariate VAR.

4.5 *Interpretations of episodes*

These different identifications also lead to dramatically different interpretations of the postwar business-cycle history in the United States. Figures 5–7 show the 24-month-ahead forecast errors in the unemployment rate, the rate of inflation, and the price level. The price-level forecast error is the

[17]Mankiw [1990] estimates a sacrifice ratio of 1.4 for the Volker disinflation over 1981-85 by assuming that aggregate demand was responsible for the entire decrease in inflation and the increase in unemployment over a 6% natural rate. Ball [1993] estimated average output sacrifice ratios for the U.S. of 2.4 by assuming the movements in trend inflation between "inflation peaks" and "inflation troughs" was attributed to aggregate demand. Dividing his estimate by an Okun's law coefficient of 2 yields an unemployment sacrifice ratio of 1.2. Gordon and King [1982] estimate a 48-month output/inflation sacrifice ratio of 3.0, which again corresponds to an unemployment-inflation sacrifice ratio of approximately 1.5. Their estimates are constructed from a model that determines real and nominal output, given exogenous movements in the money supply.

Table 6:
Sacrifice Ratios for Permanent Reductions in Inflation

A. Results for the Bivariate VAR

Horizon	TK			REM			RBC		
(Months)	u	π	SR	u	π	SR	u	π	SR
0	0.49	-0.31	0.04	0.22	-3.17	0.02	0.00	-5.55	0.00
	(0.10)	(0.07)	(0.01)	(0.03)	(0.40)	(0.00)	(0.00)	(0.77)	(0.00)
12	0.90	-1.42	0.87	0.36	-1.12	0.36	-0.08	-0.88	-0.07
	(0.16)	(0.24)	(0.17)	(0.06)	(0.16)	(0.04)	(0.09)	(0.22)	(0.06)
24	0.64	-0.96	1.60	0.26	-0.96	0.66	-0.05	-0.96	-0.12
	(0.12)	(0.08)	(0.28)	(0.05)	(0.05)	(0.09)	(0.06)	(0.05)	(0.13)
36	0.73	-0.98	2.28	0.30	-0.99	0.94	-0.06	-1.00	-0.18
	(0.13)	(0.04)	(0.40)	(0.04)	(0.02)	(0.14)	(0.07)	(0.02)	(0.20)
48	0.71	-1.01	3.01	0.29	-1.01	1.23	-0.06	-1.00	-0.23
	(0.13)	(0.03)	(0.53)	(0.05)	(0.01)	(0.18)	(0.07)	(0.01)	(0.26)
60	0.71	-1.00	3.71	0.29	-1.00	1.52	-0.06	-1.00	-0.29
	(0.12)	(0.01)	(0.65)	(0.05)	(0.00)	(0.23)	(0.07)	(0.00)	(0.33)

B. Simulations from the DRI Model

Horizon (Years)	u	Core Inflation	SR
0	0.2	0.0	0.2
1	0.9	-0.1	1.1
2	1.4	-0.3	2.5
3	1.6	-0.6	4.1
4	1.8	-0.9	5.9
5	2.1	-1.0	8.0

Notes: TK denotes the model $\lambda = -1.56$, REM the model with $\lambda = -0.07$ and RBC the model with $\lambda = 0$. The results in Panel A are calculated from the impulse responses summarized in Figure 3. They show the estimated responses for unemployment (u) and the inflation rate (π) corresponding to an impulse in ϵ_{dt} that eventually lowers inflation by 1%. The sacrifice ratio (SR) is the accumulated number of annual unemployment percentage points attributed to this shock. Estimated standard errors are shown in parentheses. The results in Panel B are taken from Eckstein (1981, Table 6.2, page 46).

percentage error; since the horizon is 24 months, this means that the average inflation error over the forecast period can be determined by dividing the price forecast error by 2. The figures also show the component of the forecast error associated with realizations of the demand disturbance.

Table 7 summarizes the total and demand shock components of the forecast error for the NBER dated business-cycle peaks and troughs. It is instructive to focus on two episodes on which one may bring some prior knowledge to bear, namely, (i) the recession that began in 73:11 and ended in 75:3; and (ii) the recession that began in 81:7 and ended in 82:11. Many macroeconomists would argue that the former recession was dominated by a "supply shock" in the form of energy price increases, and the latter was dominated by a "demand shock" originating in monetary policy. As expected from the results summarized in Figure 3, the RBC identification associates essentially all movements in u_t to supply shocks in both of these episodes. In contrast, the TK identification attributes essentially all movements in u_t to demand disturbances in both episodes. Further, under this identification, the demand shock explains 25% of the variance in inflation over the entire period. But, in the two episodes of interest, this demand-based theory predicts too much disinflation in 1974 and too little in 1982.

Table 7:
Twenty-four Month Ahead Forecast Errors
and Demand Shock Component
NBER Cyclical Turning Points

	Unemployment				Inflation				Price Level			
	Total	——Demand——			Total	——Demand——			Total	——Demand——		
Date		TK	REM	RBC		TK	REM	RBC		TK	REM	RBC
A. Cyclical Peaks												
57:8	-0.62	-0.77	-1.08	0.15	2.02	1.11	2.31	0.91	5.45	1.52	5.00	3.93
60:4	-2.07	-1.96	0.64	-0.31	0.78	3.61	-0.02	-2.83	0.54	4.56	2.70	-4.02
69:12	-0.32	-0.39	-0.46	0.07	1.82	0.38	0.73	1.44	1.38	0.57	1.35	0.81
73:11	-1.24	-1.46	-1.70	0.22	5.75	2.33	4.74	3.41	3.94	1.70	3.71	2.23
80:1	-0.37	-0.63	-1.76	0.26	8.55	1.10	5.99	7.45	7.69	1.63	6.45	6.05
81:7	1.28	1.37	0.77	-0.09	-1.12	-2.07	-1.95	0.95	-1.59	-2.55	-2.43	1.19
90:7	-0.25	-0.28	-0.03	0.03	2.51	0.72	0.38	1.78	0.10	0.63	0.36	-0.53
B. Cyclical Troughs												
58:4	3.18	3.00	0.51	0.18	1.84	-3.04	0.05	4.88	5.52	-0.70	2.81	6.21
61:2	0.87	0.88	0.64	-0.02	-1.75	-0.60	-1.65	-1.15	-0.38	0.53	-0.49	-0.90
70:11	2.38	2.30	0.70	0.09	0.65	-2.14	-2.10	2.80	0.57	-1.44	-0.96	2.00
75:3	3.37	3.06	0.20	0.32	-1.11	-3.46	-0.22	2.35	6.20	-0.80	3.55	7.00
80:7	1.18	0.85	-0.56	0.34	-1.54	-1.08	0.63	-0.47	6.16	1.44	5.01	4.71
82:11	3.64	3.74	2.84	-0.09	-10.99	-4.63	-10.94	-6.36	-8.25	-3.59	-8.81	-4.66
91:4	1.19	1.20	0.58	-0.01	-3.67	-1.08	-3.54	-2.59	-2.19	-0.66	-2.33	-1.53

Notes: TK denotes the model with $\lambda = -1.56$, REM denotes the model with $\lambda = 0.07$, and RBC denotes the model with $\lambda = 0.09$. Inflation is average of forecast errors (% AR) with ± 1 month of trough. Price level is in percentage points.

Figure 5

24-Month Ahead Forecast Errors
RBC Identification (λ = 0.00)

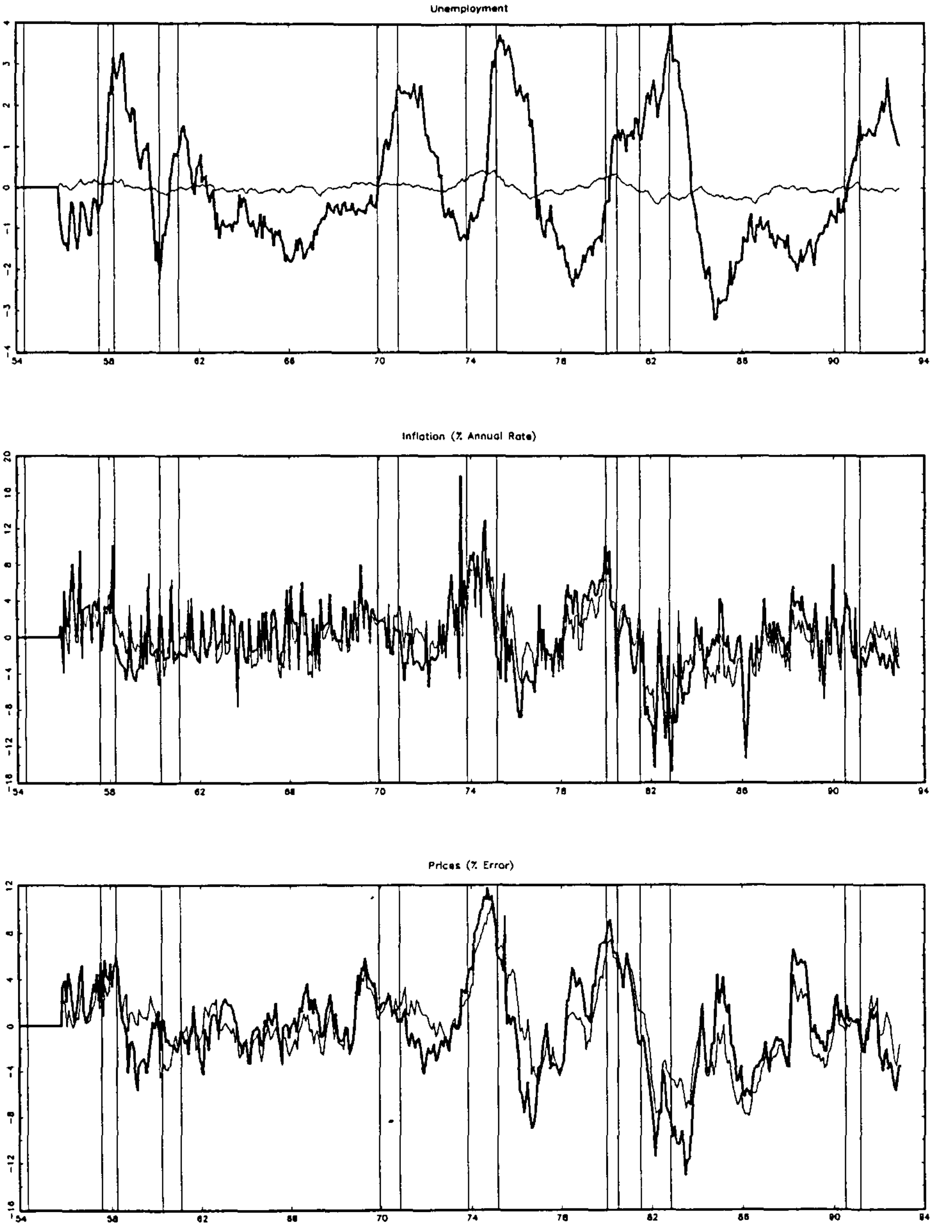

Notes: Total Error (thick line), Demand Shock Component (thin line)

