# Inflation and monetary policy in the twentieth century

**Lawrence J. Christiano and Terry J. Fitzgerald**

## Introduction and summary

Economists continue to debate the causes of inflation. One reason for this is that bad economic outcomes are frequently accompanied by anomalous inflation behavior. The worst economic performance in the U.S. in the twentieth century occurred during the Great Depression of the 1930s, and there was a pronounced deflation at that time. Economic performance in the U.S. in the 1970s was also weak, and that was associated with a pronounced inflation.

So, what is it that makes inflation sometimes high and sometimes low? In one sense, there is widespread agreement. Most economists think that inflation cannot be unusually high or low for long, without the fuel of high or low money growth.[1] But, this just shifts the question back one level. What accounts for the anomalous behavior of money growth?

Academic economists attempting to understand the dynamics of inflation pursue a particular strategy. They start by studying the dynamic characteristics of inflation data, as well as of related variables. These characteristics represent a key input into building and refining a model of the macroeconomy. The economist's model must not only do a good job in capturing the behavior of the private economy, but it must also explain the behavior of monetary authorities. The hoped-for final product of this research is a model that fits the facts well. Implicit in such a model is an "explanation" of the behavior of inflation, as well as a prescription for what is to be done to produce better outcomes.[2]

To date, much research has focused on data from the period since World War II. For example, considerable attention and controversy have been focused on the apparent inflation "inertia" in these data: the fact that inflation seems to respond only with an extensive delay to exogenous shifts in monetary policy.[3] We argue that much can be learned by incorporating data from the first half of the century into the analysis. The data from the early part of the century behave quite differently in many ways from the data we are accustomed to studying. In particular, we emphasize four differences between the pre- and post-war data:[4]

- Inflation is much more volatile, and less persistent, in the first half of the twentieth century.
- Average inflation is lower in the first half of the century.
- Money growth and inflation are coincident in the first half of the century, while inflation lags money by about two years in the second half.
- Finally, inflation and unemployment are strongly negatively related in the first half of the century, while in the second half a positive relationship emerges, at least in the lower frequency components of the data.

These shifts in the behavior of inflation constitute potentially valuable input in the quest for a good model.

The outline of our article is as follows. To set the background, we begin with a brief, very selective, overview of existing theories about inflation. We divide the set of theories into two groups: those that focus on "people" and those that focus on "institutions." We describe the very different implications that each group of theories has for policy. We then turn to documenting the facts listed above. After that, we review the implications of the facts for theories. We focus in particular on the institution view. According to this view, what

*Lawrence J. Christiano is a professor of economics at Northwestern University, a research fellow at the National Bureau of Economic Research (NBER), and a consultant to the Federal Reserve Bank of Chicago. He acknowledges support from a National Science Foundation grant to the NBER. Terry J. Fitzgerald is a professor of economics at St. Olaf College and a consultant to the Federal Reserve Bank of Minneapolis.*

is crucial to achieving good inflation outcomes is the proper design of monetary policy institutions. Our discussion reviews ideas initially advanced by Kydland and Prescott (1977) and later developed further by Barro and Gordon (1983a, b), who constructed a beautifully simple model for expositing the ideas. We show that the Barro–Gordon model does very well at understanding the second and fourth facts above concerning inflation in the twentieth century.[5] We also discuss the well-known fact that that model has some difficulty in addressing the disinflation that occurred in the U.S. in the 1980s. This and other considerations motivate us to turn to modern representations of the ideas of Kydland–Prescott and Barro–Gordon. While this work is at an early stage, it does contain some surprises and may lead to improved theories that provide a better explanation of the inflation facts.

## Ideas about inflation: People versus institutions

Economists are currently pursuing several theories for understanding inflation behavior. However, the theories are still in their infancy and are best thought of as "prototypes": They are too simple to be credibly thought of as fitting the facts well. Although these research programs are still at an early stage, it is possible to see two visions emerging. Each has different implications for what needs to be done to achieve better inflation outcomes. To understand what is at stake in this research, it is interesting to sketch the different visions. Our loose names for the competing visions are the *people* vision on the one hand and the *institution* vision on the other. Although it is not the case that all research neatly falls into one or the other of these categories, they are nevertheless useful for spelling out the issues.

Under the people vision, bad inflation outcomes of the past reflect the honest mistakes of well-meaning central bankers trying to do what is inherently a very difficult job. For example, Orphanides (1999) has argued that the high inflation of the 1970s reflects that policymakers viewed the low output of the time as a cyclical phenomenon, something monetary policy could and should correct. However, in retrospect we now know that the poor economic performance of the time reflected a basic productivity slowdown that was beyond the power of the central bank to control. According to Orphanides, real-time policymakers under a mistaken impression about the sources of the slowdown did their best to heat up the economy with high money growth. To their chagrin, they got only high inflation and no particular improvement to the economy. From this perspective, the high inflation of the 1970s was a blunder.

Another explanation of the high inflation of the 1970s that falls into what we call the people category appears in Clarida, Gali, and Gertler (1998). They characterize monetary policy using a framework advocated by Taylor (1993): Fed policy implements a "Taylor rule" under which it raises interest rates when expected inflation is high, and lowers them when expected inflation is low. According to Clarida, Gali, and Gertler, the Fed's mistake in the 1970s was to implement a version of the Taylor rule in which interest rates were moved too little in response to movements in expected inflation. They argue that this type of mistake can account for the inflation take-off that occurred in the U.S. in the 1970s.[6] In effect the root of the problem in the 1970s lay in a bad Taylor rule. According to the institution view, limitations on central bankers' technical knowledge about the mechanics of avoiding high inflation are not the key reason for the bad inflation outcomes that have occurred in the past. This view implicitly assumes that achieving a given inflation target over the medium run is not a problem from a technical standpoint. The problem, according to this view, has to do with central bankers' incentives to keep inflation on track and the role of government institutions in shaping those incentives.

The institution view—initiated by Kydland and Prescott (1977) and further developed by Barro and Gordon (1983a, b)—focuses on a particular vulnerability of central banks in democratic societies (see figure 1). If people expect inflation to be high (A), they may take protective actions (B), which have the effect of placing the central bank in a dilemma. On the one hand, it can accommodate the inflationary expectations with high money growth (C). This has the cost of producing inflation, but the advantage of avoiding a recession. On the other hand, the central bank can keep money growth low and prevent the inflation that people expect from occurring (D). This has the cost of producing a recession, but the benefit that inflation does not increase. Central bankers in a democratic society will be tempted to accommodate (that is, choose C) when confronted with this dilemma. If people think this is the sort of central bank they have, this increases the likelihood that A will occur in the first place.

So, what is at stake in these two visions, the people vision versus the institution vision? Each has different implications for what should or should not be done to prevent bad inflation outcomes in the future. The people vision implies that more and better research is needed to reduce the likelihood of repeating past mistakes. This research focuses more on the technical, operational aspect of monetary policy. For example, research motivated by the Clarida, Gali, and Gertler argument

FIGURE 1

**Central banker in a democratic society**

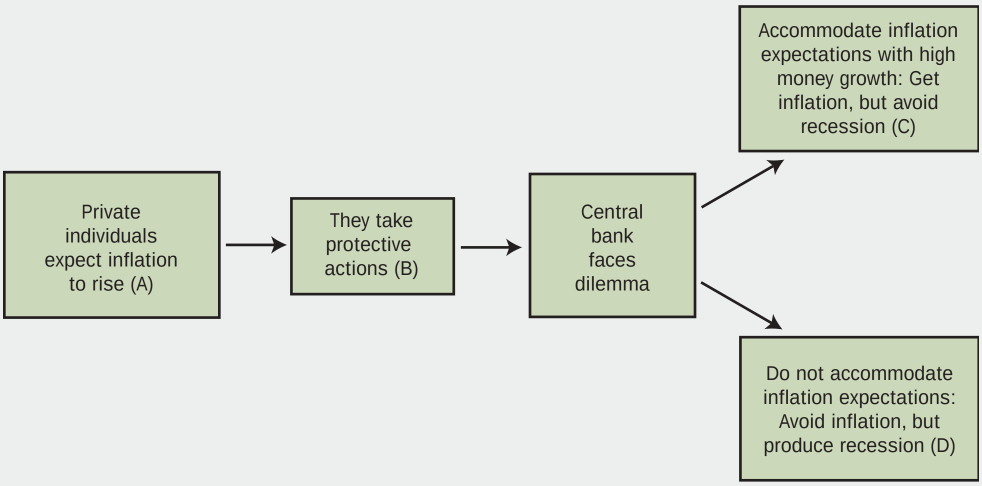


focuses on improvements in the design of the Taylor rule to ensure that it does not become part of the problem. The institutional perspective, not surprisingly, asks how better to design the institutions of monetary policy to achieve better outcomes. This type of work contemplates the consequences of, say, a legal change that makes low inflation the sole responsibility of the Federal Reserve. Other possibilities are the type of employment contracts tried in New Zealand, which penalize the central bank governor for poor inflation outcomes. The basic idea of this literature is to prevent scenarios like A in figure 1 from occurring, by convincing private individuals that the central bank would not choose C in the event that A did occur.

In this article, we start by presenting data on inflation and unemployment and documenting how those data changed before and after the 1960s. We argue that these data are tough for standard versions of theories that there is a time consistency problem in monetary policy. We then discuss whether there may be other versions of these theories that do a better job at explaining the facts.

## The data

This section describes the basic data on inflation and related variables and documents the observations listed in the introduction. First, we study the relationship between unemployment and inflation; then we turn to money growth and inflation.

### *Unemployment and inflation*

To show the difference between data in the first and second parts of the twentieth century, we divide the dataset into the periods before and after 1960. To better characterize the movements in the data, we break the data down into different frequency components. The techniques for doing this, reviewed in Christiano and Fitzgerald (1998), build on the observation that any data series of length, say $T$, can be represented exactly as the sum of $T/2$ artificial data series exhibiting different frequencies of oscillation. Each data series has two parameters: One controls the amplitude of fluctuation and the other, phase. The parameters are chosen so that the sum over all the artificial data series precisely reproduces the original data. Adding over just the data series whose frequencies lie inside the business cycle range of frequencies yields the *business cycle component* of the original data. We define the business cycle frequencies as those that correspond to fluctuations with period between two and eight years. We also consider a lower frequency component of the data, corresponding to fluctuations with period between eight and 20 years. We consider a very low frequency component of the data, which corresponds to fluctuations with period of oscillation between 20 and 40 years. Finally, for the post-1960 data when quarterly and monthly observations are available, we also consider the high frequency component of the data, which is composed of fluctuations with period less than two years.[7]

We begin by analyzing the data from the first part of the century. The raw data are displayed in figure 2, panel A. That figure indicates that there is a negative relationship between inflation and unemployment. This is confirmed by examining the scatter plot of inflation and unemployment in figure 2, panel B, which also shows a negative relationship (that is, a Phillips curve).[8] The regression line displayed in figure 2, panel B highlights this negative relationship.[9] Figure 2, panels C, D, and E exhibit the different frequency components of the data. Note that a negative relationship is apparent at all frequency components. The contemporaneous correlations between different frequency components of the inflation and unemployment data are reported in table 1. In each case, the number in parentheses is a *p*-value for measuring whether the indicated correlation is statistically different from zero. For example, a *p*-value less than 0.05 indicates that the indicated correlation is statistically different from zero at the 5 percent level.[10] The negative correlation in the business cycle frequencies is particularly significant.

