

GEORGE A. AKERLOF

University of California, Berkeley

WILLIAM T. DICKENS

Brookings Institution

GEORGE L. PERRY

Brookings Institution

Near-Rational Wage and Price Setting and the Long-Run Phillips Curve

OVER THIRTY YEARS ago, in his presidential address to the American Economic Association, Milton Friedman asserted that in the long run the Phillips curve was vertical at a natural rate of unemployment that could be identified by the behavior of inflation.¹ Unemployment below the natural rate would generate accelerating inflation, and unemployment above it, accelerating deflation. Five years later the New Classical economists posed a further challenge to the stabilization orthodoxy of the day. In their models with rational expectations, not only was monetary policy unable to alter the long-term level of unemployment, it could not even contribute to stabilization around the natural rate.² The New Keynesian economics has shown that, even with rational expectations, small amounts of wage and price stickiness permit a stabilizing monetary policy.³ But

The authors would like to thank Peter Kimball, Kathleen Withers, Marc-Andreas Muendler, and Megan Monroe for invaluable research assistance. We also thank Zarina Durrani for her patient instruction in how compensation professionals make decisions and how they take inflation into account. The authors wish to thank Robert Akerlof, Pierre Fortin, David Levine, Maurice Obstfeld, David Romer, Robert Solow, and participants in seminars at Williams College, Georgetown University, the University of California at Berkeley, the Levy Institute, and the Bank of Canada for valuable comments. George Akerlof is grateful to the Canadian Institute for Advanced Research, the MacArthur Foundation, and the Brookings Institution for financial support. William Dickens and George Akerlof also wish to thank the National Science Foundation for financial support under research grant number SBR 97-09250.

1. Friedman (1968); see also Phelps (1968) for an analysis very similar to that of Friedman.

2. See, for example, Lucas (1972); Sargent (1973).

3. See Akerlof and Yellen (1985); Mankiw (1985).

the idea of a natural unemployment rate that is invariant to inflation still characterizes macroeconomic modeling and informs policymaking.

The familiar empirical counterpart to the theoretical natural rate is the nonaccelerating-inflation rate of unemployment, or NAIRU. Phillips curves embodying a NAIRU are estimated using lagged inflation as a proxy for inflationary expectations. NAIRU models appear in most textbooks, and estimates of the NAIRU—which is assumed to be relatively constant—are widely used by economic forecasters, policy analysts, and policymakers. However, the inadequacy of such models has been demonstrated forcefully in recent years, as low and stable rates of inflation have coexisted with a wide range of unemployment rates. If there were a single, relatively constant natural rate, we should have seen inflation slowing significantly when unemployment was above that rate, and rising when it was below. Instead, the inflation rate has remained fairly steady, with annual inflation as measured by the urban consumer price index (CPI-U) ranging from 1.6 to 3.0 percent since 1992, while the unemployment rate has ranged from 6.8 to 3.9 percent. In this paper we present a model that can accommodate relatively constant inflation over a wide range of unemployment rates.

Another motivation is a recent finding by William Brainard and George Perry.⁴ Estimating a Phillips curve in which all the parameters are allowed to vary over time, they find that the coefficient on the proxy for expected inflation in the Phillips curve has changed considerably, while other parameters of that model have been relatively constant. In particular, Brainard and Perry found that the coefficient on expected inflation was initially low in the 1950s and 1960s, grew in the 1970s, and has fallen since then. The model we present can explain both why the coefficient on expected inflation might be expected to change over time and, to some extent, the time pattern of changes observed by Brainard and Perry.

Our paper also allows an interpretation of the findings of Robert King and Mark Watson and of Ray Fair.⁵ Both find a long-run trade-off between inflation and unemployment. In addition, King and Watson find that the amount of inflation that must be tolerated to obtain a given reduction of unemployment rose considerably after 1970. Our model allows a trade-off, but only at low rates of inflation such as those that prevailed in the 1950s,

4. Brainard and Perry (2000).

5. King and Watson (1994); Fair (2000).

1960s, and 1990s. At higher rates of inflation, the trade-off is reduced, and at high enough rates of inflation, it disappears.