Figure 6

24-Month Ahead Forecast Errors
REM Identification (λ = -0.07)

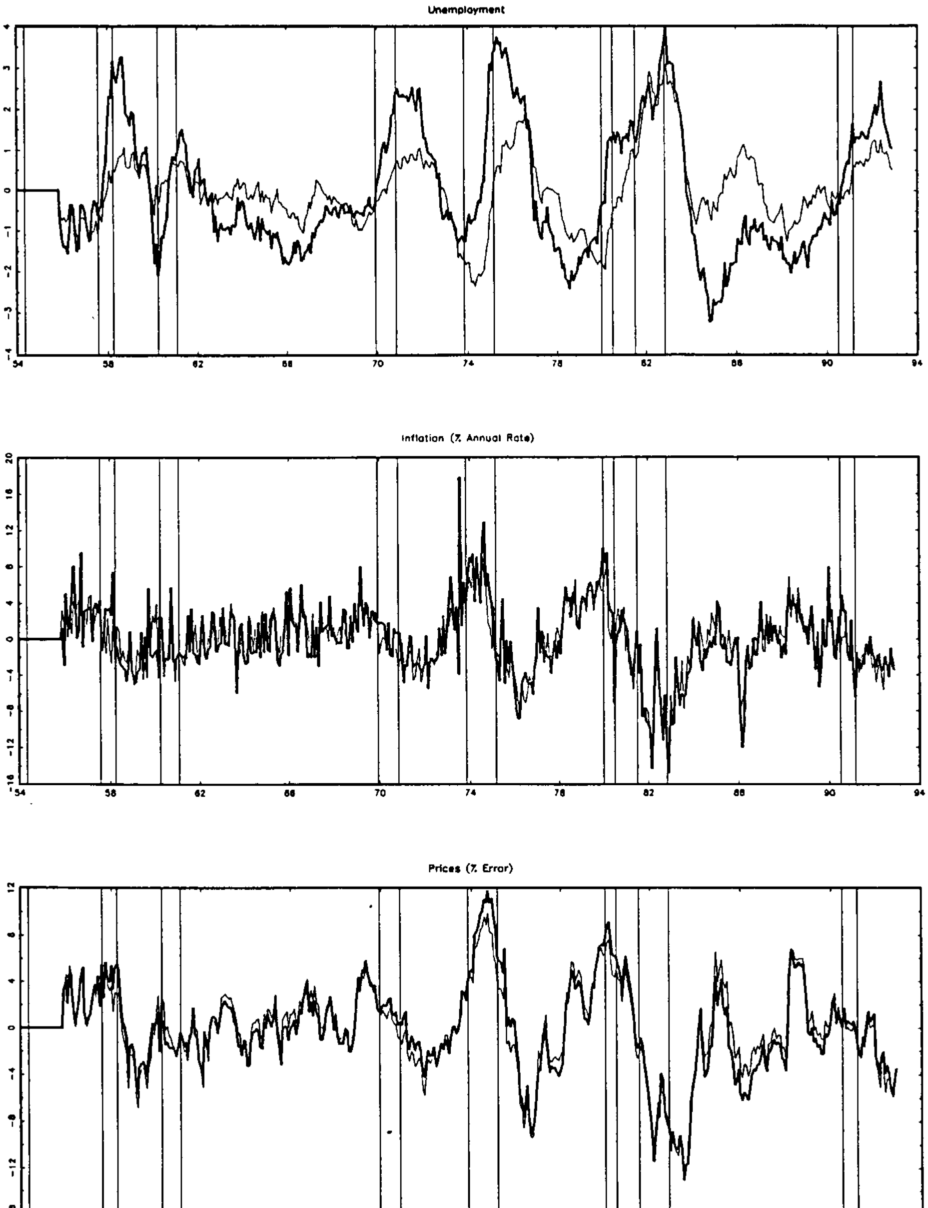

Notes: Total Error (thick line), Demand Shock Component (thin line)

Figure 7

24-Month Ahead Forecast Errors
TK Identification (λ = -1.56)

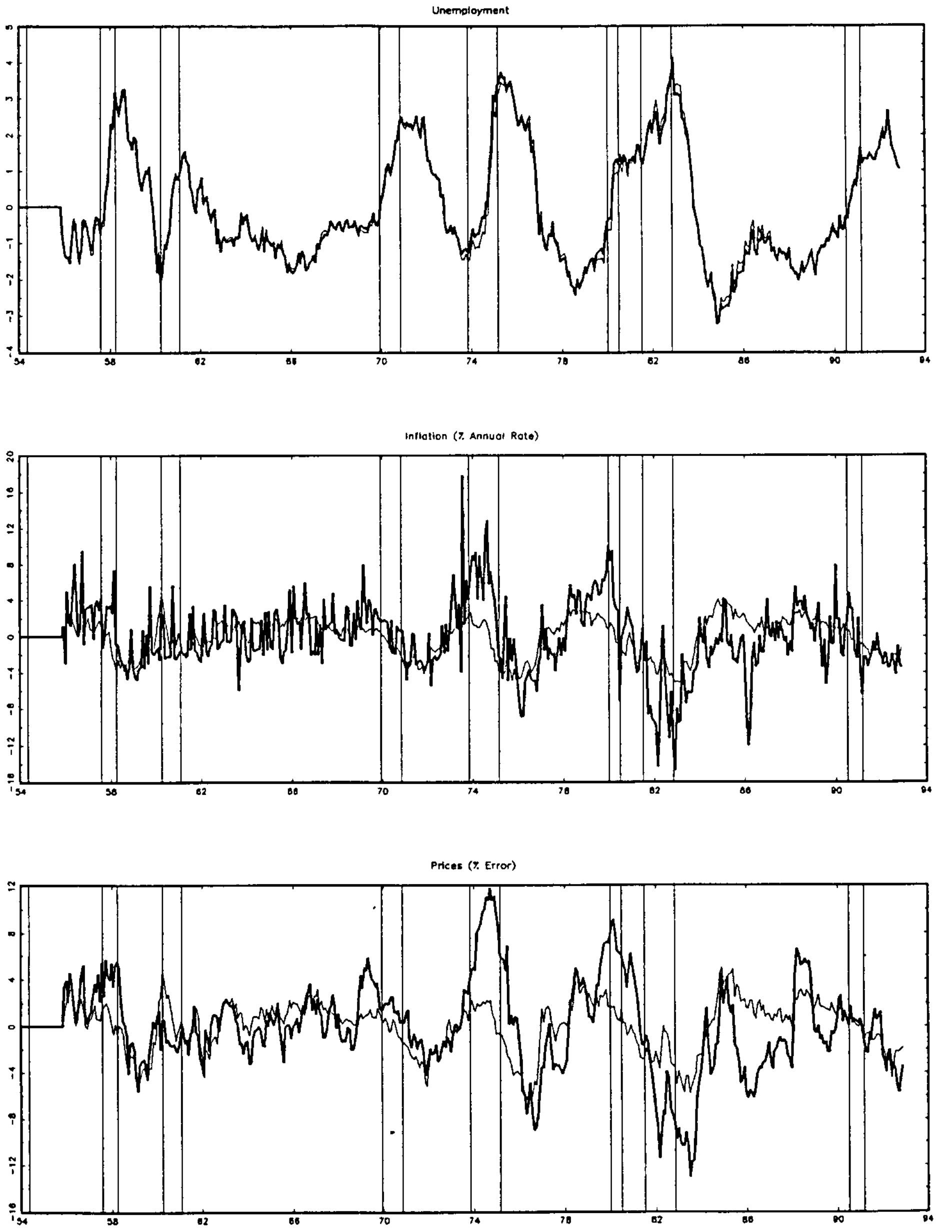

Notes: Total Error (thick line), Demand Shock Component (thin line)

The REM identification yields an interpretation of the U.S. postwar business cycle between these two extremes. The demand shock accounts for a negligible part of the unemployment that occurred during the 1974 recessions, but the bulk of unemployment during the 1982 recession. That is, the total 24-month-ahead forecast error for unemployment in 1975:3 represents the "surprise" in unemployment as of 1973:3: it was 3.37% and the monetarist model indicates that only .20% of this increase was attributable to demand. Comparably, the forecast error for unemployment in 1982:11 is the surprise in unemployment as of 1980:11: it is 3.64% and the monetarist model indicates that 2.84% of this was demand-induced. This model also indicates that the bulk of the surprise disinflation of the 1982 recession was due to demand factors and that there was an important effect of supply shocks on inflation and real activity during the 1974 recession.

4.6 *Summary of findings*

In summary, this section has demonstrated that different ways of interpreting the dynamic correlations between unemployment and inflation lead to radically different estimates of the unemployment-inflation (Phillips curve) trade-off, the costs of disinflation, and the interpretation of the U.S. postwar business cycle. Of course, since our versions of these different models are "just identified" in a econometric sense, they each fit the data on unemployment and inflation equally well. Additional information-economic theory or knowledge of the source of changes in unemployment or inflation in specific episodes–is needed to discriminate between the models. We think that most macroeconomists are likely to find the REM identification most compelling: its long-run predictions square better with the natural rate theory, it provides a balanced decomposition of the influence of supply and demand shocks on economic fluctuations, and it performs better in explaining episodes in which prior knowledge can plausibly be applied.

5 Stability of the Phillips curve

In this section we investigate the stability of the bivariate relation between unemployment and inflation. We do this in three ways. First, we examine the stability of the reduced form VAR using a variety of tests for time-varying coefficients. Second, we examine the stability of forecasting equations and forecast performance. Finally, we compare estimates of the structural models and their implied Phillips curve trade-off over the 1954-69 and 1970-92 periods.