We analyze the post-1960 monthly inflation and unemployment data in figure 3, panels A–F.[11] There is a sense in which these data look similar to what we saw for the early period, but there is another sense in which their behavior is quite different. To see the similarity, note from the raw data in figure 3, panel A that for frequencies in the neighborhood of the business cycle, inflation and unemployment covary negatively. That is, the Phillips curve seems to be a pronounced feature of the higher frequency component of the data. At the same time, the Phillips curve appears to have vanished in the very lowest frequencies. The data in figure 3, panel A show a slow trend rise in unemployment throughout the 1960s and 1970s, which is reversed starting in early 1983. A similar pattern occurs in inflation, though the turnaround in inflation begins in April 1980, roughly three years before the turnaround in unemployment. The low frequency component of the data dominates in the scatter plot of inflation versus unemployment, exhibited in figure 3, panel B. That figure suggests that the relationship between inflation and unemployment is positive, in contrast with the pre-1960s data, which suggest otherwise (see figure 2, panel B).[12]

We can formalize and quantify our impressions based on casual inspection of the raw data using frequency components of the data, as reported in figure 3, panels C–F. Thus, the frequency ranges corresponding to periods of oscillation between two months and 20 years (see figure 3, panels C–E) are characterized by a noticeable Phillips curve. Table 1 shows that the correlation in the range of high frequencies (when available) and in the business cycle frequencies is significantly negative. The correlation between inflation and unemployment is also negative in the 8–20 year range, but it is not statistically significantly different from zero in this case. Presumably, this reflects the relative paucity of information about these frequencies in the post-1960s data. Finally, figure 3, panel F indicates that the correlation between 20 and 40 year components is now positive, with unemployment lagging inflation. These results are consistent with the hypothesis that the Phillips curve changed relatively little in the 2–20 year frequency range, and that the changes that did occur are primarily concentrated in the very low frequencies. Formal tests of this hypothesis, shown in table B1 in box 1, fail to reject it.

Some of the observations reported above have been reported previously. For example, the low-frequency observations on unemployment have been documented using other methods in Barro (1987, Chapter 16). Also, similar frequency extraction methods have been used to detect the presence of the Phillips curve in the business cycle frequency range.[13] What has not been documented is how far the Phillips curve extends into the lowest frequencies. In addition, we show that inflation leads unemployment in the lowest frequency range.

Finally, we noted in the introduction that inflation in the early part of the century was more volatile and less persistent than in the second part. We can see this by comparing figure 2, panel A with figure 3, panel A. We can see the observation on volatility by comparing the scales on the inflation portion of the graphs.

TABLE 1

**CPI inflation and unemployment correlations**

| Sample | High frequency | Business cycle frequency | 8–20 years | 20–40 years |
|---|---|---|---|---|
| 1900–60 (annual) | | –0.57 (0.00) | –0.32 (0.19) | –0.51 (0.23) |
| 1961–97 (annual) | | –0.38 (0.11) | –0.16 (0.41) | 0.45 (0.32) |
| 1961:Q2–97:Q4 (quarterly) | –0.37 (0.00) | –0.65 (0.00) | –0.30 (0.29) | 0.25 (0.34) |
| 1961, Jan.–97, Dec. (monthly) | –0.24 (0.00) | –0.69 (0.00) | –0.27 (0.30) | 0.23 (0.40) |

Notes: Contemporaneous correlation over indicated sample periods and frequencies. Numbers in parentheses are *p*-values, in decimals, against the null hypothesis of zero correlation at all frequencies. For further details, see the text and notes 7 and 10.

FIGURE 2

**Unemployment and inflation, 1900–60**

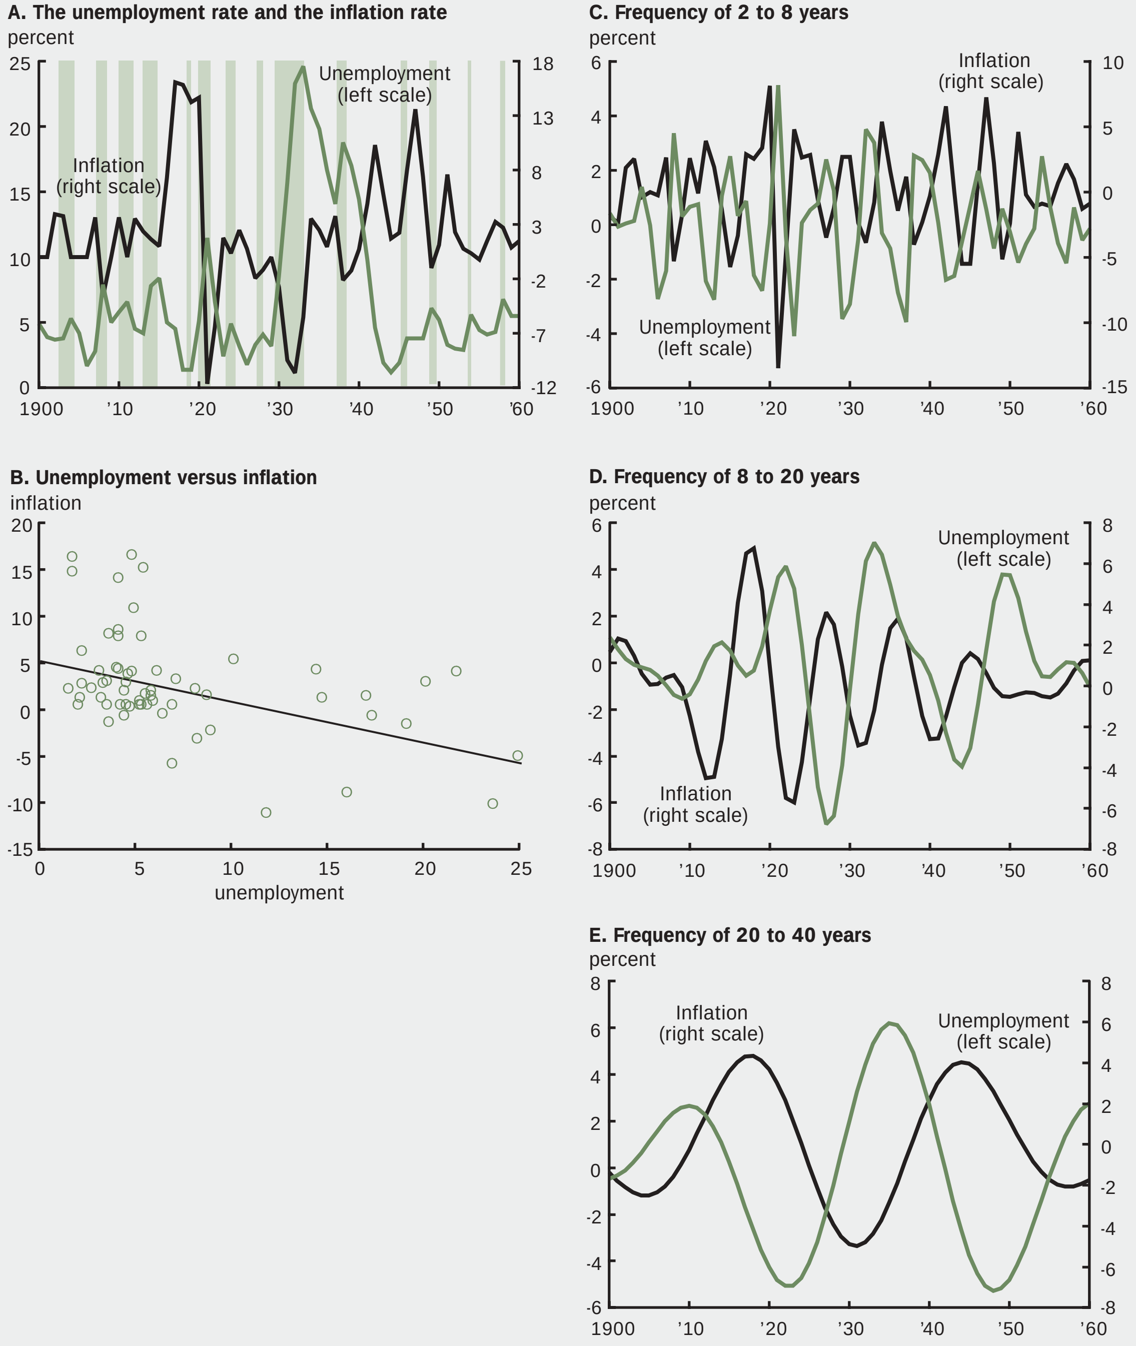


Note: Shaded areas indicate recessions as defined by the National Bureau of Economic Research. The black line indicates inflation and the green line indicates unemployment.

Source: Authors' calculations based upon data from the U.S. Department of Labor, Bureau of Labor Statistics.

FIGURE 3

**Unemployment and inflation, 1960–99**

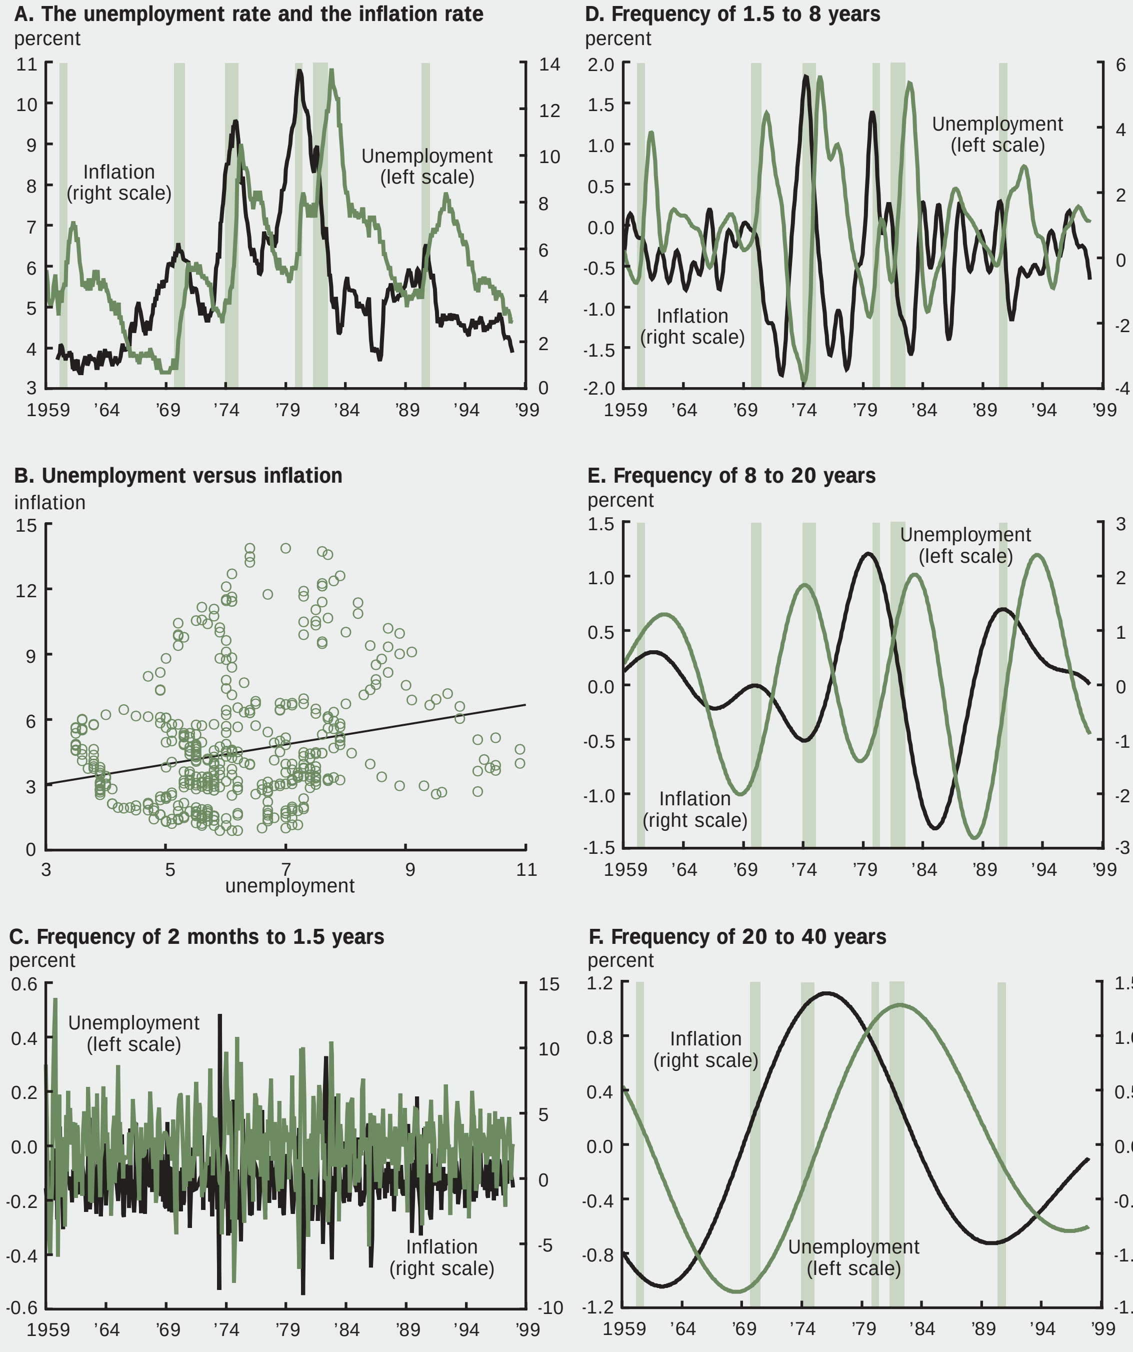


Note: Shaded areas indicate recessions as defined by the National Bureau of Economic Research. The black line indicates inflation and the green line indicates unemployment.
Source: Authors’ calculations based upon data from the U.S. Department of Labor, Bureau of Labor Statistics.