Much of the empirical controversy surrounding the relationship between inflation and unemployment has focused on how people form expectations. This may be neither the most important theoretical nor the most important empirical issue. Instead, this paper suggests that it is not how people form expectations but how they use them—and even whether they use them at all—that is the issue. Economists typically assume that economic agents make the best possible use of the information available to them. But psychologists who study how people make decisions have a different view. They see individuals as acting like intuitive scientists, who base their decisions on simplified, abstract models.⁶ However, these simple intuitive models can be misleading; indeed, sometimes they are incorrect. Psychologists have studied the use of these simplified abstractions, often called mental frames or decision heuristics, and the mistakes that result from them. Economists should not assume absence of cognitive error in economic decisions, nor should they assume that their own models and those of the public exactly coincide.

We propose that there are three important ways in which the treatment of inflation by real-world economic agents diverges from the treatment assumed in economic models. First, when inflation is low, a significant number of people may ignore inflation when setting wages and prices. Second, even when they take inflation into account, they may not treat it as economists assume. In particular, we hypothesize that the informal use of inflationary expectations in wage and price decisions leads to less than complete projection of anticipated inflation, with consequences for the aggregate relation between inflation and unemployment. Finally, we believe that workers have a different view of inflation from that of trained economists. Workers see inflation as increasing prices and reducing their real earnings; they do not fully, if at all, appreciate that inflation increases the nominal demand for their services. Thus they tend to view the nominal wage increases they receive at low rates of inflation as a sign that their work is appreciated, and to be happier in their jobs as a result. They may also be unaware of the extent to which inflation is increasing the pay available to them in alternative jobs. Even fully rational employers, who must solve the typical efficiency wage problem, can exploit workers' misper-

6. See Nisbett and Ross (1980).

ceptions by paying wages that are less than what would be required if workers fully incorporated inflation into their mental frames.

If *any* of these three departures from the fully rational use of information about inflation are important, then at low rates of inflation, prices and wages will be set consistently lower relative to nominal aggregate demand than they would be at zero inflation. As a result, operating the macroeconomy with a low but positive rate of inflation will permit a higher level of employment to be sustained. We will show that at low rates of inflation the behaviors that we posit, which depart from the fully rational decisions of typical economic models, impose very small costs on those who practice them. Because there may be subjective or objective costs associated with fully rational behavior, or because implementing fully rational behavior may require overcoming some perception threshold or behavioral inertia, it is plausible that the costs of nonrational behavior may be too small to induce rational behavior from all economic agents. However, if inflation increases, the costs of being less than perfectly rational about it will also rise, and people will switch their behavior to take inflation into full account. Thus, although increasing inflation modestly above zero will permit lower unemployment, there is a rate of inflation above which the sustainable unemployment rate rises as more and more people adopt fully rational behavior. This rate of inflation thus minimizes the sustainable rate of unemployment.

The remainder of the paper proceeds in three steps. First, we describe departures from perfect rationality at low rates of inflation and present some evidence that supports our view. Second, we formally derive our model of near-rational wage and price setting, show that the costs of near rationality are small, derive short- and long-run Phillips curves from the model, and present a calibration exercise that shows that, even when only a fraction of wages and prices are influenced by near-rational behavior, there can still be substantial long-run gains in employment from moderate, rather than very low or zero, inflation. Finally, we estimate the theoretical model using postwar quarterly U.S. data. The results support the theoretical model and are surprisingly robust.

Near-Rational Behavior Toward Inflation

As noted above, psychologists who study decisionmaking approach it differently from the way economists do. Psychologists have identified

many ways in which real-world decisionmaking departs from economic rationality. Here we describe three ways in which we suspect that behavior toward inflation departs from the economist's rational model.