Table 8:
Stability Tests for the Autoregressions

A. Chow Tests For A Single Break in 1970 – VAR

Equation:	Wald Statistic	(df)	P-Value
Unemployment Equation	41.3	(25)	.021
Inflation Equation	47.9	(25)	.004
Both Equations	88.7	(50)	.001

B. Chow Tests For A Single Break in 1970
Univariate Autoregressions

Unemployment Equation	16.3	(13)	.231
Inflation Equation	28.4	(13)	.008

C. Alternative Tests For Instability In Autoregressions

	Stability Test						
Specification	Nyblom	PK1	PK2	Q	AP1	AP2	Qdate
Unem. (Univ.)	-	-	-	**	-	*	59:3
Unem. (Biv.)	**	-	-	***	***	***	59:3
Infl. (Univ.)	-	-	-	***	***	***	74:8
Infl. (Biv.)	-	-	-	***	***	***	74:8

Notes: All regressions contained a constant term and twelve lags. The Wald statistics in Panels A and B allowed difference error covariance matrices in the two sub-samples. The entries in Panel C represent: not significant at the 10% level (-), and significant at the 10% level (*), 5% level (**), and 1% level (***). Nyblom denotes Nyblom's (1989) test, robustified as in Hansen (1991); PK1 and PK2 are the CUSUM and CUSUM^2 tests from Ploberger and Kramer (1992). Q is the Quandt (1960) likelihood ratio test for discrete change in coefficients at an unknown time, calculated as the maximum of the standard Wald test statistic calculated over all dates in the middle 70% of the sample. AP1 and AP2 are the mean Wald and mean exponential Wald tests over all possible break dates in the middle 70% of the sample period. (See Andrews, Lee, and Ploberger [1992].) Q date is the date corresponding to the maximum for the Quandt statistic.

5.1 *Stability of the reduced form*

Table 8 summarizes results from a variety of time-varying coefficient tests. Panels A and B contain results from split-sample Chow tests, using a break data of 1970:1. In panel A we present tests for each of the two equations in the unemployment-inflation VAR and for the system as a whole. The null hypothesis of stability is rejected, and more instability is evident in the inflation equation than in the unemployment equation. Panel B presents tests for stability of univariate autoregressions for unemployment and inflation. Stability is strongly rejected for the inflation process but not for the unemployment process.

Our choice of 1970 as a break date is predicated in part on literature from the early 1970s documenting shifts in the Phillips curve, and so statistical inference in panels A and B suffers from pretesting problems. These problems are remedied in panel C which summarizes results from a variety of statistical tests that are not predicated on a specific break date. The first three tests, labeled Nyblom, PK1, and PK2 are OLS versions of tests developed for stochastically varying regression coefficients (e.g., the CUSUM and CUSUM-squared tests of Brown, Durbin and Evans [1975]). The final tests, labeled Q, AP, and AP1 are modifications of the Wald (Chow) split-sample test for the case of an unknown break date. The first is the Quandt (1960) test, which is formed as the maximum of the split-sample Wald tests over all possible break dates. (Here, the maximum is chosen over all possible dates in the middle 70% of the sample.) The AP1 and AP2 tests are the Andrews-Ploberger average and average exponential Wald tests over the same possible break dates. (See Andrews, Lee and Ploberger [1992]). Since all of these tests have different nonstandard null distributions, we do not report the value of the statistics. Instead, the table lists the statistics that are significant at the 1%, 5%, and 10% levels.

These statistics tell much the same story as the split-sample Chow tests in panels A and B. First, there is significant evidence of a shift in the inflation process. Moreover, it appears as if this shift occurred in the early 1970s. There is also evidence of instability in the unemployment equation in the VAR. For this equation, the Quandt test finds a maximum of the Wald statistic in 1959:3; yet the value of the statistic in the early 1970s is nearly as large as the 1959:3 value. Thus, these tests suggest a shift in the VAR occurring around 1970. Finally, there is limited evidence of instability in the univariate autoregression of unemployment.

5.2 *Stability of forecasting models and performance*

Table 9 and Figure 8 examine the stability of forecasting performance over the sample period. In Table 9, the root mean square forecasting error (RMSE) is shown for three forecasting models over three periods, and for four forecasting horizons. For example, panel A shows the results for one-step-ahead forecasts. The first row of the panel shows results for a VAR with coefficients estimated over 1954-69; the next row shows results from a VAR estimated over 1970-92, and the final row shows results over 1954-92. The first two columns show the forecasting performance over 1954-69 for unemployment and inflation, respectively; the next two columns show that forecasting performance over 1970-92; and the last two columns summarize the forecasting performance over the entire 1954-92 period.

Two major conclusions follow from this table. First, the 1954-69 model forecasts the 1970-92 period nearly as well as the 1970-92 model, and the 1970-92 model forecasts the 1954-69 period nearly as well as the 1954-69 model. For example, looking at the one-step-ahead forecasts over 1970-92, the RMSEs for the unemployment rate are 0.21 for the 1954-69 model versus 0.17 for the 1970-92 model; the corresponding RMSEs for inflation are 3.18 and 2.56. While the differences in RMSEs are statistically significantly different from one another, they do not signal an overwhelming failure of the 1954-69 model relative to the 1970-92 model for this period. This conclusion obtains for both periods and all forecasting horizons considered in the table. This relative stability of the forecasting models is evident in Figure 8 which plots the 24-month-ahead forecast error for both the 1954-69 and the 1970-92 models. While some differences stand out (notably the unemployment forecasts in 1956 and 1960), the forecast errors for the two models are remarkably similar.

The second main conclusion from Table 9 is that forecasts for horizons 12 months and longer were significantly less accurate in the 1970-92 period than in the 1954-69 period. This result obtains regardless of the VAR model used for forecasting. For example, using the full sample VAR, the 24-month-ahead RMSE was 2.71 over 1954-69 and nearly doubled to 4.59 over the 1970-92 period. This forecast deterioration is also evident for the unemployment rate at the 24-month horizon. For horizons less than twelve months, there is little apparent deterioration in the forecasts for either series.

These forecasting results suggest two broad conclusions about changes in the impulse-propagation mechanism characterizing the unemployment-inflation VAR. First, since there is no deterioration in the short-run forecasts, there appears to have been little change in the variance of shocks across the two periods, at least when shocks are limited to the bivariate process examined here. Second, the deterioration in medium to longer-run forecasts in the second period suggests more persistent effects of shocks in the latter period. (By an increase in persistence we mean an increase in the size of moving average coefficients at long lags. This increase in the magnitude of moving average coefficients leads directly to an increase in medium to longer-run

Table 9:
Root Mean Squared Error For Models Estimated
Over Different Sample Periods

A. 1-Step-Ahead Forecast Error RMSE

	Forecasting Period					
	1954:1-1969:12		1970:1-1992:12		1954:1-1992:12	
Estimation Period	Unemp.	Infl.	Unemp.	Infl.	Unemp.	Infl.
1954:1-1969:12	0.182	2.459	0.208	3.177	0.198	2.904
1970:1-1992:12	0.214	3.002	0.174	2.560	0.192	2.750
1954:1-1992:12	0.192	2.643	0.180	2.643	0.185	2.750

B. 6-Step-Ahead Forecast Error RMSE

1954:1-1969:12	0.560	2.515	0.748	3.203	0.677	2.940
1970:1-1992:12	0.700	2.666	0.609	3.125	0.648	2.945
1954:1-1992:12	0.605	2.557	0.639	3.150	0.625	2.945

C. 12-Step-Ahead Forecast Error RMSE

1954:1-1969:12	0.905	2.625	1.344	3.804	1.184	3.371
1970:1-1992:12	1.182	3.033	1.104	3.527	1.137	3.333
1954:1-1992:12	0.977	2.730	1.178	3.597	1.100	3.333

D. 24-Step-Ahead Forecast Error RMSE

1954:1-1969:12	1.050	2.680	1.851	4.747	1.573	4.029
1970:1-1992:12	1.277	2.881	1.680	4.482	1.528	3.906
1954:1-1992:12	1.101	2.708	1.751	4.592	1.519	3.906

Notes: The entries in the table refer to the root mean square forecast error for unemployment and inflation for the forecasting period shown. For example, the forecast error dated 1954:1 is the forecast error for 1954:1 using forecasts computed in earlier periods. The forecasts were formed using VAR(12) models (including a constant) estimated over the periods given in the first column of the table.