In the early period, the scale extends from –12 percent to +18 percent, at an annual rate. In the later sample, the scale extends over a smaller range, from 0 percent to 14 percent. In addition, the inflation data in the early period are characterized by sharp movements followed almost immediately by reversals in the other direction. By contrast, in the later dataset, movements in inflation in one direction are less likely to be reversed immediately by movements in the other direction.

***Money growth and inflation***

We report our results for money growth and inflation in detail in Christiano and Fitzgerald (2003), so here we just summarize the findings. We display these results in figure 4, panels A–E and figure 5, panels A–F. The style of analysis is much the same as for the unemployment and inflation data.

Consider the data from the early part of the century first. Figure 4, panel A shows that money growth (M2) and inflation move together very closely. The relationship appears to be essentially contemporaneous. This impression of a positive relationship is confirmed by the scatter plot between inflation and money growth in figure 4, panel B. To the eye, the positive relationship in figure 4, panel A appears to be a feature of all the frequency components of the data. This is confirmed in figure 4, panels C–E. Here we see the various frequency components of the data and how closely the data move together in each of them.

Now consider the data from the later part of the century. The raw data are reported in figure 5, panel A. The differences between these data in the early and late parts of the century are dramatic. At first glance, it may appear that the two variables, which moved together so closely in the early sample, are totally unrelated in the late sample. On closer inspection, the differences do not seem so great after all. Thus, in the very low frequencies there does still appear to be a positive relationship. Note how money growth generally rises in the first part of the late sample, and then falls in the second part. Inflation follows a similar pattern. It is in the higher frequencies that the relationship seems to have changed the most. Whereas in the early sample, the relationship between the two variables appeared to be contemporaneous, now there seems to be a significant lag. High money growth is not associated immediately with high inflation, but instead is associated with high inflation several years later. These observations, which are evident in the raw data, are confirmed by figure 5, panels B–F. Thus, panel B shows the scatter plot between money growth and inflation, which exhibits a positive relationship. Clearly, this positive relationship is dominated by the low frequency behavior of the data. It masks the very different behavior that we observe in the higher frequencies. Figure 5, panels D and E show how the variables are so far out of phase in the business cycle and lower frequencies that they actually have a negative relationship. The strong positive and contemporaneous relationship between the very low frequency components of the data that we noticed in figure 5, panel A, is quite evident in panel F.

## Implications of the evidence for macroeconomic models

The differences in the time series behavior of inflation in the first and second parts of the last century offer a potentially valuable source of information on the underlying mechanisms that drive inflation. For example, in the introduction, we talked about the recent literature that focuses on explaining the apparent *inertia* in inflation: the tendency for inflation to respond slowly to shocks. These findings are based on analysis of data from the second half of the century. We suspect that similar analysis of data for the first part of the century would find less inertia. This is because we saw that inflation is less persistent in the early sample, and its movements are more contemporaneous with movements in money. These observations provide a potentially important clue about how the private economy is put together: Whatever accounts for inflation inertia in the second part of the century must be something that was absent in the first part. For example, some have argued that frictions in the wage-setting process and variability in the rate of utilization of capital have the potential to account for the inflation inertia in postwar data.[14] If this is right, then wage-setting frictions must be smaller in the early sample, or there must have been greater limitations on the opportunities to achieve short-term variation in the utilization rate of capital.

The remainder of this section focuses on the change in the relationship between inflation and unemployment. At first glance, the change appears to lend support to the institutions view of inflation, as captured in the work of Kydland and Prescott (1977) and Barro and Gordon (1983a, b). A second glance suggests the evidence is not so supportive after all. Therefore, we begin with a brief review of the Barro-Gordon model.

***Barro–Gordon model***

The model comprises two basic relationships. The first summarizes the private economy. The second summarizes the behavior of the monetary authority. The private economy is captured by the *expectations-augmented Phillips curve*, originally associated with Friedman (1968) and Phelps (1967):

1) $u - u^N = -\alpha(\pi - \pi^e),\ \alpha > 0.$

BOX 1

## Formally testing our hypothesis about the Phillips curve

Formal tests of the hypothesis that the Phillips curve changed relatively little in the 2–20 year frequency range fail to reject it. Table B1 displays *p*-values for the null hypothesis that the post-1960s data on inflation and unemployment are generated by the bivariate vector autoregression (VAR) that generated the pre-1960s data. We implement the test using 2,000 artificial post-1960s datasets obtained by simulating a three-lag VAR and its fitted residuals estimated using the pre-1960s unemployment and inflation data.[1] In each artificial dataset, we compute correlations between filtered inflation and unemployment just like we did in the actual post-1960s data. Table B1 indicates that 9 percent of correlations between the business cycle component of inflation and unemployment exceed the –0.38 value reported in table 1 for the post-1960s data, so that the null hypothesis fails to be rejected at the 5 percent level. The *p*-value for the 8–20 year correlation is quite large and is consistent with the null hypothesis at any standard significance level.

The statistical evidence against the null hypothesis that there has been no change in the 20–40 year component of the data is also not strong. This may in part reflect a lack of power stemming from the relatively small amount of information in the sample about the 20–40 year frequency component of the data. But, the *p*-value may also be overstated for bias reasons. The table indicates that there is a small sample bias in this correlation, since the small sample mean, –0.35, is substantially larger than the corresponding probability limit of –0.45. A bias-adjustment procedure would adjust the coefficients of the estimated pre-1960s VAR so that the implied small sample mean lines up better with the pre-1960s empirical estimate of –0.51. Presumably, such an adjustment procedure would shift the simulated correlations to the *left*, reducing the *p*-value. It is beyond the scope of our analysis to develop a suitable bias adjustment method.[2] However, we suspect that, given the large magnitude of the bias, the bias-corrected *p*-value would be substantially smaller than the 14 percent value reported in the table.[3]

[1]We redid the calculations in table B1 using a five-lag VAR and found that the results were essentially unchanged. The only notable differences in the results are that the *p*-value for the business cycle correlations between inflation and unemployment is 0.06 and the *p*-value for these correlations in the 20–40 year range is 0.11.

[2]One could be developed along the lines pursued by Kilian (1998).

[3]To get a feel for the likely quantitative magnitude of the effects of bias adjustment, we redid the bootstrap simulations by adjusting the variance-covariance matrix of the VAR disturbances used in the bootstrap simulations. Let $V = [V]_{ij}$ denote the variance-covariance matrix. In the pre-1960s estimation results, $V_{1,2} = -0.1024$, $V_{1,1} = 0.0018$, $V_{2,2} = 6.0653$. When we set the value of $V_{1,2}$ to –0.0588 and recomputed the entries in table B1 in box 1, we found that the mean correlations were as follows: business cycle, –0.75 (0.01); 8–20 year, –0.54 (0.09); and 20–40 year, –0.51 (0.06). The numbers in parentheses are the analogs of the *p*-values in table B1. Note how the mean correlation in the 20–40 year frequency coincides with the empirical estimate reported in the first row of table 1, and that the *p*-value has dropped substantially, from 0.23 to 0.06. This is consistent with our conjecture that bias adjustment may have an important impact on the *p*-value for the 20–40 year correlation. However, the other numbers indicate that the bias adjustment procedure that we applied, by varying $V_{1,2}$ only, is not a good one. Developing a superior bias adjustment method is clearly beyond the scope of this article.

### TABLE B1
### Testing null hypothesis that post-1960s equal pre-1960s correlations

| Frequency | Plim | Small sample mean | Standard deviation, small sample mean | *p*-value |
|---|---|---|---|---|
| 2–8 year | –0.66 | –0.61 | $0.0036 \times \sqrt{2000}$ | 0.09 |
| 8–20 year | –0.36 | –0.38 | $0.0079 \times \sqrt{2000}$ | 0.25 |
| 20–40 year | –0.45 | –0.35 | $0.0129 \times \sqrt{2000}$ | 0.14 |

Notes: Data-generating mechanism in all cases is a three-lag, bivariate VAR fit to pre-1960s data. *p*-value: frequency, in 2,000 artificial post-1960s datasets, that contemporaneous correlation between the indicated frequency components of *x* and *y* exceeds, in absolute value, the corresponding post-1960s estimate. Plim: mean, over 1,000 artificial samples of length 2,000 observations each, of correlation. Small sample mean: mean of correlation, across 2,000 artificial post-1960s datasets. Standard deviation, small sample (product of Monte Carlo standard error for mean and $\sqrt{2000}$ ): standard deviation of correlations across 2,000 artificial post-1960s datasets.