FIGURE 8 24-Month Ahead Forecast Errors

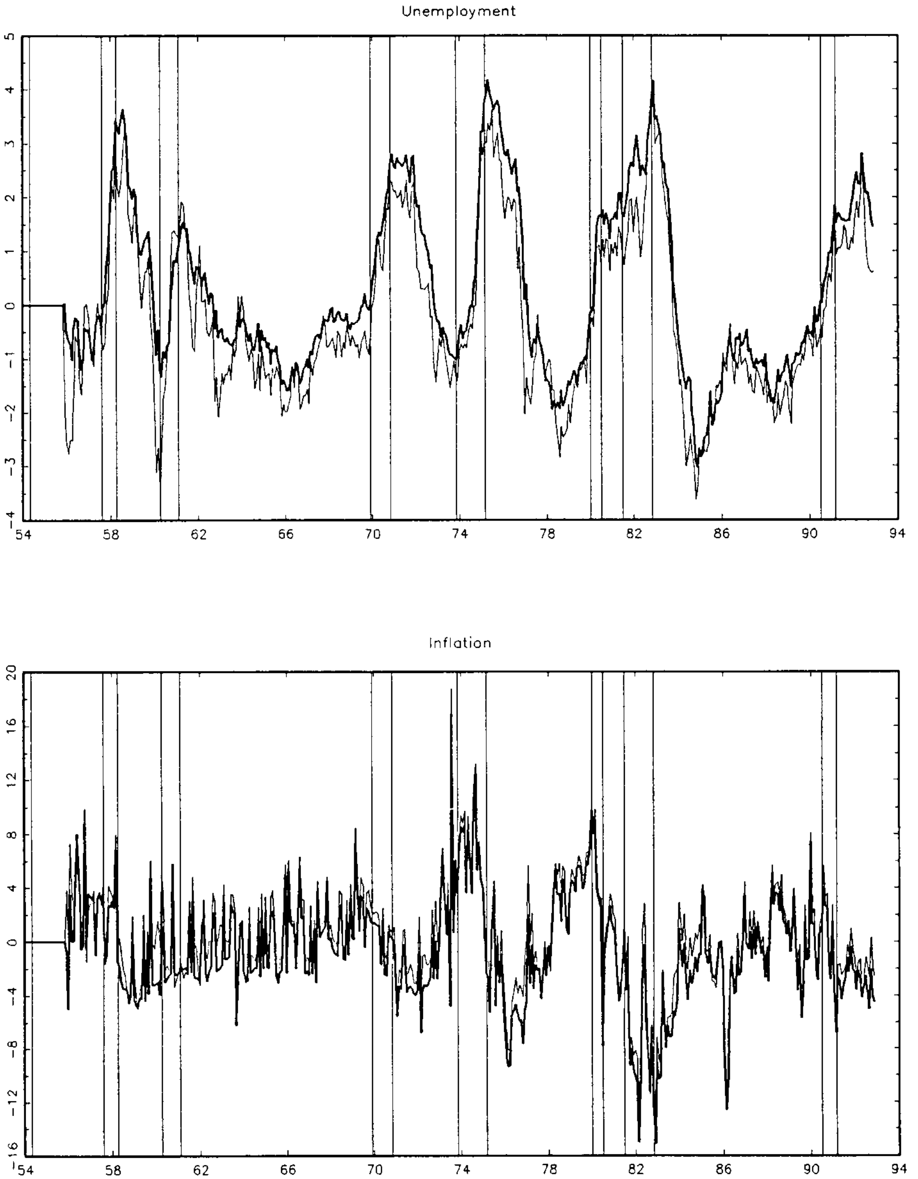

Notes: 1954-69 VAR (thick line), 1970-72 VAR (thin line) forecast error variances.)

In spite of these changes in the post-1969 period, the 1954-69 VAR forecasts perform nearly as well as the 1970-92 VAR, even at the 24-month horizon. This means that the intrinsic forecast error in the model dominates the error arising from model misspecification. This is shown in Table 10 which shows population RMSEs for both the 1954-69 and 1970-92 forecasting models, assuming first that the data are generated by the 1954-69 model and then by the 1970-92 model. The conclusions from these population RMSEs are the same as the sample RMSEs in Table 9: the 1954-69 model forecasts data generated by the 1970-92 model nearly as well as the optimal forecasting model.

Table 10:
Population Standard Errors For Forecasts From
Estimated Sub-Sample Models

	DGP: 54-69 Model		DGP: 70-92 Model	
Forecast	Forecast Model		Forecast Model	
Horizon	54-69	70-92	54-69	70-92
	A. Unemployment Rate			
1	0.18	0.21	0.21	0.17
6	0.58	0.69	0.73	0.61
12	0.94	1.14	1.28	1.11
24	1.12	1.21	1.77	1.72
	B. Inflation Rate			
1	2.46	2.98	3.17	2.56
6	2.55	2.65	3.19	3.12
12	2.77	3.02	3.79	3.61
24	3.05	3.18	4.61	4.53

Notes: Each entry in the table is the population standard error of the forecast constructed from models estimated over either 1954-69 or 1970-92. The first two columns of entries assume that the data are generated by the 54-69 model (so that the 54-69 model is the optimal forecasting model and the 70-92 model is sub-optimal); the last two columns assume that the data are generated by the 70-92 model (so that the 70-92 model is the optimal forecasting model and the 54-69 model is sub-optimal.)

Table 11
Sub-Sample Stability of Phillips Curve Models
A. Traditional Keynesian Model ($\lambda = -1.56$)

Lag	Demand Shock Impulse Response 54-69	Demand Shock Impulse Response 70-92	Forecast Error RMSE's Demand Shock 54-69	Forecast Error RMSE's Demand Shock 70-72	Forecast Error RMSE's Supply Shock 54-69	Forecast Error RMSE's Supply Shock 70-72
			A.1 Unemployment			
1	-0.18	-0.17	0.18	0.17	0.01	0.01
	(0.31)	(0.61)	(0.31)	(0.61)	(0.01)	(0.01)
12	-0.28	-0.41	0.93	1.09	0.07	0.22
	(0.47)	(1.44)	(1.57)	(3.86)	(0.41)	(0.20)
24	-0.16	-0.33	1.12	1.67	0.09	0.41
	(0.27)	(1.17)	(1.89)	(5.92)	(0.37)	(0.27)
36	-0.22	-0.32	1.32	1.99	0.10	0.51
	(0.37)	(1.12)	(2.24)	(7.06)	(0.38)	(0.33)
48	-0.19	-0.33	1.49	2.29	0.11	0.59
	(0.32)	(1.15)	(2.52)	(8.10)	(0.38)	(0.39)
∞	-0.20	-0.32				
	(0.34)	(1.15)				
			A.2 Inflation			
1	0.12	0.11	0.12	0.11	2.46	2.56
	(0.20)	(0.39)	(0.20)	(0.39)	(0.13)	(0.21)
12	0.25	0.80	0.75	1.78	2.67	3.14
	(0.46)	(2.74)	(1.26)	(6.15)	(0.20)	(0.36)
24	0.09	0.63	0.95	2.87	2.90	3.51
	(0.17)	(2.23)	(1.60)	(9.93)	(0.27)	(0.64)
36	0.18	0.54	1.08	3.47	3.10	3.86
	(0.30)	(1.90)	(1.83)	(12.04)	(0.31)	(0.82)
48	0.14	0.57	1.21	3.99	3.30	4.20
	(0.25)	(1.93)	(2.06)	(13.80)	(0.35)	(0.95)
∞	0.15	0.57				
	(0.27)	(1.98)				

A.3 Phillips Curve Trade-off

Lag	Sample Period 1954-69		1970-92	
0	-1.56	(0.00)	-1.56	(0.00)
12	-1.15	(0.94)	-0.54	(0.14)
24	-1.26	(0.55)	-0.54	(0.14)
36	-1.31	(0.45)	-0.58	(0.21)
48	-1.29	(0.39)	-0.57	(0.15)
60	-1.30	(0.40)	-0.56	(0.15)
72	-1.29	(0.39)	-0.57	(0.12)
∞	-1.30	(0.39)	-0.57	(0.13)

Note: Estimated standard errors are shown in parentheses.

Table 11 (Continued)
Sub-Sample Stability of Phillips Curve Models
B. Rational Expectations Monetarist Model ($\lambda = -0.07$)

	Demand Shock Impulse Response		—— Forecast Error RMSE's —— Demand Shock		Supply Shock	
Lag	54-69	70-92	54-69	70-72	54-69	70-72
			B.1 Unemployment			
1	-0.12	-0.13	0.12	0.13	0.13	0.12
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
12	-0.21	-0.23	0.71	0.64	0.61	0.91
	(0.05)	(0.06)	(0.12)	(0.12)	(0.12)	(0.16)
24	-0.12	-0.17	0.86	0.92	0.72	1.45
	(0.03)	(0.07)	(0.17)	(0.23)	(0.17)	(0.32)
36	-0.17	-0.17	1.02	1.09	0.85	1.75
	(0.04)	(0.05)	(0.19)	(0.30)	(0.19)	(0.40)
48	-0.15	-0.17	1.15	1.24	0.95	2.01
	(0.04)	(0.05)	(0.23)	(0.35)	(0.22)	(0.47)
∞	-0.15	-0.17				
	(0.03)	(0.06)				
			B.2 Inflation			
1	1.75	1.86	1.75	1.86	1.72	1.76
	(0.08)	(0.10)	(0.08)	(0.10)	(0.17)	(0.24)
12	0.31	1.04	2.13	3.09	1.77	1.87
	(0.18)	(0.19)	(0.14)	(0.32)	(0.18)	(0.22)
24	0.30	0.75	2.43	4.10	1.83	1.94
	(0.07)	(0.15)	(0.21)	(0.53)	(0.21)	(0.23)
36	0.34	0.71	2.68	4.81	1.89	1.95
	(0.06)	(0.13)	(0.26)	(0.68)	(0.24)	(0.26)
48	0.33	0.74	2.92	5.45	1.94	1.96
	(0.06)	(0.13)	(0.32)	(0.81)	(0.28)	(0.29)
∞	0.33	0.74				
	(0.06)	(0.13)				

B.3 Phillips Curve Trade-off

	Sample Period			
Lag	1954-69		1970-92	
0	-0.07	(0.00)	-0.07	(0.00)
12	-0.62	(0.35)	-0.24	(0.07)
24	-0.40	(0.10)	-0.24	(0.07)
36	-0.50	(0.10)	-0.23	(0.06)
48	-0.45	(0.09)	-0.23	(0.06)
60	-0.47	(0.09)	-0.23	(0.06)
72	-0.46	(0.09)	-0.23	(0.06)
∞	-0.47	(0.09)	-0.23	(0.06)

Note: Estimated standard errors are shown in parentheses.

Table 11 (Continued)
Sub-Sample Stability of Phillips Curve Models
C. Real Business-Cycle Model ($\lambda = 0.00$)

Lag	Demand Shock Impulse Response 54-69	Demand Shock Impulse Response 70-92	Forecast Error RMSE's — Demand Shock 54-69	Demand Shock 70-72	Supply Shock 54-69	Supply Shock 70-72
			C.1 Unemployment			
1	0.00	0.00	0.00	0.00	0.18	0.17
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.01)
12	-0.04	0.11	0.12	0.26	0.93	1.08
	(0.05)	(0.06)	(0.12)	(0.12)	(0.12)	(0.15)
24	-0.02	0.11	0.17	0.48	1.11	1.65
	(0.02)	(0.06)	(0.17)	(0.25)	(0.17)	(0.30)
36	-0.03	0.10	0.19	0.59	1.31	1.97
	(0.03)	(0.05)	(0.20)	(0.31)	(0.19)	(0.40)
48	-0.03	0.10	0.22	0.68	1.48	2.26
	(0.03)	(0.05)	(0.23)	(0.35)	(0.23)	(0.46)
∞	-0.03	0.10				
	(0.03)	(0.05)				
			C.2 Inflation			
1	2.46	2.55	2.46	2.55	0.08	0.22
	(0.13)	(0.21)	(0.13)	(0.21)	(0.17)	(0.15)
12	0.21	0.64	2.69	3.10	0.68	1.85
	(0.18)	(0.19)	(0.16)	(0.30)	(0.17)	(0.34)
24	0.35	0.40	2.93	3.43	0.83	2.96
	(0.08)	(0.13)	(0.21)	(0.41)	(0.24)	(0.63)
36	0.33	0.44	3.15	3.75	0.93	3.59
	(0.06)	(0.12)	(0.26)	(0.53)	(0.29)	(0.82)
48	0.33	0.46	3.35	4.07	1.04	4.12
	(0.06)	(0.12)	(0.31)	(0.64)	(0.35)	(0.96)
∞	0.33	0.45				
	(0.06)	(0.12)				

C.3 Phillips Curve Trade-off

Lag	Sample Period 1954-69		1970-92	
0	0.00	(0.00)	0.00	(0.00)
12	-0.17	(0.21)	0.25	(0.18)
24	-0.06	(0.07)	0.28	(0.23)
36	-0.10	(0.10)	0.21	(0.15)
48	-0.08	(0.08)	0.22	(0.16)
60	-0.09	(0.09)	0.23	(0.16)
72	-0.09	(0.09)	0.22	(0.16)
∞	-0.09	(0.09)	0.22	(0.16)

Note: Estimated standard errors are shown in parentheses.