FIGURE 4

**Measuring money growth and inflation, 1900–60**

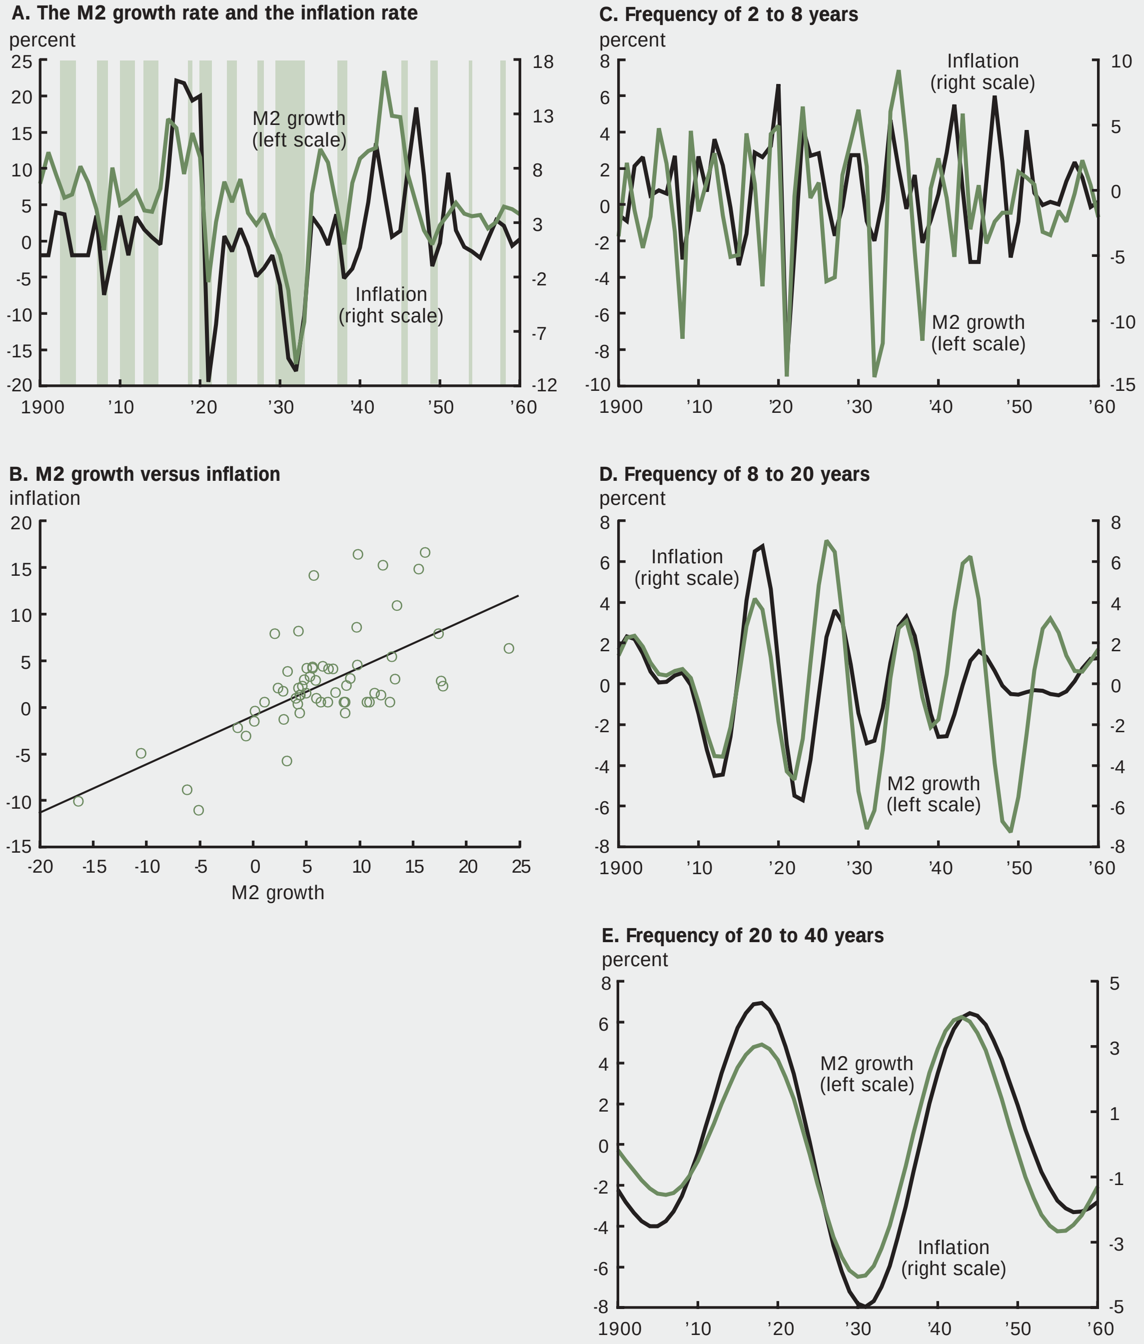


Note: Shaded areas indicate recessions as defined by the National Bureau of Economic Research.
Source: Authors' calculations based upon data from the Federal Reserve System and the U.S. Department of Labor, Bureau of Labor Statistics.

FIGURE 5

**Measuring money growth and inflation, 1960–99**

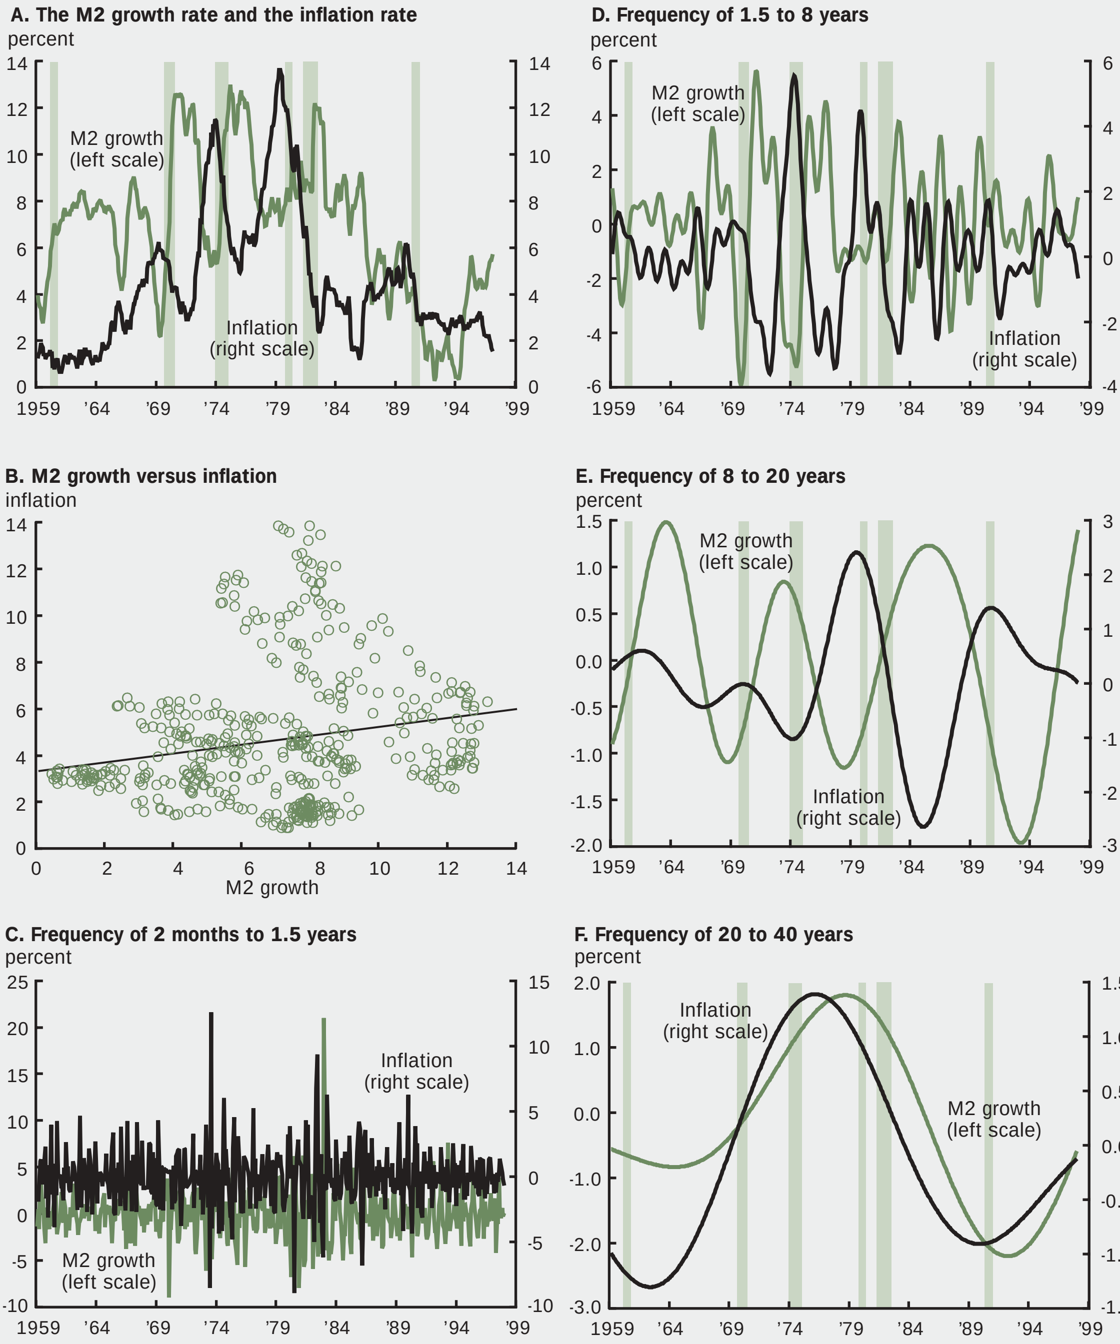


Note: Shaded areas indicate recessions as defined by the National Bureau of Economic Research.
Source: Authors' calculations based upon data from the Federal Reserve System and the U.S. Department of Labor, Bureau of Labor Statistics.

Here, $u$ is the actual rate of unemployment, $u^N$ is the *natural rate of unemployment*, $\pi$ is the actual rate of inflation, and $\pi^e$ is the rate of inflation expected by the private sector. The magnitude of $\alpha$ controls how much the actual rate of unemployment falls below its natural rate when inflation is higher than expected. The natural rate of unemployment is the unemployment rate that would occur if there was no surprise in inflation. The natural rate of unemployment is exogenous to the model, evolving in response to developments in unemployment insurance, social attitudes toward the unemployed, and other factors.

Note that according to the expectations augmented Phillips curve, if the monetary authority raises inflation above what people expected, then unemployment is below its natural rate. The mechanism by which this occurs is not explicit in the model, but one can easily imagine how it might work. For example, $\pi^e$ might be the inflation rate that is expected at the time wage contracts are set. Suppose that expectations of inflation are low, so that firms and workers agree to low nominal wages. Suppose that the monetary authority decides—contrary to expectations at the time wage contracts are written—to increase inflation by raising money growth. Given that wages in the economy have been pre-set at a low level, this translates into a low real wage, which encourages firms to expand employment and thereby reduce unemployment.[15]

The second part of the Barro–Gordon model summarizes the behavior of the monetary authority, which chooses $\pi$. Although the model does not specify the details of how this control is implemented, we should think of it happening via the monetary authority's control over the money supply. At the time that the monetary authority chooses $\pi$, the value of $\pi^e$ is predetermined. If the monetary authority can move $\pi$ above $\pi^e$, then, according to the expectations-augmented Phillips curve, unemployment would dip below the natural rate. It is assumed that the monetary authority wishes to push the unemployment rate below its natural rate, and this is captured by the notion that it would like to minimize:

2) $\frac{1}{2} [(u - ku^N)^2 + \gamma\pi^2], \gamma > 0, k < 1.$

The first term in parentheses indicates that, ideally, the monetary authority would like $u = ku^N < u^N$. The model does not specify exactly why the monetary authority wants unemployment below the natural rate. In principle, there are various factors that could rationalize this. For example, the presence of distortionary taxes or monopoly power could make the level of economic activity inefficiently low, and this might translate into a natural rate of unemployment that is suboptimally high.

In practice, the monetary authority would not necessarily go for the ideal level of unemployment, because the increase in $\pi$ that this requires entails costs. These are captured by the $\gamma\pi^2$ term in the objective. According to this term, the ideal level of inflation is zero.[16] The higher the level of inflation, the higher the marginal cost.

The Barro–Gordon model views the monetary authority as choosing $\pi$ to optimize its objective, subject to the expectations-augmented Phillips curve and to the given value of $\pi^e$. The optimal choice of $\pi$ reflects a balancing of the benefits and costs summarized in the monetary authority's objective. A graph of the best response function appears in figure 6, where $\pi^e$ appears on the horizontal axis, and $\pi$ appears on the vertical. The 45-degree line in the figure conveniently shows the level of inflation that the policymaker would select if it chose to validate private expectations of inflation.

Note how the best response function is flatter than the 45-degree line. This reflects the increasing marginal cost of inflation at higher levels of inflation. At low levels of expected inflation, the marginal cost of inflation is low, so the benefits outweigh the costs. At such an inflation rate, the monetary authority would try to surprise the economy by moving to a higher level. On the other hand, if expected inflation were very high, then the marginal cost of going even higher would outweigh the benefits, and the monetary authority would choose to violate expectations by choosing a lower inflation rate. Not surprisingly, there is an inflation rate in the middle, $\pi^*$, where the monetary authority chooses not to surprise the economy at all. This is the inflation rate where the best response function crosses the 45-degree line. Because of the linear nature of the expectations-augmented Phillips curve and the quadratic form of monetary authority preferences, the best response function is linear, guaranteeing that there is a single crossing.