5.3 *Stability of structural models*

The change in the VARs propagation mechanism uncovered by the forecasting comparison suggests potential instability in the Phillips curve trade-off. This possibility is examined in Table 11 for the three identifying restrictions discussed in Section 4. Panel A presents results for the TK identification. Consistent with the forecasting comparisons, there is little change in the very short-run properties of the model. The point estimates of the impulse responses suggest an increase in the persistence of demand disturbances: at the 48-month horizon the impulse results for unemployment increases by 50% in the latter period, and the impulse response for inflation increases by a factor of 4. This change in the relative persistence across unemployment and inflation reduces the estimated long-run Phillips trade-off in the 1970-92 period. The trade-off falls from -1.3 in the early period to -0.6 in the latter period. Standard errors are large, however, and the t-statistic for this change is only -1.8. The model's supply shocks have more persistence effects on the second period, and this is particularly true for their effect on unemployment. Yet, they still explain only a small fraction of unemployment forecast errors: 6% at the 48-month horizon. Panel B shows the results for the REM identification. In this model, the response of unemployment to demand shocks is remarkably stable across the two periods for all horizons. The increased persistence in the inflation rate is accounted for almost entirely by demand shocks. Again, the change in the relative persistence of demand shocks across unemployment and inflation leads to a change in the estimated Phillips trade-off; in this identification it falls from -.47 in the early period to -.23 in the latter period. Finally, panel C shows the results for the RBC identification; in both periods this model behaves like the TK model with the interpretation of the shocks reversed.

Taken together, these results point to a decline in the effect of demand shocks on unemployment relative to their effect on inflation. In both the TK and REM identifications, the point estimates suggest this relative trade-off has decreased by approximately 50% in the second part of the sample. As usual in work with VARs, these results are tempered by large standard errors and a resulting lack of statistical precision.

6 Robustness of results

In the previous section we provided a detailed examination of the effect of sample periods on estimates of the Phillips trade-off. In this section we consider the effects of other changes in our basic specification. We begin by considering uncertainty in the value of λ, and present results for a wide range of values of λ. We then consider the robustness of the main findings to: (i) changes in the lag length in the VAR, (ii) relaxation of the unit root constraint, (iii) increases in the sampling interval from monthly to quarterly, (iv) changes in the measure of the aggregate price level, and (v) incorporation of additional indicators of aggregate supply disturbances in the model.

In the previous sections we have discussed structural VAR estimates conditioned on three values of λ that served to identify the VAR. Yet, as discussed in Section 3, there is some uncertainty in the precise value of λ most indicative of the TK and REM models. Specifically, the standard error associated with the estimate of λ constructed from the TK identification is large (see footnote 12), and the point estimates or λ from the studies of Barro and Rush [1980] and Sargent [1976] ranged from -.07 to -.17, and were constructed using quarterly rather than monthly data (see Appendix B). We now discuss how uncertainty about this parameter affects the conclusions reached in the previous sections. Table 12 summarizes results for values of λ ranging from 0.0 to -3.0, using both monthly and quarterly data. Panel A of the table shows the estimated Phillips trade-off at impact and for the 1-year and the 3-year horizon. Panel B shows the resulting estimates of the fraction of the forecast errors attributable to the identified demand shock. Looking first at the results for the monthly data, the results change little as λ varies from -0.5 to -3.0. For example, as λ increases from -3.0 to -0.5, the estimated 3-year trade-off increases from -.77 to -.66 and the identified demand shock continues to be nearly perfectly correlated with unemployment. On the other hand, the results change dramatically as λ varies from -.15 to -0.03, with estimated 3-year Phillips trade-offs changing from -.46 to -.15 and correlation between the identified demand shock and the unemployment forecast error falling from .85 to .16. Table 12 puts our choice of λ in perspective: the TK identification (λ = -1.56) equates ϵ_{dt} with the forecast error in unemployment; the RBC identification (λ = 0.00) equates ϵ_{st} with the forecast error in unemployment; and the REM identification (λ = -0.07) attributes 50% of the variance in unemployment forecast errors with ϵ_{dt} and 50% with ϵ_{st}. Table 12 also includes results for quarterly data for the range of point estimates of λ constructed from Barro and Rush [1980]. The quarterly results with λ = -.07 correspond roughly to the monthly results with λ = -.03, and the quarterly results with λ = -.17 correspond roughly to the monthly results with λ = -.07. Thus, our choice of λ = -.07 corresponds to the upper range of the point estimates from Barro and Rush [1980].

Tables 13 and 14 summarize results for eight different specifications of an empirical model; Table 13 shows the estimated Phillips trade-off and Table 14 shows the forecast error variance decomposition. The first row of each panel shows the baseline specification used in Sections 4 and 5. The other seven rows show results from modifications to this baseline specification. Specification 2 relaxes the unit root constraint, and presents results for the VAR estimated using the levels of inflation and the unemployment rate.

While this specification makes it impossible to calculate the long-run Phillips trade-off, the impact and medium-run trade-offs can be calculated. The only notable difference between the results for this specification and the baseline specification is the decrease in precision of estimates at the 3-year horizon when λ = -1.56. Specification 3 modifies the baseline model by increasing the number of lags in the VAR from 12 to 18. The results are robust to this change.

The high-frequency variability in the inflation rate evident in Figure 1 suggests that the baseline estimates might be contaminated by measurement error in the index of the price level. The next three specifications investigate this possibility. Specifications 4 and 5 use quarterly averages of the price index and unemployment rate to help attenuate any measurement errors in the levels of these series. (Specification 4 uses 4 quarterly lags in the VAR and specification 5 uses 6 lags.) Specification 6 replaces the quarterly consumer price index with the quarterly gross domestic product price deflator. The results from these three specifications are very similar to one another, and to the baseline monthly results when λ = -1.56 or λ = 0.0. They differ from the baseline monthly results when λ = -0.07, and as discussed above are similar to the results for a monthly model with λ = -0.03.

Finally, specifications 7 and 8 add an additional supply indicator to the quarterly and monthly specifications. Specifically, following Gordon (1982, 1990b) we add a measure of the relative price of food and energy. Letting π_t^{pe} denote the inflation rate for food and energy, the VAR is now specified with $(\pi_t^{pe} - \pi), \Delta u_t$ and $\Delta \pi_t$, and is identified with the additional assumption that $(\pi_t^{pe} - \pi_t)$ is contemporaneously exogenous (or, equivalently, ordered first in a Wold causal chain).[18] As shown in the table, the baseline results are robust to this modification.

7 Summary and conclusions

In this paper, we study the postwar U.S. Phillips correlations and Phillips curve. That is, we consider the joint (bivariate) time series behavior of U.S. inflation and unemployment. We use monthly data over the postwar period, focusing on two subperiods (1954-1969 and 1970-present) as well as the full period. The results of this econometric investigation can usefully be broken into two parts: time series interactions of inflation and unemployment; and results from bivariate structural models.

We use two econometric methods to determine some central features of

[18]The food and energy price index is formed as a weighted average of the food and energy components of the crude material producer price index (PPI). The weights are the relative importance of these indexes in the December 1992 PPI. (Specifically, we used Citibase series PW1100 and PW1300 with weights 0.66 and 0.34, respectively.)

Table 12:
Sensitivity Of Results To Changes in Impact Phillips Trade-off (λ)

A. Phillips Trade-off

	λ	— Horizon — 1 Year		3 Year	
Monthly Data					
	-3.00	-0.64	(.17)	-0.77	(.41)
	-1.56	-0.63	(.16)	-0.74	(.13)
	-0.75	-0.61	(.15)	-0.70	(.12)
	-0.50	-0.58	(.14)	-0.66	(.11)
	-0.25	-0.52	(.11)	-0.56	(.08)
	-0.15	-0.45	(.09)	-0.46	(.06)
	-0.07	-0.32	(.07)	-0.30	(.05)
	-0.03	-0.17	(.07)	-0.15	(.05)
	0.00	0.09	(.11)	.06	(.07)
Quarterly Data					
	-0.17	-0.29	(.06)	-0.27	(.05)
	-0.07	-0.13	(.11)	-0.12	(.12)

B. Contribution of Demand Shock to Forecast Error in u and π
Variance Decomposition

	Horizon —Impact—				—1 Year—				—3 Year—			
λ	u		π		u		π		u		π	
Monthly Data												
-3.00	1.00	(.33)	0.00	(.00)	0.99	(1.6)	0.15	(3.8)	0.99	(1.5)	0.26	(10.)
-1.56	1.00	(.01)	0.00	(.00)	0.99	(.12)	0.15	(.30)	0.99	(.12)	0.27	(.82)
-0.75	1.00	(.00)	0.01	(.00)	0.98	(.03)	0.17	(.05)	0.98	(.03)	0.30	(.10)
-0.50	0.99	(.01)	0.02	(.00)	0.96	(.03)	0.20	(.05)	0.96	(.04)	0.34	(.09)
-0.25	0.95	(.02)	0.07	(.01)	0.90	(.05)	0.28	(.05)	0.89	(.06)	0.44	(.09)
-0.15	0.85	(.03)	0.19	(.02)	0.77	(.07)	0.40	(.05)	0.77	(.09)	0.57	(.08)
-0.07	0.52	(.04)	0.52	(.04)	0.43	(.09)	0.68	(.04)	0.42	(.10)	0.81	(.04)
-0.03	0.16	(.02)	0.87	(.03)	0.10	(.05)	0.88	(.03)	0.09	(.06)	0.89	(.04)
0.00	0.00	(.00)	1.00	(.00)	0.01	(.02)	0.85	(.04)	0.01	(.02)	0.73	(.08)
Quarterly Data												
-0.17	0.40	(.06)	0.60	(.06)	0.34	(.08)	0.83	(.04)	0.30	(.11)	0.92	(.03)
-0.17	0.10	(.02)	0.90	(.05)	0.07	(.04)	0.89	(.04)	0.05	(.04)	0.80	(.10)

Notes: The quarterly results are constructed from a VAR(4) using quarterly averages of the monthly CPI and unemployment rate data. Standard errors are shown in parentheses.