What is equilibrium in the model? We assume everyone—the monetary authority and the private economy—is rational. In particular, the private economy understands the monetary authority's policymaking process. It knows that if it were to have expectations, $\pi^e < \pi^*$, then actual inflation would be higher than $\pi^e$. So, it cannot be rational to have an expectation like this. It also understands that if it were to have expectations, $\pi^e > \pi^*$, the monetary authority would choose an inflation rate lower than $\pi^e$. So, this expectation cannot be rational either. The only rational thing for the private economy to expect is $\pi^*$. So, this is equilibrium in the model. The formula for this is

3) $$\pi = \psi u, \psi = \frac{\alpha(1-k)}{\gamma} > 0.$$

FIGURE 6

**Effect on macroeconomic models**

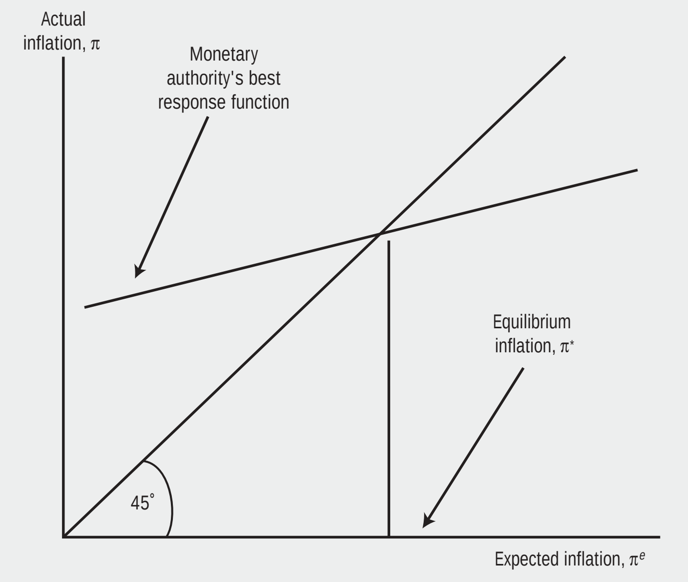


According to the model, inflation is predicted to be proportional to the actual level of unemployment. There are several crucial things to note here. First, the actual level of unemployment is equal to the natural rate, because in equilibrium the monetary authority cannot surprise the private economy. So, monetary policy in practice does not succeed in driving unemployment below the natural rate at all. Second, inflation is positive, being proportional to unemployment. This is higher than its ideal level, here presumed to be zero. These two observations imply that in equilibrium, all the monetary authority succeeds in doing is producing an inflation rate above its ideal level. It makes no headway on unemployment. That is, this optimizing monetary authority simply succeeds in producing suboptimal outcomes. How is this possible?

The problem is that the monetary authority lacks the ability to commit to low inflation. At the time the monetary authority makes its decision, the private economy has already formed its expectation about inflation. The private economy knows that if it expects inflation to occur at the socially optimal level, $\pi^e = 0$, then the monetary authority has an incentive to deviate to a higher level of inflation (see figure 6).[17]

Eggertsson (2001) has recently drawn attention to one of Aesop's fables, which captures aspects of the situation nicely. Imagine a lion that has fallen into a deep pit. Unless it gets out soon, it will starve to death. A rabbit shows up and the lion implores the rabbit to push a stick lying nearby into the hole, so that the lion can climb out. The lion cries out from the depths of its soul, with a most solemn commitment not to eat the (juicy-looking) rabbit once it gets out. But, the rabbit is skeptical. It understands that the intentions announced by the lion while in the hole are not time consistent. While in the hole, the lion has the incentive to declare, with complete sincerity, that it will not eat the rabbit when it gets out. However, that plan is no longer optimal for the lion when it is out of the hole. At this point, the lion's optimal plan is to eat the rabbit after all. The rational rabbit, who understands the time inconsistency of the lion's optimal plan, would do well to leave the lion where it is. What the lion would like while it is in the hole is a commitment technology: something that convinces the rabbit that the lion will have no incentive or ability to change the plan it announces from the hole after it is out.

In some respects, the rabbit and the lion resemble the private economy and the monetary authority in the Barro–Gordon model. Before $\pi^e$ is chosen, the monetary authority would like people to believe that it will choose $\pi = 0$. The problem is that after the private economy sets $\pi^e = 0$, the monetary authority has an incentive to choose $\pi > \pi^e$ (see figure 6). As in the fable, what the monetary authority needs is some sort of commitment technology, something that convinces private agents that if they set $\pi^e = 0$, the monetary authority has no incentive or ability to deviate to $\pi > 0$. Rational agents in an economy where the monetary authority has no such commitment technology do well to set $\pi^e = \pi^* > 0$. This puts the monetary authority in the dilemma discussed in the introduction. Its optimal choice in this case is to validate expectations by setting $\pi = \pi^*$ (that is, it chooses C in figure 1).

The crucial point of Kydland–Prescott and Barro–Gordon is that if the monetary authority has a credible commitment to low inflation, then better outcomes would occur than if it has no such ability to commit. In both cases, the same level of unemployment occurs (that is, the natural rate), but the authority with commitment achieves the ideal inflation rate, while the monetary authority without commitment achieves a socially suboptimal higher inflation rate. The problem, as with the lion in the fable, is coming up with a credible commitment technology. The commitment technology must be such that the monetary authority actually has no incentive to select a high inflation rate after the private economy selects $\pi^e$.

What makes adopting a commitment technology particularly difficult is that the monetary authority's preferences in Barro–Gordon (unlike the lion's preferences in the fable) are fundamentally democratic preferences: They reflect actual social costs and benefits. Credible commitment technologies must involve basic changes in monetary institutions, which make them, in effect, less democratic. Changes that have been adopted in practice are the legal and other mechanisms that make central banks independent from the administrative and legislative branches of government. The classic institutional arrangement used to achieve commitment has been the gold standard. Tying the money supply to the quantity of gold greatly limits the ability of the central bank to manipulate $\pi$.

***Barro–Gordon and the data***

The Barro–Gordon model is surprisingly effective at explaining key features of the inflation–unemployment relationship during the twentieth century. It is perhaps reasonable to suppose that the U.S. monetary authorities more closely resembled the monetary authority with commitment in the Barro–Gordon model in the early part of the last century and more closely resembled the monetary authority without commitment in the last part of the century. After World War II, the U.S. government resolved that all branches of government—including the Federal Reserve—should be committed to the objective of full employment. This commitment reflected two views. The first view, apparently validated by the experience of the Great Depression, is that activist stabilization policy is desirable. It was codified into law by the Full Employment Act of 1946. The second view, associated with the intellectual revolution of John Maynard Keynes, is that successful activist stabilization policy is feasible. This view was firmly entrenched in Washington, DC, by the time of the arrival of the Kennedy administration in 1960. Kennedy's Council of Economic Advisors resembles a "who's who" of Keynesian economics.[18]

The notion that policymakers were committed to low inflation in the early part of the century and relatively more concerned with economic stabilization later implies, via the Barro–Gordon model, that inflation in the late period should have been higher than it was in the early period. Comparison of figure 4, panel A and figure 5, panel A shows that this is indeed the case. Another implication of the model is that inflation should have been constant at zero in the early period, and this most definitely was *not* the case (see figure 4, panel A).[19] But, this is not a fundamental problem for the model. There is a simple, natural timing change in the model that eliminates this implication, without changing the central message of the analysis in the previous section. In particular, suppose that the actions of the central bank have an impact on inflation only with a $p$-period delay with $p > 0$. In this way, the monetary authority is not able to eliminate the immediate impact of shocks to the inflation rate. The policymaker with commitment sets the $p$-period-ahead expected inflation rate to zero. Suppose that the analogous timing assumption applies to the private sector, so that there are movements in inflation that are not expected at the time it sets $\pi^e$. Under the expectations-augmented Phillips curve, this introduces a source of negative correlation between inflation and unemployment. This sort of delay in the private sector could be rationalized if wage contracts extended over $p$ periods of time. Under these timing assumptions, the prediction of the model under commitment is that the actual inflation rate fluctuates, and inflation and unemployment covary negatively, as was actually observed over the early part of the twentieth century. (The appendix analyzes the model with time delays.) When the monetary authorities drop their commitment to low inflation in the later part of the century, the model predicts that unemployment and inflation move together more closely and that the relationship will actually be positive in the lowest frequencies. In the higher frequencies, the correlation might still be negative, for the reason that it is negative in all frequencies when there is commitment: Inflation in the higher frequencies is hard to control when there are implementation delays.[20] In this sense, the Barro–Gordon models seems at least qualitatively consistent with the basic facts about what happened to the inflation–unemployment relationship between the first and second parts of the past century. It is hard not to be impressed by this.[21]

But, there is one shortcoming of the model that may be of some concern. Recall from figure 3, panel A that inflation in the early 1980s dropped precipitously, just as unemployment soared to a postwar high. This behavior in inflation and unemployment is so pronounced that it has a substantial impact on the very low frequency component of the data. According to figure 3, panel F, the 20–40 year component of unemployment lags the corresponding component of inflation by several years. As a technical matter, it is possible to square this with the model. The version of the model discussed in the previous paragraph allows for the possibility that a big negative shock to the price level—one that was beyond the control of the monetary authority—occurred that drove actual unemployment up above the natural rate of unemployment. But the explanation rings hollow. The model itself implies that, on average, the low frequency component of unemployment *leads* inflation, not the other way around (see the appendix for an elaboration). This is because unemployment is related to the incentives to inflate, so when unemployment rises, one expects inflation to rise in response. In fact, with the implementation and observation delays, one expects the rise in inflation to occur with a delay after a rise in unemployment.

In sum, the Barro–Gordon model seems to provide a way to understand the change in inflation–unemployment dynamics between the first and second parts of the last century. However, the disinflation of the early 1980s raises some problems for the model. That experience appears to require thinking about the deflation of the early 1980s as an accident. Yet, to all direct appearances it was no accident at all. Conventional wisdom takes it for granted that the disinflation was a direct outcome of intentional efforts taken by the Federal Reserve, beginning with the appointment of Paul Volcker as chairman in 1979. Many observers interpret this experience as a fundamental embarrassment to the Barro–Gordon model. Some would go further and interpret this as an embarrassment to the ideas behind it: the notion that time inconsistency is important for understanding the dynamics of U.S. inflation. They argue that, according to the model, the only way inflation could fall precipitously absent a drop in unemployment is with substantial institutional reform to implement commitment. There was no institutional reform in the early 1980s, so the institutional perspective must, at best, be of second-order importance for understanding U.S. inflation.

***Alternative representation of the notion that commitment matters***

By the standards of our times, the Barro–Gordon model must be counted a massive success. Its two simple equations convey some of the most profound ideas in macroeconomics. In addition, it accounts nicely for broad patterns in twentieth century data: the fact that inflation on average was higher in the second half, and the changed nature of the unemployment–inflation relationship.

Yet, the model encounters problems understanding the disinflation of the 1980s. Perhaps this is a problem for the specific equations of the model. But, is it a problem for the *ideas* behind the model? We just do not know yet, because the ideas have not been studied in a sufficiently wide range of economic models. Efforts to incorporate the basic ideas of Kydland–Prescott and Barro–Gordon into modern models have only just begun. This process has been slow, in part because the computational challenge of this task is enormous. Indeed, the computational difficulties of these models serve as another reminder of the power of the original Barro–Gordon model: With it, the reader can reach the core ideas armed simply with a sheet of paper and a pencil.

Why should we incorporate the ideas into modern models? First, the ideas have proved enormously productive in helping us understand the broad features of inflation in the twentieth century. This suggests that they deserve further attention. Second, as we will see below, when we do incorporate the ideas into modern models, unexpected results occur. They may provide additional possibilities for understanding the data. Third, because modern models are explicitly based on microfoundations, they offer opportunities for econometric estimation and testing that go well beyond what is possible with the original Barro–Gordon model. In modern models, crucial parameters like $\alpha$, $k$, and $\gamma$ are related explicitly to production functions, to features of labor and product markets, to properties of utility functions, and to the nature of information transmission among agents. These linkages make it possible to bring a wealth of data to bear, beyond data on just inflation and unemployment. In the original Barro–Gordon model, $\alpha$, $k$, and $\gamma$ are primitive parameters, so the only way to obtain information on them is using the data on inflation and unemployment itself.