Table 13:
Sensitivity of Estimated Phillips Trade-off to Changes in Specification

Specification	Phillips Trade-off for Horizon: Impact	1 year	3 year
		$\lambda = -1.56$	
1. Baseline Model	-1.56	-0.63 (.16)	-0.74 (.13)
2. Levels	-1.56	-0.59 (.16)	0.22 (36.)
3. 18 Monthly Lags	-1.56	-0.64 (.16)	-0.64 (.12)
4. 4 Quarterly Lags	-1.56	-0.62 (.10)	-0.61 (.09)
5. 6 Quarterly Lags	-1.56	-0.57 (.09)	-0.55 (.10)
6. GDP Deflator	-1.56	-2.67 (1.6)	-1.91 (.82)
7. π^{fe} (Monthly)	-1.56	-0.69 (.19)	-0.82 (.18)
8. π^{fe} (Quarterly)	-1.56	-0.68 (.13)	-0.64 (.14)
		$\lambda = -0.07$	
1. Baseline Model	-0.07	-0.32 (.07)	-0.30 (.05)
2. Levels	-0.07	-0.22 (.07)	0.54 (.35)
3. 18 Monthly Lags	-0.07	-0.31 (.08)	-0.27 (.06)
4. 4 Quarterly lags	-0.07	-0.13 (.07)	-0.12 (.06)
5. 6 Quarterly Lags	-0.07	-0.09 (.08)	-0.04 (.10)
6. GDP Deflator	-0.07	-0.35 (.11)	-0.26 (.12)
7. π^{fe} (Monthly)	-0.07	-0.34 (.09)	-0.32 (.05)
8. π^{fe} (Quarterly)	-0.07	-0.11 (.11)	-0.10 (.11)
		$\lambda = 0.00$	
1. Baseline Model	0.00	0.09 (.11)	0.06 (.07)
2. Levels	0.00	0.18 (.11)	0.84 (.49)
3. 18 Monthly Lags	0.00	0.21 (.18)	0.18 (.14)
4. 4 Quarterly Lags	0.00	0.10 (.12)	0.08 (.09)
5. 6 Quarterly Lags	0.00	0.16 (.14)	0.22 (.19)
6. GDP Deflator	0.00	-0.09 (.10)	0.02 (.13)
7. π^{fe} (Monthly)	0.00	0.20 (.17)	0.11 (.09)
8. π^{fe} (Quarterly)	0.00	0.22 (.20)	0.22 (.21)

Description of Specifications:
1. The VAR(12) used Section 4
2. VAR(12) with levels used in place of first differences
3. VAR(18)
4. Quarterly VAR(4) using quarterly averages of the CPI and unemployment rate
5. Quarterly VAR(6) using quarterly averages of the CPI and unemployment rate
6. Quarterly VAR(6) using the GDP deflator instead of the CPI
7. VAR(12) with the relative price of food and energy included
8. Quarterly VAR(6) with the relative price of food and energy included

Table 14:
Sensitivity of Estimated Variance Decompositions to Changes in Specification Contribution of Demand Shock to Forecast Error Variance Decomposition

	Horizon					
	—Impact—		—1 year—		—3 year—	
Spec.	u	π	u	π	u	π
			$\lambda = -1.56$			
1	1.00 (.01)	0.00 (.00)	0.99 (.12)	0.15 (.30)	0.99 (.12)	.027 (.82)
2	1.00 (.01)	0.00 (.00)	0.99 (.80)	0.16 (.32)	0.73 (3.0)	0.24 (.72)
3	1.00 (.01)	0.00 (.00)	0.98 (.24)	0.17 (.37)	0.95 (.55)	0.39 (1.5)
4	0.98 (.03)	0.02 (.00)	0.96 (.04)	0.30 (.08)	0.95 (.06)	0.54 (.13)
5	0.99 (.02)	0.02 (.00)	0.96 (.04)	0.31 (.08)	0.89 (.10)	0.62 (.13)
6	1.00 (.01)	0.02 (.00)	1.00 (.00)	0.07 (.04)	0.99 (.02)	0.13 (.09)
7	1.00 (.01)	0.00 (.00)	0.95 (.53)	0.13 (.19)	0.91 (1.0)	0.20 (.40)
8	0.94 (.04)	0.02 (.00)	0.89 (.06)	0.22 (.07)	0.85 (.10)	0.42 (.14)
			$\lambda = -0.07$			
1	0.52 (.04)	0.52 (.04)	0.43 (.09)	0.68 (.04)	0.42 (.10)	0.81 (.04)
2	0.53 (.04)	0.52 (.04)	0.41 (.08)	0.67 (.04)	0.36 (.08)	0.75 (.05)
3	0.52 (.04)	0.52 (.04)	0.38 (.09)	0.67 (.04)	0.32 (.11)	0.80 (.04)
4	0.10 (.02)	0.90 (.05)	0.07 (.04)	0.89 (.04)	0.05 (.04)	0.80 (.10)
5	0.10 (.02)	0.91 (.05)	0.06 (.03)	0.88 (.05)	0.02 (.02)	0.72 (.13)
6	0.12 (.02)	0.92 (.05)	0.16 (.07)	0.94 (.03)	0.12 (.09)	0.97 (.02)
7	0.50 (.04)	0.52 (.04)	0.35 (.08)	0.57 (.07)	0.32 (.09)	0.55 (.14)
8	0.09 (.02)	0.86 (.05)	0.05 (.03)	0.70 (.08)	0.02 (.03)	0.45 (.13)
			$\lambda = 0.00$			
1	0.00 (.00)	1.00 (.00)	0.01 (.02)	0.85 (.04)	0.01 (.02)	0.73 (.08)
2	0.00 (.00)	1.00 (.00)	0.01 (.02)	0.84 (.05)	0.27 (.10)	0.76 (.08)
3	0.00 (.00)	1.00 (.00)	0.02 (.02)	0.83 (.05)	0.04 (.05)	0.61 (.10)
4	0.00 (.00)	1.00 (.00)	0.00 (.01)	0.78 (.08)	0.01 (.02)	0.57 (.13)
5	0.00 (.00)	1.00 (.00)	0.01 (.01)	0.77 (.09)	0.05 (.06)	0.47 (.14)
6	0.00 (.00)	1.00 (.01)	0.01 (.01)	0.95 (.05)	0.01 (.01)	0.91 (.09)
7	0.00 (.00)	0.99 (.01)	0.02 (.02)	0.76 (.07)	0.02 (.03)	0.52 (.14)
8	0.00 (.00)	0.99 (.02)	0.01 (.01)	0.66 (.09)	0.03 (.04)	0.31 (.11)

Description of Specifications:
1. The VAR(12) used Section 4
2. VAR(12) with levels used in place of first differences
3. VAR(18)
4. Quarterly VAR(4) using quarterly averages of thc CPI and unemployment rate
5. Quarterly VAR(6) using quarterly averages of the CPI and unemployment rate
6. Quarterly VAR(6) using the GDP deflator instead of the CPI
7. VAR(12) with the relative price of food and energy included
8. Quarterly VAR(6) with the relative price of food and energy included

the joint behavior of inflation and unemployment. In Section 2 of the paper, we explore how the links between inflation and unemployment depend on whether we look at low frequencies (trend behavior), intermediate frequencies (business-cycle behavior), or high frequencies (irregular behavior). In Sections 4, 5, and 6 of the paper, we use linear time series methods to study the reduced-form interaction of inflation and unemployment as well as the structural Phillips curve.

We begin by documenting that there is a pronounced negative correlation of inflation and unemployment at business-cycle frequencies, which is remarkably stable over the postwar period. Lower frequency comovements of inflation and unemployment, however, display links that are very unstable across time. When we turn to a more detailed time series characterization of the bivariate process, there are three notable results. First, there is evidence of I(1) behavior in inflation and unemployment, but no evidence of cointegration. This corresponds to the idea that there are "stochastic trends" in inflation and unemployment; it sets the stage for structural estimates relating these trends. Second, there is close to a one-way Granger-causal structure. That is, in the dynamic reduced form (forecasting VAR), unemployment depends mainly on its own one-step-ahead forecast errors while inflation depends on errors in both inflation and unemployment. Third, there is important evidence of econometric instability over subsamples. We begin by using a battery of tests to document general evidence of changing structure. We then provide a characterization of how the structure changes through time, focusing on differences between the pre-1970 and later sample periods. In terms of short-term forecasting, we find that there is little difference between subperiods: the standard deviation of one-step-ahead forecast error is essentially the same over subperiods, and there are only small differences in forecast performance out to a horizon of twelve months. By contrast, there are major changes that affect the ability to forecast inflation and unemployment in the longer term. In particular, the standard deviation of long-horizon forecast error (two years and beyond) is much larger in the later period, most strikingly for inflation. This is an indication of increased persistence of the effects of shocks. It is this increased persistence, not more volatile shocks, that makes the post-1970 interval intrinsically more difficult to forecast. Moreover, medium to long-term forecasts are affected by parameter instability, but this instability — which would suggest the value of reestimating earlier Phillips curve models over the 1970 to 1990 period — is swamped by the general increase in the difficulty of forecasting in the post-1970 period.