To see the sort of things that can happen when the ideas of Kydland–Prescott and Barro–Gordon are incorporated into modern models, we briefly summarize some recent work of Albanesi, Chari, and Christiano (2002).[22] They adapt a version of the classic monetary model of Lucas and Stokey (1983), so that it incorporates benefits of unexpected inflation and costs of inflation that resemble the factors Barro and Gordon appeal to informally to justify the specification of their model. However, because the model is derived using standard specifications of preferences and technology, there is no reason to expect that the monetary authority's best response function is linear, as in the Barro–Gordon model (recall figure 6). Indeed, Albanesi, Chari, and Christiano find that for almost all parameterizations for the model, if there is any equilibrium at all there must be two. That is, the best response function is nonlinear, and has the shape indicated in figure 7.

In one respect, it should not be a surprise that there might be multiple equilibriums in a Barro–Gordon type model. Recall that an equilibrium is a level of inflation where benefits of additional unexpected inflation just balance the associated costs. But we can expect that these costs and benefits change nonlinearly for higher and higher levels of inflation. If so, then there could be multiple levels of inflation where equilibrium occurs, as in figure 7.

There is one version of the Albanesi–Chari–Christiano model in which the intuition for the multiplicity is particularly simple. In that version, private agents can, at a fixed cost, undertake actions to protect themselves against inflation. In principle, such actions may involve acquiring foreign currency deposits for use in transactions. Or, they may involve fixed costs of retaining professional assistance in minimizing cash balances when inflation is high. Although these efforts are costly for individuals, they do mean that on the margin, the costs of inflation are reduced from the perspective of a benevolent monetary authority. Turning to figure 7, one might imagine that at low levels of inflation, the basic Barro–Gordon model applies. People do not undertake fixed costs to protect themselves against inflation, and the best response function looks roughly linear, cutting the 45-degree line at the lower level of inflation indicated in the figure. At higher levels of inflation, however, people do start to undertake expensive fixed costs to insulate themselves. By reducing the marginal cost of inflation, this has the effect of increasing the incentive for the monetary authority to raise inflation. Of course, this assumes that the benefits of inflation do not simultaneously decline. In the Albanesi–Chari–Christiano model, in fact they do not decline. This is why in this version of their model, the best response function eventually begins to slope up again and, therefore, to cross the 45-degree line at a higher level of inflation.

FIGURE 7

**Variant of the Lucas–Stokey model**

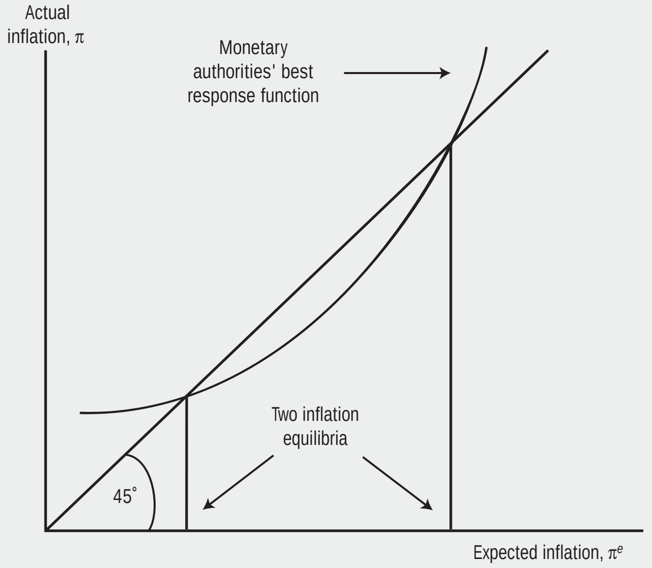


The previous example is designed to just present a flavor of the Albanesi–Chari–Christiano results. In fact, the shape of the best response function resembles qualitatively the picture in figure 7, even in the absence of opportunities for households to protect themselves from inflation.

What are the implications of this result? Essentially, there are new ways to understand the fact that inflation is sometimes persistently high and at other times (like now) persistently low. In the Barro–Gordon model, this can only be explained by appealing to a fundamental variable that shifts the best response function. The disinflation of the early 1980s suggests that it may be hard to find such a variable in practice.

But, is a model with multiple equilibriums testable? Perhaps. Inspection of figure 7 suggests one possibility. Shocks to the fundamental variables that determine the costs and benefits of inflation from the perspective of the monetary authority have the effect of shifting the best response curve up and down. Notice how the high-inflation equilibrium behaves differently from the low-inflation equilibrium as the best response function, say, shifts up. Inflation in the low-inflation equilibrium rises, and in the high-inflation equilibrium it falls. Thus, these shocks have an opposite correlation with inflation in the two equilibriums. This sign switch in equilibriums is an implication of the model that can, in principle, be tested. For example, Albanesi–Chari–Christiano explore the model's implication that interest rates and output covary positively in the low-inflation equilibrium and negatively in the high-inflation equilibrium. Using data drawn from over 100 countries, they find evidence in support of this hypothesis.

But, the Albanesi–Chari–Christiano model is still too simple to draw final conclusions about the implications of lack of commitment for the dynamics of inflation. The model has been kept very simple so that—like the Barro-Gordon model—it can be analyzed with a sheet of paper and a pencil (well, perhaps one would need two sheets of paper!). We know from separate work on problems with a similar logical structure that when models are made truly dynamic, say with the introduction of investment, the properties of equilibriums can change in fundamental ways (see, for example, Krusell and Smith, 2002). It still remains to explore the implications of lack of commitment in such models. In particular, it is important to explore whether the disinflation experience of the early 1980s, which appears to be a problem for the Barro–Gordon model, can be reconciled with modern models.

## Conclusion

We characterized the change in the nature of inflation dynamics before and after the 1960s. We reviewed various theories about inflation, but put special focus on the *institutions* view: theories that focus on lack of commitment in monetary policy as the culprit behind bad inflation outcomes. We argued that this view, as captured in the famous model of Barro and Gordon (1983a, b), accounts well for the broad outlines of the data. Not only does it capture the fact that inflation was, on average, lower in the early period of the twentieth century than in the later period, but it also accounts for the shift that occurred in the unemployment–inflation dynamics. In the early period, inflation and unemployment exhibit a negative relationship at all frequency bands. In the later period, the negative relationship persists in the higher frequency bands, while a positive relationship emerges in the low frequencies. We show how the Barro–Gordon model can account for this shift as reflecting the notion that monetary policy was credibly committed to low inflation in the early period, while it abandoned that commitment in the later period.

Although the model does well on these broad facts, it has some well-known difficulties addressing the disinflation in the U.S. in the 1980s. This, among other considerations, motivates the recent research on the implications of absence of commitment in monetary policy. We show that that research uncovers some surprising—relative to the original Barro–Gordon analysis—implications of lack of commitment. These may ultimately prove helpful for achieving a better model of inflation dynamics. But that research has a long way to go, before we fully understand the implications of absence of commitment in monetary policy.

What is at stake in this work? If absence of commitment is in fact the primary reason for the poor inflation outcomes of the past, then research on ways to improve inflation outcomes needs to focus on improved design of monetary institutions.

NOTES

[1]This belief is based in part on the evidence (see, for example, Barsky and Kilian, 2000, for a discussion of the role of money growth in the 1970s inflation). But, it is also based on the view that good economic theory implies a close connection—at least over horizons as long as a decade—between money growth and inflation. Recently, some economists' confidence in the existence of a close connection between money growth and inflation has been shaken by the discovery, in seemingly well-specified economic models, that the connection can be surprisingly weak. For example, Loyo (1999) uses the "fiscal theory of the price level" to argue that it was a high nominal interest rate that initiated the rise in inflation in Brazil, and that this rise in the interest rate was in a meaningful sense not "caused" by high money growth. Loyo drives home his point that it was not high money growth that caused the high inflation by articulating it in a model in which there is no money. For a survey of the fiscal theory, and of Loyo's argument in particular, see Christiano and Fitzgerald (2000). Others argue that standard economic theories imply a much weaker link than was once thought, between inflation and money growth. For example, Benhabib, Schmitt-Grohé, and Uribe (2001a, b) and Krugman (1998) argue that it is possible for there to be a deflation even in the presence of positive money growth. Christiano and Rostagno (2001) and Christiano (2000) review these arguments, respectively. In each case, they argue that the deflation, high money growth scenario depends on implausible assumptions.

[2]This description of economists' research strategy is highly stylized. In some cases, the model is not made formally explicit. In other cases, the model is explicit, but the data plays only a small role in building confidence in the model.

[3]Prominent recent papers that draw attention to the inertia puzzle include Chari, Kehoe, and McGrattan (2000), and Mankiw (2001). Christiano, Eichenbaum, and Evans (2001) describe variants of standard macroeconomic models that can account quantitatively for the inertia.

[4]The first, second, and aspects of the fourth observations have been made before. To our knowledge the third observation was first made in Christiano and Fitzgerald (2003). For a review of the first two observations, see Blanchard (2002). For a discussion of the fourth using data on the second half of the twentieth century, see King and Watson (1994), King, Stock, and Watson (1995), Sargent (1999), Staiger, Stock, and Watson (1997), and Stock and Watson (1998).

[5]In this respect, our analysis resembles that of Ireland (1999), although his analysis focuses on data from the second half of the twentieth century only, while we analyze both halves.

[6]For a critical review of the Clarida, Gali, and Gertler argument, see Christiano and Gust (2000). Other arguments that fall into what we are calling the people category include Sargent (1999). Sargent argues that periodically, the data line up in such a way that there appears to be a Phillips curve with a favorable trade-off between inflation and unemployment. High inflation then results as the central bank attempts to exploit this to reduce unemployment. As emphasized in Sargent (1999, chapter 9), the high inflation of the 1970s represents a challenge for this argument. This is because the dominant fact about the early part of this decade was the apparent "death" of the Phillips curve: Policymakers and students of the macroeconomy were stunned by the fact that inflation and unemployment both increased at the time.

[7]The different frequency components of the data are extracted using the band pass filter method summarized in Christiano and Fitzgerald (1998) and explained in detail in Christiano and Fitzgerald (2003).

[8]It is worth emphasizing that, by "Phillips curve," we mean a statistical relationship, and not necessarily a relationship exploitable by policy.

[9]The slope of the regression line drawn through the scatter plot of points in figure 2, panel B is –0.42, with a *t*-statistic of 3.77 and an $R^2$ of 0.20.

[10]Specifically, they are *p*-values for testing the null hypothesis that there is no relationship at any frequency between the two variables, against the alternative that the correlation is in fact the one reported in the table. These *p*-values are computed using the following bootstrap procedure. We fit separate *q*-lag scalar autoregressive representations to the level of inflation (first difference, log CPI) and to the level of the unemployment rate. We used random draws from the fitted disturbances and actual historical initial conditions to simulate 2,000 artificial datasets on inflation and unemployment. For annual data, $q = 3$; for monthly, $q = 12$; and for quarterly, $q = 8$. The datasets on unemployment and inflation are independent by construction. In each artificial dataset, we compute correlations between the various frequency components, as we did in the actual data. In the data and the simulations, we dropped the first and last three years of the filtered data before computing sample correlations. The numbers in parentheses in table 1 are the frequency of times that the simulated correlation is greater than the estimated correlation is positive. If it is negative, we compute the frequency of times that the simulated correlation is less than the simulated value. These are *p*-values under the null hypothesis that there is no relationship between the inflation and unemployment data.

[11]Figure 3 exhibits monthly observations on inflation and unemployment. To reduce the high frequency fluctuations in inflation, figure 3, panel A exhibits the annual average of inflation, rather than the monthly inflation rate. The scatter plot in figure 3, panel B is based on the same data used in figure 3, panel A. Figure 3, panels C–F are based on monthly inflation, that is, $1{,}200\log(CPI_{t-1}/CPI_{t-1})$, and unemployment. The line in figure 3, panel B represents a regression line drawn through the scatter plot. The slope of that line, based on monthly data covering the period 1959:Q2–98:Q1, is 0.47 with a *t*-statistic of 5.2.

[12]Consistent with these observations, when inflation and unemployment are detrended using a linear trend with a break in slope (not level) in 1980:Q4 for inflation and 1983:Q1 for unemployment, the scatter plots of the detrended variables show a negative relationship. The regression of detrended inflation on detrended unemployment has a coefficient of –0.31, with *t*-statistic of –4.24 and $R^2 = 0.037$. The slope coefficient is similar to what was obtained in note 9 for the pre-1960s period, but the $R^2$ is considerably smaller.

[13]See King and Watson (1994), Stock and Watson (1998), and Sargent (1999, p. 12), who apply the band-pass filtering techniques proposed in Baxter and King (1999). The relationship between the Baxter–King band-pass filtering methods and the method used here is discussed in Christiano and Fitzgerald (2003).

[14]See, for example, Christiano, Eichenbaum, and Evans (2001).

[15]In the years since the expectations-augmented Phillips curve was first proposed, evidence has accumulated against it. For example, Christiano, Eichenbaum, and Evans (2001) display evidence that suggests that inflation surprises are not the mechanism by which shocks, including monetary policy shocks, are transmitted to the real economy. Although the details of the mechanism underlying the expectations-augmented Phillips curve seem rejected by the data, the basic idea is still very much a part of standard models. Namely, it is the unexpected component of monetary policy that impacts on the economy via the presence of some sort of nominal rigidity.

[16]Extending the analysis to the case where the socially optimal level of inflation is non-zero (even, random) is straightforward.

[17]In later work, Barro and Gordon (1983a) pointed out that there exist equilibriums in which reputational considerations play a role. In such equilibriums, a monetary authority might choose to validate $\pi^e = 0$ out of concern that if it does not do so, then in the next period $\pi^e$ will be an extremely large number with the consequence that whatever they do then, the social consequences will be bad. In this article, we do not consider these “trigger strategy” equilibriums, and instead limit ourselves to Markov equilibriums, in which decisions are limited to be functions only of the economy’s current state. In the present model, there are no state variables, and so decisions, $\pi^e$ and $\pi$, are simply constants. A problem with allowing the presence of reputational considerations is that they support an extremely large set of equilibriums. Essentially, anything can happen and the theory becomes vacuous.

[18]It would be interesting to understand why earlier monetary authorities were relatively less concerned with stabilizing the economy and more committed, for example, to the gold standard.

[19]As mentioned in an earlier note, the model does not require that the optimal level of inflation is literally zero. Implicitly, what we are assuming is that the optimal level of inflation, $\pi^o$ in the note, is much smoother than the inflation rate actually observed in the early sample.

[20]These observations are established in the appendix.

[21]The argument we have just made is similar in spirit to the one that appears in Ireland (1999).

[22]This builds on previous work by Chari, Christiano, and Eichenbaum (1998).

APPENDIX: INFLATION–UNEMPLOYMENT COVARIANCE FUNCTION IN THE IMPLEMENTATION-DELAY VERSION OF BARRO–GORDON MODEL

This appendix works out the covariance implications of a version of the Barro–Gordon model with implementation delays (implementation delays are discussed in Barro and Gordon [1983, pp. 601–602]). The particular version we consider is the one proposed in Ireland (1999). We work out the model's implications for the type of frequency-domain statistics analyzed in the text. In particular, we seek the covariance properties of inflation and unemployment, when we consider only a specified subset of frequency components (high, business cycle, low, and very low) of these variables.

We obtain two sets of results. One pertains to the commitment version of the model and the other to the no-commitment version:

- It is possible to parameterize the commitment version of the model so that the covariance between inflation and unemployment is negative for all subsets of frequency components.
- In the no-commitment version of the model, the covariance between inflation and unemployment can be positive in the very low frequency components of the data and negative in the higher frequency components. Unemployment does not lag inflation in the very low frequency data, and it may actually lead, depending on parameter values.

The idea is that policymakers can only influence the $p$-period ahead forecast of inflation, not actual inflation. With this change, the objective of the policymaker is $E_{-p}\,[(u - ku^N)^2 + \gamma\pi^2]/2$. Actual inflation, $\pi$, is $\pi = \hat{\pi} + \theta * \eta$, where $\hat{\pi}$ is a variable chosen $p \geq 0$ periods in the past by the policymaker, and $\theta * \eta$ captures the shocks that impact $\pi$ between the time $\hat{\pi}$ is set and $\pi$ is realized. Here,

$$\theta * \eta_t = \begin{cases} \left(\theta_0 + \theta_1 L + \ldots + \theta_{p-1} L^{p-1}\right)\eta_t & p \geq 1 \\ 0 & p = 0 \end{cases},$$

where $\eta_t$ is white noise and $L$ is the lag operator, $L^j\eta_t = \eta_{t-j}$. The policymaker's problem is optimized by setting $\hat{\pi} = \psi\hat{u}^N$, where $\hat{u}^N$ is the forecast of the period $t$ natural rate of unemployment, made $p$ periods in the past, $\hat{u}_t^N = E_{t-p}u_t^N$, computable at the time $\hat{\pi}$ is selected and $x$ is defined in the text. Following Ireland (1999), we suppose that $u^N$ has a particular unit root time series representation:

$$(1-L)u_t^N = \lambda(1-L)u_{t-1}^N + \nu_t \quad -1 < \lambda < 1.$$

With this representation,

$$\hat{u}_t^N = g * \nu_t = g(L)\nu_t,$$

where

$$g(L) = L^p\left[\frac{1-\lambda^{p+1}}{1-\lambda}\frac{1}{1-\lambda L} + \frac{L}{(1-\lambda L)(1-L)}\right].$$

We suppose that $\pi^e$ in the expectations augmented Phillips curve is the *p*-period ahead forecast of inflation made by private agents. We impose rational expectations, $\pi^e = \hat{\pi}$. Then, it is easy to verify that when there is no commitment, inflation and unemployment evolve in equilibrium according to

$$\pi_t = \psi g(L)\nu_t + \theta(L)\eta_t,\ u_t = \frac{1}{(1-\lambda L)(1-L)}\nu_t - \alpha\theta(L)\eta_t,$$

respectively. We make the simplifying assumption that all shocks are uncorrelated with each other. Outcomes when there is commitment are found by replacing $\psi$ in the above expression with 0. In this case, it is easy to see that the covariance between inflation and unemployment is unambiguously negative. Under no commitment, it is possible for this correlation to be positive.

It is convenient to express the joint representation of the variables as follows:

$$x_t \equiv \begin{pmatrix} u_t \\ \pi_t \end{pmatrix} = F(L)\begin{pmatrix} \nu_t \\ \eta_t \end{pmatrix},$$

where

$$F(L) = \begin{bmatrix} \frac{1}{(1-\lambda L)(1-L)} & -\alpha\theta(L) \\ \psi g(L) & \theta(L) \end{bmatrix}.$$

Denote the covariance function of $x_t$ by

$$c(k) = Ex_t x'_{t-k} = \begin{bmatrix} Eu_t u_{t-k} & Eu_t \pi_{t-k} \\ E\pi_t u_{t-k} & E\pi_t \pi_{t-k} \end{bmatrix},$$

for $k = 0, \pm 1, \pm 2, \ldots$ . We want to understand the properties of the covariance function, $E\tilde{\pi}_t \tilde{u}_{t-1}$, where $\tilde{\pi}_t$ is the component of $\pi_t$ in a subset of frequencies, and $\tilde{u}_t$ is the component of $u_t$ in the same subset of frequencies.

For this, some results in spectral analysis are useful (see Sargent [1987, chapter 11], or, for a simple review, see Christiano and Fitzgerald [1998]). The spectral density of a stochastic process at frequency $\omega \in (-\pi, \pi)$ is the Fourier transform of its covariance function:

$$S(\omega) = \sum_{j=-\infty}^{\infty} c(k) e^{-\omega j}.$$

The covariances can then be recovered applying the inverse Fourier transform to the spectral density:

$$c(l) = \frac{1}{2\pi}\int_{-\pi}^{\pi} S(\omega)\, e^{\omega l}.$$

It is trivial to verify the latter relationship, using the definition of the spectral density and the fact

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} e^{\omega l} d\omega = \begin{cases} 0, l \neq 0 \\ 1, l = 0 \end{cases}.$$

The inverse Fourier transform result is convenient for us, because in practice there exists a very simple, direct way to compute $S(\omega)$.

Let $S(\omega)$ denote the spectral density of $x_t$, after a band-pass filter has been applied to $x_t$ to isolate a subset of frequencies. Then,

$$S(\omega) = F(e^{-i\omega})VF(e^{i\omega})',$$

where $V$ is the variance-covariance matrix of $(\nu_t, \eta_t)$. Here, $V = [V_{ij}]$ and $V_{11} = E\nu_t^2$, $V_{12} = E\nu_t\eta_t$, $V_{22} = \eta_t^2$. Evaluating the 2,1 element of $S(\omega)$, which we denote $S_{\pi u}(\omega)$:

$$S_{\pi u}(\omega) = \frac{\psi g(e^{-i\omega})}{(1-\lambda e^{i\omega})(1-e^{i\omega})} V_{11} + \left[\frac{\theta(e^{-i\omega})}{(1-\lambda e^{i\omega})(1-e^{i\omega})} - \psi\alpha g(e^{-i\omega})\theta(e^{i\omega})\right]V_{12} - \alpha\theta(e^{-i\omega})\theta(e^{i\omega})V_{22}.$$

Then,

$$E\pi_t u_{t-l} = \frac{1}{2\pi}\int_{-\pi}^{\pi} S_{\pi u}(\omega)\, e^{i\omega l} d\omega$$

$$= \frac{1}{2\pi}\int_{0}^{\pi} s(\omega, l)\, d\omega,$$

where

$$s(\omega,l) = S_{\pi u}(\omega)e^{i\omega l} + S_{\pi u}(-\omega)e^{-i\omega l},\ \omega \in (-\pi, \pi).$$

There are two features of the covariance function, $E\pi_t u_{t-l}$, that we wish to emphasize. First, in the case of commitment, when $\psi$ is replaced by 0 in $S_{\pi u}(\omega)$, it is possible to choose parameters so that $E\pi_t u_{t-l} \leq 0$ for all $l$, over all possible subsets of frequencies. Consider, for example, $V_{12} = 0$, $p = 1$ and $\theta(e^{-i\omega}) = \theta > 0$, so that

$$E\pi_t u_{t-1} = -\frac{1}{2\pi}\int_{0}^{\pi} a\theta V_{22}\left[e^{i\omega l} + e^{-i\omega l}\right] d\omega$$

$$= \begin{cases} -a\theta V_{22} & l=0 \\ 0 & l\neq 0 \end{cases}.$$

Second, when there is commitment so that $\psi = \alpha(1-k)/\gamma$, then the covariance in the very low frequency components of inflation and unemployment is positive over substantial leads and lags. Also, there unemployment may lead inflation, if only by a small amount. We establish these things by first noting that for the very lowest frequency bands,

$$s(\omega, l) \approx \left[\frac{g(e^{-i\omega})e^{i\omega l}}{(1-\lambda e^{i\omega})(1-e^{i\omega})} + \frac{g(e^{i\omega})e^{-i\omega l}}{(1-\lambda e^{-i\omega})(1-e^{-i\omega})}\right]\psi V_{11}.$$