In terms of results from structural models of inflation and unemployment, we find some results that are surprising. We work in the style of researchers like Gordon and Solow, who interpreted the Phillips curve structurally, but we do this using structural VAR techniques. These procedures require that we specify how to map from the forecast errors of the VAR into economically interpretable shocks, *i.e.*, that we undertake an identification of the Phillips curve system. Alternative identifying assumptions have important implications for the magnitude of long-run multipliers, for the sources of business-cycle fluctuations, for the interpretation of episodes, and for sacrifice ratios, defined as the unemployment cost of moving to a permanently lower rate of inflation. We compute long-run trade-offs between inflation and unemployment, despite the arguments of Lucas [1972] and Sargent [1971]. As discussed in Fisher and Seater [1993] and King and Watson [1992], when inflation is an I(1) process this is a valid exercise; we provide evidence that the data are consistent with the I(1) restriction for all of the sample periods that we study.

A traditional Keynesian identification yields: (i) large estimated long-run trade-offs between inflation and unemployment, although these fall by 50% in the latter sample period; (ii) 2-year-ahead forecast errors (one measure of business-cycle fluctuations) in which demand shocks explain essentially all of unemployment and only 25% of inflation; (iii) long-run variability in inflation with a source that is approximately 50% demand shocks and 50% other (price, supply) shocks. Further, under this identification, every major postwar recession is fully explained by demand shocks (even the oil-price interval) and most recession intervals involve a decline in inflation. Finally, sacrifice ratios — specified as the cumulative loss in unemployment over a five-year period — of a permanent disinflation (induced by demand changes) are large (3.7), but not as large as those in the DRI model (8.0).

By contrast, a real business-cycle identification yields very small estimated trade-offs between inflation and unemployment at all horizons and 2-year-ahead forecast errors in unemployment that are dominated by identified disturbances. In addition, all postwar recessions are explained by supply shocks (even the 1981-82 recession). Thus, both the traditional Keynesian and RBC models provide coherent — although very different — explanations of the postwar data.

An alternative identifying assumption — suggested by the work of Sargent [1976] and Barro and Rush [1980] — yields a very different picture of economic fluctuations in the short and long term: (i) there are much smaller estimated long-run trade-offs between inflation and unemployment than in the traditional Keynesian model; (ii) demand shocks explain 43% of the two-year forecast error in unemployment and 75% of that in inflation; (iii) the long-run variability of inflation is nearly entirely due to demand shocks; and (iv) supply shocks have little long-run effect on inflation but have important effects on real activity in both the short and long run. Further, major postwar recessions appear to be a result of a mix of supply and demand shocks (even the oil-price interval); most recessions continue to involve declines in inflation. Finally, sacrifice ratios at the five-year horizon are much smaller (1.52) than in the traditional Keynesian model (3.71). Arguably, this identification yields a mainstream interpretation of the postwar U.S. data.

Our results reinforce some beliefs of the neoclassical/monetarist and Keynesian schools, but they provide many more challenges. We highlight three of these results. First, there is evidence for the Lucas-Sargent hypothesis: increased persistence of inflation reduces the long-run Phillips curve slope. However, the inflation-unemployment trade-off slope is also affected by short-run identifications in a way that is at least as important quantitatively. Second, the time series evidence indicates why Keynesian macroeconometricians have seen little reason to change their practices, except for potentially "patching up" the long-run slope of the Phillips curve. That is, they have seen relatively little evidence of instability over horizons of interest (forecasts of up to a year). While the quality of longer-run forecasts of both inflation and unemployment have deteriorated in the post-1970 period, the deterioration reflects an increase in the intrinsic uncertainty surrounding longer-run forecasts of these series. Little improvement is achieved by updating the forecasting equations estimated through 1969. Third, the Phillips curve at every horizon is more unstable across identifications than it is across time: at shorter horizons, the identification is essentially all that matters.

Thus, as a result of our investigation, it is hard for a neoclassical monetarist economist to argue that the Phillips correlations are absent from the U.S. data or that a structural Phillips curve is unstable in the short run. While traditional Keynesian macroeconometricians may take comfort in the stability of the short-term Phillips curve, we think that few other macroeconomists will find that their short-run identifications generate a plausible description of postwar U.S. business cycles.

A. Derivation of theoretical results

In this appendix we provide background for results stated in Section 3 of the main test. In that section, we considered the structural model:

$$u_t = f\pi_t - \sum_{i=1}^{q} g_i E_{t-i}\pi_t + s_t, \tag{22}$$

$$\pi_t = \rho_1\pi_{t-1} + ... + \rho_n\pi_{t-n} + m_t, \tag{23}$$

where s_t and m_t are *iid* zero-mean random variables. The model has a rational expectations reduced form

$$u_t = \sum_{i=0}^{q+n} \beta_{u\pi,i}\pi_{t-i} + s_t = \beta_{u\pi}(L)\pi_t + s_t, \tag{24}$$

and our first goal is to show that $\beta_{u\pi}(1) = \sum_{i=0}^{q+n} \beta_{u\pi,i} = f - \sum_{i=1}^{q} g_i$ if there is a single unit root in the inflation generating process.

To begin, write (23) as $\rho(L)\pi_t = m_t$, and assume that $\rho(z)$ has one unit root and all other roots outside the unit circle. Define $\phi(z)$ by $\rho(z) = (1-z)\phi(z)$, so that $\Delta\pi_t = \pi_t - \pi_{t-1}$ is stationary with moving average representation $\Delta\pi_t = \phi(L)^{-1}m_t = \mu(L)m_t$. Then,

$$\begin{aligned} E_{t-j}\Delta\pi_t &= \mu_j m_{t-j} + \mu_{j+1}m_{t-j+1} + ... \\ &= [L^{-j}\mu(L)]_+ m_{t-j} \\ &= [L^{-j}\mu(L)]_+ \phi(L)\Delta\pi_{t-j} \end{aligned} \tag{25}$$

where $[\cdot]_+$ means "ignore negative powers of L" as in Hansen and Sargent [1980]. Further,

$$\begin{aligned} E_{t-j}\pi_t &= E_{t-j}\{\pi_{t-j} + \Delta\pi_{t-j+1} + ...\Delta\pi_t\} \\ &= \{1 + \sum_{i=1}^{j}[L^{-i}\mu(L)]_+(1-L)\phi(L)\}\pi_{t-j}. \\ &= \kappa_j(L)\pi_{t-j}, \end{aligned} \tag{26}$$

where $\kappa_j(1) = 1$ for all j. Thus, in (24):

$$\begin{aligned} \beta_{u\pi}(L)\pi_t &= f\pi_t - \sum_{i=1}^{q} g_i E_{t-i}\pi_t \\ &= f\pi_t - \sum_{i=1}^{q} g_i\kappa_i(L)\pi_{t-i}. \end{aligned} \tag{27}$$

So that $\beta_{u\pi}(z) = \{f - \sum_{i=1}^{q} g_i\kappa_i(z)z^i\}$ and $\beta_{u\pi}(1) = \{f - \sum_{i=1}^{q} g_i\}$ as required.

In Section 3, we also wrote (24) as

$$u_t = \{f - \sum_{i=1}^{q} g_i\}\mu(1)M_t + \psi(L)m_t + s_t \tag{28}$$

where $M_t = M_{t-1} + m_t$ is a martingale, $\psi(L)m_t$ is a stationary component of unemployment arising from demand shocks, and where $\Delta\pi_t = \mu(L)m_t$ is the moving average representation for $\Delta\pi_t$. We will derive (28) under the assumption that $\mu(L)$ is 1-summable; thus (28) is valid when π_t is I(1), so that $\mu(L) = \phi(L)^{-1}$ as above, or when π_t is I(0), so that $\mu(L) = (1 - L)\rho(L)^{-1}$.

To derive (28), let $\bar{\pi}_t$ denote the permanent component of inflation as in Beveridge and Nelson [1981]: $\bar{\pi}_t = lim_{k\to\infty}E_t\pi_{t+k}$. Since $\pi_{t+j} = \pi_{t-1} + \{\Delta\pi_t + \Delta\pi_{t+1} + ...\Delta\pi_{t+j}\}$,

$$\begin{aligned}
\bar{\pi}_t &= \pi_{t-1} + lim_{j\to\infty} \sum_{k=0}^{j} E_t\Delta\pi_{t+k} \\
&= \pi_{t-1} + lim_{j\to\infty} \sum_{k=0}^{j}[L^{-k}\mu(L)]_+ m_t \\
&= \bar{\pi}_{t-1} + \mu(1)m_t \\
&= \mu(1)M_t,
\end{aligned} \tag{29}$$

from the 1-summability of $\mu(L)$ and the definitions of $\bar{\pi}_t$ and M_t. Notice that if inflation is I(0), then $\mu(1) = 0$ and there is no variation in trend inflation. By contrast, if inflation contains a unit root, then there is variation in $\bar{\pi}_t$ since $\mu(1) = 1/\phi(1)$ is nonzero.

Now, write (22) as:

$$u_t = \{f - \sum_{i=1}^{q} g_i\}\pi_t + \{\sum_{i=1}^{q} g_i(\pi_t - E_{t-i}\pi_t)\} + s_t. \tag{30}$$

Decompose $\pi_t = \mu(1)M_t + \mu^*(L)m_t$, where $\mu^*(z) = [\mu(z) - \mu(1)]/1 - z$. When $\mu(L)$ is 1-summable, $\mu^*(1) = 0$, so that $\mu^*(L)m_t$ can be interpreted as a "temporary component," Combining these expressions, we can write:

$$u_t = \{f - \sum_{i=1}^{q} g_i\}\bar{\pi}_t + \psi(L)m_t + s_t. \tag{31}$$

where $\psi(L)m_t = \{\sum_{i=1}^{q} g_i(\pi_t - E_{t-i}\pi_t)\} + \{f - \sum_{i=1}^{q} g_i\}\mu^*(L)m_t$. Equation (28) follows from (29) and (31).

B. REM estimates of the short-run Phillips curve

As our value of the short-run Phillips curve slope from the rational expectations monetarist (REM) literature, we use $\lambda = -.07$. In this appendix, we show how estimates in this range may be derived from the REM studies of Sargent [1976] and Barro and Rush [1980].

B.1 Sargent's short-run Phillips curve

The short-run Phillips curve estimated by Sargent [1976, Table 9] takes the form:

$$u_t = -0.287(P_t - E_{t-1}P_t) + \text{predetermined variables} + \epsilon_t$$

That is, unemployment innovations are attributable to a price-level forecast error and a shock term. (In this expression, P_t denotes the logarithm of the price level at date t, represented empirically by the GNP deflator.) The predetermined part of unemployment, $E_{t-1}u_t$, is explained by its own lags as well as by a constant plus deterministic trend. Since his model suggested that u_t and $P_t - E_{t-1}P_t$ would be jointly determined, Sargent estimated this specification by instrumental variables. He estimated his model on quarterly U.S. data from 1947 to 1978.

To relate Sargent's specification to ours, we simply note that the annualized inflation rate, $\pi_t = 4(P_t - P_{t-1})$. Further, forecasting errors for the price-level and inflation rate coincide (up to a scaling factor) if P_{t-1} is an element of the information set on which $E_{t-1}P_t$ is based, as it is in Sargent's analysis and ours. Thus, $\pi_t - E_{t-1}\pi_t = 4(P_t - E_{t-1}P_t)$. Hence, Sargent's estimator of λ is $-.287/4 = -.07$.

B.2. The Barro and Rush slope estimate

One can derive an implicit instrumental variables estimator of λ from the work of Barro and Rush [1980] on the effects of unanticipated money growth on unemployment and the price level.

These authors' estimates may be interpreted as pertaining to the unemployment rate and the price level as in:

$$u_t = \lambda_{uM}(M_t - E_{t-1}M_t) + \text{predetermined variables} + \epsilon_{ut}$$

$$P_t = \lambda_{PM}^1 M_t + \lambda_{PM}^2(M_t - E_{t-1}M_t) + \text{predetermined variables} + \epsilon_{Pt}$$

The price level is positively affected by the level of the money stock ($\lambda_{PM}^1 \cong 1$) and negatively by unanticipated money when that raises output.

The overall price-level effect of a money shock is $\lambda_{PM} = (\lambda^1_{PM} + \lambda^2_{PM}) > 0$. Thus, we can form an implicit instrumental variables estimator of our λ parameter as:

$$\lambda = \frac{\lambda_{uM}}{4\lambda_{PM}}.$$

That is, the Barro-Rush estimator determines the short-run slope as the effect of monetary-induced price changes on unemployment. As in the Sargent analysis, we must scale price-level surprises by 4 to convert them into surprises in our annual inflation rate.

A complication arises from the fact that the Barro-Rush unemployment equation is actually estimated using a dependent variable of the form $x_t = log(u_t/(1-u_t))$. Hence, we must scale estimates of λ_{xM} by $\frac{du}{dx} = u(1-u) \simeq .05$.

There are a battery of estimates in Barro and Rush [1980], which are differentiated by various assumptions about whether the constraint that $\lambda^1_{PM} = 1$ is imposed, serial correlation correction, lag length, etc. The largest magnitude of the short-run slope comes about when $\lambda_{xm} = -4.0$ and $\lambda_{PM} = \lambda^1_{PM} + \lambda^2_{PM} = .30$: these are the point estimates in the final columns of Tables 2.1 and 2.2

$$\lambda = \frac{du}{dx}\frac{\lambda_{xM}}{4\lambda_{PM}} = .05\frac{-4.0}{4(.30)} = -.17.$$

This λ estimate involves the smallest estimated short-run effect of money on prices. Using another estimate of $\lambda_{PM} = \lambda^1_{PM} + \lambda^2_{PM} = .63$ provides a value of λ that is less than one-half as large in magnitude and, hence, is very close to that of Sargent. Perhaps this coincidence is not surprising since one of Sargent's instrumental variables was the unexpected component of money.

References

Anderson, L.C. and Carlson, K.M., (1972). An Econometric Analysis of the Relation of Monetary Variables to the Behavior of Prices and Unemployment, *The Econometrics of Price Determination*, (ed.) O. Eckstein. Washington, D.C.: Board of Governors of the Federal Reserve System.

Andrews, D.W.K., Lee, I., and Ploberger, W., (1992). Optimal Changepoint Tests for Normal Linear Regression, Cowles Foundation working paper no. 1016, Yale University.

Ball, L., (1993). What Determines the Sacrifice Ratio?, manuscript, Princeton University.

Barro, R.J. and Rush, M., (1980). Unanticipated Money and Economic Activity, *Rational Expectations and Economic Policy*, (ed.) S. Fischer, Chicago, IL: University of Chicago Press.

Baxter, M., and King, R.G., (1993). Measuring Business Cycles, Approximate Band-pass Filters for Economic Time Series, manuscript, University of Virginia.

Beveridge, S. and Nelson, C.R., (1981). A New Approach to the Decomposition of Time Series in Permanent and Transitory Components with Particular Attention to the Measurement of the Business Cycle, *Journal of Monetary Economics*, **7**: 151–74.

Brown, R.L., Durbin, J., and Evans, J.M., (1975). Techniques for Testing the Constancy of Regression Relationships Over Time, *Journal of the Royal Statistical Society*, Series *B*, **37**: 149–63.

Eckstein, O., (1972). *The Econometrics of Price Determination*, Washington, D.C.: Board of Governors of the Federal Reserve System

Eckstein, O. and Fromm, G., (1968). The Price Equation, *American Economic Review*, **58**: 1159–1183.

Fischer, S., (1977). Long-term Contracts, Rational Expectations, and the Optimal Money Supply Rule, *Journal of Political Economy*, **85**: 191–205.

Fisher, M., and Seater, J., (1993). Long-Run Neutrality and Superneutrality in an ARIMA Framework, *American Economic Review*, **83**: 402–415.

Friedman, M., (1968). The Role of Monetary Policy, *American Ecnomic Review*, **58**: 1–17.

Gordon, R.J., (1970). The Recent Acceleration of Inflation and Its Lessons for the Future, *Brookings Papers on Economic Activity*, **1**: 8–41.

Gordon, R.J., (1971). Prices in 1970: The Horizontal Phillips Curve, *Brookings Papers on Economic Activity*, **3**: 449–58.

Gordon, R.J., (1982). Price Inertia and Ineffectiveness in the United States, 1890–1980, *Journal of Political Economy*, **90**: 1087–1117.

Gordon, R.J., (1990a). Comments: The Phillips Curve Now and Then, *Growth/Productivity/Unemployment*, (ed.) P. Diamond Cambridge, MA: MIT Press.

Gordon, R.J., (1990b). What is New-Keynesian Economics? *Journal of Economic Literature*, **28**: 1116–1171.

Gordon, R.J. and King, S., (1982). The Output Cost of Disinflation in Traditional and Vector Autoregressive Models, *Brookings Papers on Economic Activity*, **1**: 205–44.

Granger, C., (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods, *Econometrica*, **34**: 150–161.

Gray, J., (1976). Wage Indexation: A Macroeconomic Approach, *Journal of Monetary Economics*, **2**: 221–235.

Hansen, B.E., (1991). Tests for Parameter Instability in Regressions with I(1) Processes, *Journal of Business and Economic Statistics*, **10**: 321–36.

Hansen, L.P. and Sargent, T.J., (1980). Formulating and Estimating Dynamic Linear Rational Expectations Models, *Journal of Economic Dynamics and Control*, **2**: 7–46.

Hicks, J., (1937). Mr. Keynes and the Classics, *Econometrica*, **5**: 147–159.

King, R.G., and Watson, M.W., (1992). Testing Long Run Neutrality, National Bureau of Economic Research working paper no. 4156.

Klein, L., (1950). *Economic Fluctuations in the United States*, New York: John Wiley and Sons.

Lucas, R.E., Jr., (1972). Econometric Testing of the Natural Rate Hypothesis, *The Econometrics of Price Determination*, (ed.) O. Eckstein. Washington, D.C.: Board of Governors of the Federal Reserve System.

Lucas, R.E., Jr. and Sargent, T.J., (1979). After Keynesian Macroeconometrics, *After the Phillips Curve: Persistence of High Inflation and Unemployment*, Federal Reserve Bank of Boston, Conference Series No. 19.

Mankiw, N.G., (1990). A Quick Refresher Course in Macroeconomics, *Journal of Economic Literature*, **28**: 1645-60.

Modigliani, F., (1944). Liquidity Preference and the Theory of Interest and Money, *Econometrica*, **12**: 45–88.

Nyblom, J., (1989). Testing for the Constancy of Parameters Over Time, *Journal of the American Statistical Association*, **84**: 223–30.

Okun, A.M., (1978). Efficient Disinflationar Policies, *American Economic Review*, **68**: 348–57.

Ploberger, W. and Kramer, W., (1992). The CUSUM Test with OLS Residuals, *Econometrica*, **60**: 271–86.

Phelps, E., (1967). Phillips Curves, Expectations of Inflation, and Optimal Inflation Over Time, *Economica*, NS **135**: 254-281.

Phelps, E. and Taylor, J.B., (1977). Stabilization Powers of Monetary Policy Under Rational Expectations, *Journal of Political Economy*, **85**: 163–90.

Phillips, A.W.H., (1958). The Relation Between Unemployment and the Rate of Change of Money Wages in the United Kingdom, 1861–1957, *Economica*, NS **25**: 283–299.

Roberts, J., (1993). The Sources of Business Cycles: A Monetarist Interpretation, *International Economic Review*, **34**: 923–934.

Quandt, R.E., (1960). Tests of the Hypothesis that a Linear Regression System Obeys Two Separate Regimes, *Journal of the American Statistical Association*, **55**: 324–330.

Samuelson, P.A. and Solow, R.M., (1960). Analytical Aspects of Anti-Inflation Policy, *American Economic Review, Papers and Proceedings*, **50**: 177–194.

Sargent, T.J., (1971). A Note on the Accelerationist Controversy, *Journal of Money, Credit and Banking*, **3**: 50–60.

Sargent, T.J., (1976). A Classical Macroeconometric Model for the United States, *Journal of Political Economy*, **84**: 207–37.

Sargent, T.J., (1979). *Macroeconomic Theory*, Orlando, FL: Academic Press.

Solow, R.M., (1969). *Price Expectations and the Phillips Curve*, Manchester, UK: Manchester University Press.

Solow, R.M., (1970). Discussion of R.J. Gordon's The Recent Acceleration of Inflation and Its Lessons for the Future, *Brookings Papers on Economic Activity*, **1**: 42–6.

Stock, J.H., (1991). Confidence Intervals of the Largest Autoregressive Root in U.S. Macroeconomic Time Series, *Journal of Monetary Economics*, **28**: 435–460.