To see this, note that $s(\omega, l)$ can be broken into three parts, corresponding to the coefficients on $V_{11}$, $V_{12}$, and $V_{22}$, respectively. For $\omega$ in the neighborhood of zero, the coefficient on $V_{22}$ is obviously bounded, since $\theta(e^{-i\omega})\,\theta(e^{i\omega})$ is bounded for all $\omega \in (-\pi, \pi)$. The same is true for the coefficient on $V_{12}$, although this requires more algebra to establish. Finally, the coefficient on $V_{11}$ is not bounded. For $\omega$ close enough to zero, this expression is arbitrarily large. For this reason, for $\omega$ close enough to zero this expression dominates the whole covariance. To establish the remainder of the second result, we now examine more closely the expression in the previous equation. Substituting out for $g$ from above:

$$\frac{s(\omega,l)}{xV_{it}} \approx \frac{\dfrac{\frac{1-\lambda^{p+1}}{1-\lambda}\left(1-e^{-i\omega}\right)+e^{-i\omega}}{\left(1-\lambda e^{-i\omega}\right)\left(1-e^{i\omega}\right)}e^{i\omega(l-p)}}{\left(1-\lambda e^{i\omega}\right)\left(1-e^{-i\omega}\right)} + \frac{\dfrac{\frac{1-\lambda^{p+1}}{1-\lambda}\left(1-e^{i\omega}\right)+e^{i\omega}}{\left(1-\lambda e^{i\omega}\right)\left(1-e^{i\omega}\right)}e^{-i\omega(l-p)}}{\left(1-\lambda e^{-i\omega}\right)\left(1-e^{-i\omega}\right)}$$

$$= \frac{\left[\frac{1-\lambda^{p+1}}{1-\lambda}\left(1-e^{-i\omega}\right)+e^{-i\omega}\right]e^{i\omega(l-p)} + \left[\frac{1-\lambda^{p+1}}{1-\lambda}\left(1-e^{i\omega}\right)+e^{i\omega}\right]e^{-i\omega(l-p)}}{2\left[1+\lambda^2-2\lambda\cos(\omega)\right]\left[1-\cos(\omega)\right]}$$

$$= \frac{\frac{1-\lambda^{p+1}}{1-\lambda}\left[\left(1-e^{-i\omega}\right)e^{i\omega(l-p)}+\left(1-e^{i\omega}\right)e^{-i\omega(l-p)}\right]+e^{-i\omega}e^{i\omega(l-p)}+e^{i\omega}e^{-i\omega(l-p)}}{2\left[1+\lambda^2-2\lambda\cos(\omega)\right]\left[1-\cos(\omega)\right]}$$

$$= \frac{\frac{1-\lambda^{p+1}}{1-\lambda}\left[\cos\left(\omega(l-p)\right)-\cos\left(\omega(l-p-1)\right)\right]+\cos\left(\omega(l-p-1)\right)}{\left[1+\lambda^2-2\lambda\cos(\omega)\right]\left[1-\cos(\omega)\right]}$$

$$= \frac{\left[1-\lambda^{p+1}\right]\cos\left(\omega(l-p)\right)+\left[\lambda^{p+1}-\lambda\right]\cos\left(\omega(l-p-1)\right)}{(1-\lambda)\left[1+\lambda^2-2\lambda\cos(\omega)\right]\left[1-\cos(\omega)\right]}.$$

Since $E\pi_t u_{t-l}$ is just the integral of $s(\omega, l)$, we can understand the former by studying the latter. Consider first the case, $\lambda = 0$, when $u_t^N$ is a pure random walk. In this case, $s(\omega, l)$, viewed as a function of $l$, is a cosine function that achieves its maximum value at $l = p$.[1] A rough estimate, based on the results in Christiano, Eichenbaum, and Evans (2001), of the time it takes for monetary policy to have its maximal impact on the price level is two years. This suggests that a value of $p$ corresponding to two years is sensible. Our notion of very low frequencies corresponds to periods of fluctuation, 20–40 years, or $10p$–$20p$. In terms of frequencies, this translates into $\omega \in [2\pi/(20p), 2\pi/(10p)]$. If we suppose the data are quarterly, then $p = 8$. For this case, we find that $s(\omega, l)$ is positive for $l \in (-10, 30)$ when $\omega = 2\pi/(10p)$ and positive for $l \in (-30, 50)$ when $\omega = 2\pi/(20p)$. We can conclude that the covariance over the very low frequencies is positive for $l \in (-10, 30)$, with unemployment leading inflation by eight periods.

When we repeated this exercise for $\lambda = 0.999$, we found that the covariance over the very low frequencies is maximized for $l$ somewhere between $l = 0$ and $l = 1$, and it is positive in the entire range, $l \in (-20, 20)$. The empirically relevant value of $\lambda$ is smaller (Ireland, 1999, reports a value in the neighborhood of 0.6), and the results we obtained for this lie in between the reports just reported for the $\lambda = 0$ and 0.999 cases. This establishes our second set of results.

[1]In general, it achieves its maximal value for any $l$ such that $\omega(l - p) = 2\pi n$, where $n$ is an arbitrary integer. So, the full set of values for which it achieves its maximum is

$$l = p + \frac{2\pi n}{\omega}, n = 0, \pm 1, \pm 2, \ldots .$$

Since at the moment we are considering small values of $\omega$, values of $l$ not associated with $n = 0$ are not of interest.

**Albanesi, Stefania, V. V. Chari, and Lawrence Christiano,** 2002, "Expectation traps and monetary policy," National Bureau of Economic Research, working paper, No. 8912, also available at http://faculty.econ.nwu.edu/faculty/christiano/research.htm.

**Barro, Robert J.,** 1987, *Macroeconomics*, second edition, Toronto: John Wiley & Sons.

**Barro, Robert J., and David B. Gordon,** 1983a, "Rules, discretion and reputation in a model of monetary policy," *Journal of Monetary Economics*, Vol. 12, pp. 101–121.

**__________,** 1983b, "A positive theory of monetary policy in a natural rate model," *Journal of Political Economy*, Vol. 91, August, pp. 589–610.

**Barsky, Robert, and Lutz Kilian,** 2000, "A monetary explanation of the great stagflation of the 1970s," University of Michigan, manuscript.

**Baxter, Marianne, and Robert G. King,** 1999, "Measuring business cycles: Approximate band-pass filters for economic time series," *Review of Economics and Statistics*, Vol. 81, No. 4, November, pp. 575–593.

**Benhabib, J., S. Schmitt-Grohé and M. Uribe,** 2001a, "Monetary policy and multiple equilibria," *American Economic Review*, Vol 91, No.1, March, pp. 167–186.

**__________,** 2001b, "The perils of Taylor rules," *Journal of Economic Theory*, Vol. 96, No. 1/2, pp. 40–69.

**Blanchard, Olivier,** 2002, *Macroeconomics*, third edition, Upper Saddle River, NJ: Prentice-Hall.

**Chari, V. V., Lawrence J. Christiano, and Martin Eichenbaum,** 1998, "Expectation traps and discretion," *Journal of Economic Theory*, Vol. 81, No. 2, pp. 462–492.

**Chari, V. V., Patrick Kehoe, and Ellen McGrattan,** 2000, "Sticky price models of the business cycle: Can the contract multiplier solve the persistence problem?," *Econometrica*, Vol. 68, No. 5, September, pp. 1151–1179.

**Christiano, Lawrence J.,** 2000, "Comment on 'Theoretical analysis regarding a zero lower bound on nominal interest rates'," *Journal of Money, Credit, and Banking*, Vol. 32, No. 4, pp. 905–930.

**Christiano, Lawrence J., and Wouter J. den Haan,** 1996, "Small-sample properties of GMM for business cycle analysis," *Journal of Business and Economic Statistics*, Vol. 14, No. 3, July, pp. 309–327.

**Christiano, Lawrence J., Martin Eichenbaum, and Charles Evans,** 2001, "Nominal rigidities and the dynamic effects of a shock to monetary policy," National Bureau of Economic Research, working paper, No. 8403.

**Christiano, Lawrence J., and Terry J. Fitzgerald,** 2003, "The band pass filter," *International Economic Review*, Vol. 44, No. 2, also available at http://faculty.econ.nwu.edu/faculty/christiano/research.htm.

**__________,** 2000, "Understanding the fiscal theory of the price level," *Economic Review*, Federal Reserve Bank of Cleveland, Vol. 36, No. 2, Quarter 2.

**__________,** 1998, "The business cycle: It's still a puzzle," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 22, No. 4, Fourth Quarter, pp. 56–83.

**Christiano, Lawrence J., and Christopher Gust,** 2000, "The expectations trap hypothesis," *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 24, No. 2, Second Quarter, pp. 21–39.

**Christiano, Lawrence J., and Massimo Rostagno,** 2001, "Money growth monitoring and the Taylor rule," National Bureau of Economic Research, working paper, also available at http://faculty.econ.nwu.edu/faculty/christiano/research.htm.

**Clarida, Richard, Jordi Gali, and Mark Gertler,** 1998, "Monetary policy rules and macroeconomic stability: Evidence and some theory," *Quarterly Journal of Economics*, Vol. 115, No. 1, February, pp. 147–180.

**Cooley, Thomas F., and Lee Ohanian,** 1991, "The cyclical behavior of prices," *Journal of Monetary Economics*, Vol. 28, No. 1, August, pp. 25–60.

**Eggertsson, Gauti B.,** 2001, "Committing to being irresponsible: Deficit spending to escape a liquidity trap," International Monetary Fund, manuscript.

**Friedman, Milton,** 1968, "The role of monetary policy," *American Economic Review*, Vol. 58, March, pp. 1–17.

**Ireland, Peter N.,** 1999, “Does the time-consistency problem explain the behavior of inflation in the United States?,” *Journal of Monetary Economics*, Vol. 44, No. 2, October, pp. 279–291.

**Kilian, Lutz,** 1998, “Small-sample confidence intervals for impulse response functions,” *Review of Economics and Statistics*, Vol. 80, No. 2, May, pp. 218–230.

**King, Robert, James Stock, and Mark Watson,** 1995, “Temporal instability of the unemployment-inflation relationship,” *Economic Perspectives*, Federal Reserve Bank of Chicago, Vol. 19, No. 3, May/June.

**King, Robert, and Mark Watson,** 1994, “The post-war U.S. Phillips curve: A revisionist econometric history,” *Carnegie-Rochester Conference Series on Public Policy*, Vol. 41, No. 0, December, pp. 152–219.

**Krugman, Paul,** 1998, “It’s baaack! Japan’s slump and the return of the liquidity trap,” *Brookings Papers on Economic Activity*.

**Krusell, Per, and Anthony Smith,** 2002, “Consumption-savings decisions with quasi-geometric discounting,” Carnegie Mellon University, Graduate School of Industrial Administration, working paper, available at http://fasttone.gsia.cmu.edu/tony/quasi.pdf.

**Kydland, Finn, and Edward C. Prescott,** 1977, “Rules rather than discretion: The inconsistency of optimal plans,” *Journal of Political Economy*, Vol. 85, No. 3, June, pp. 473–491.

**Loyo, Eduardo,** 1999, “Tight money paradox on the loose: A fiscalist hyperinflation,” Harvard University, Kennedy School of Government, unpublished manuscript.

**Lucas, Robert E., Jr., and Nancy Stokey,** 1983, “Optimal fiscal and monetary policy in an economy without capital,” *Journal of Monetary Economics*, Vol. 12, No. 1, July, pp. 55–93.

**Mankiw, N. Gregory,** 2001, “The inexorable and mysterious tradeoff between inflation and unemployment,” *Economic Journal*, Vol. 111, No. 471, pp. C45–61.

**Orphanides, Athanasios,** 1999, “The quest for prosperity without inflation,” Board of Governors of the Federal Reserve System, unpublished manuscript.

**Phelps, Edmund S.,** 1967, “Phillips curves, expectations of inflation, and optimal employment over time,” *Economica NS*, Vol. 34, No. 3, pp. 254–281.

**Sargent, Thomas J.,** 1999, *The Conquest of American Inflation*, Princeton, NJ: Princeton University Press.

**__________,** 1987, *Dynamic Macroeconomic Theory*, Cambridge, MA: Harvard University Press.

**Staiger, Douglas, James H. Stock, and Mark W. Watson,** 1997, “The NAIRU, unemployment, and monetary policy,” *Journal of Economic Perspectives*, Vol. 11, No. 1, Winter, pp. 33–49.

**Stock, James, and Mark Watson,** 1998, “Business cycle fluctuations in U.S. macroeconomic time series,” National Bureau of Economic Research, working paper, No. 6528, forthcoming in Taylor and Woodford, *Handbook of Macroeconomics*.

**Taylor, John B.,** 1993, “Discretion versus policy rules in practice,” *Carnegie-Rochester Series on Public Policy*, Vol. 39, pp. 195–214.