First, psychologists suggest that decisionmakers, far from making the best use of available information, readily ignore potentially relevant considerations and discard potentially relevant information in order to simplify their decision problems. Daniel Kahneman and Amos Tversky have dubbed this behavior editing.⁷ When people "edit" decision problems, they rule out less important considerations in order to concentrate on the few factors that matter most. In this regard, real-world decisionmakers are no different from academic economists when they construct models: unimportant factors are ignored in order to concentrate on important factors. A related literature in the psychology of perception suggests that items must reach a threshold of salience before they are even perceived.⁸ Thus, when inflation is low, it may be at most a marginal factor in wage and price decisions, and decisionmakers may ignore it entirely.⁹

We know of no strong evidence either for or against the view that some wage and price setters ignore inflation,¹⁰ but several before us have suggested the occurrence of such behavior. For example, Otto Eckstein and

7. Kahneman and Tversky (1979). Kunreuther (1978) has used the phenomenon of editing to explain why many people do not buy disaster insurance: very low probability events are ignored in decisionmaking. His book (pp. 165–86) presents the results of experiments that demonstrate the phenomenon of editing.

8. See Gleitman (1996).

9. The behavior of cost-of-living-adjustment (COLA) clauses is consistent with increasing attention being paid to inflation at higher levels of inflation. As inflation rose in the 1970s and 1980s, coverage of union workers by COLAs in the United States increased. In the late 1960s about one-quarter of workers involved in collective bargaining were covered by COLA clauses, compared with about 60 percent for the inflationary decade from 1975 to 1985 (Hendricks and Kahn, 1985, pp. 36–37). As inflation fell in the late 1980s, the fraction covered fell to 40 percent in 1990 (Holland, 1995, p. 176). Such inflation sensitivity of COLAs is consistent with our basic idea that wage and price setters tend to ignore inflation in their wage and price setting when inflation is low, but tend to take it into account as inflation rises. But this evidence has at least two other possible explanations. It is well known (see Ball, Mankiw, and Romer, 1988, p. 56) that the variance of inflation increases with its level. COLAs may increase at higher levels of inflation as insurance against this variance. Furthermore, if at higher rates of inflation a greater fraction of inflation is due to monetary rather than to real shocks, more contracts will be indexed at higher than at lower rates of inflation (see Gray, 1978).

10. Direct attempts to assess the effects of forecast inflation on wage setting have ignored the indirect effects of inflation through other information that will be correlated with inflation. Such information includes the wages and prices of competitive and complementary

Roger Brinner based their model of a shifting Phillips curve on the assumption that inflationary expectations mattered more in determining inflation in the 1970s than in the 1960s.¹¹ One major macroeconomics textbook describes the postwar U.S. Phillips curve in terms of an early period of low inflation, which was ignored by wage and price setters, and a later period of high inflation, when the coefficient on the last period's inflation was close to one.¹² Two of the officials who over the past five years have been most responsible for achieving the Federal Reserve's goal of price stability have also suggested the possibility of inflation editing. Former Fed Vice Chairman Alan Blinder, in company with coauthors, has theorized:

A businessman who cannot keep infinite amounts of information in his head may worry about a few important things and ignore the rest. And when nationwide inflation is low, it may be a good candidate for being ignored. Indeed, one prominent definition of "price stability" is inflation so low that it ceases to be a factor in influencing decisions.¹³

Senate testimony of Federal Reserve Chairman Alan Greenspan seems also to suggest the view that, at low rates of inflation, economic agents may simply ignore it:

goods and factors. Thus the findings that wage and price setters seem to put little weight on inflation (Blinder and others, 1998; Levine, 1993) are inconclusive. For this reason we made our own attempt to solicit such information. We sent an e-mail questionnaire to randomly selected members of the American Compensation Association asking them to recommend wage and salary increases in hypothetical situations varying by respondent in a number of different dimensions. The respondents were given the type of information that personnel executives typically use to make recommendations for wage and salary changes. This information included the wage and salary increases of other firms in their labor market over the past year, the desired relative wage and salary position of their firm, expected wage and salary increases of other firms in their labor market for the next year, the increase in the CPI, the difficulty of hiring and retention, and their firm's expected net revenue growth relative to that of its industry and relative to that of the economy as a whole. The mean of expected wage increases by other firms in the sample was increased one-for-one with the rate of inflation. The total effect of changes in inflation on wage and salary increases by individual firms can be seen by regressing the recommended wage and salary increases on the expected wage and salary increases of others and the CPI. The point estimate of the change caused by a one-point change in the CPI in the wages of an individual firm, given that that firm's changes are representative of other firms facing the same increase in the CPI, is 0.738. This estimate is obtained by dividing the coefficient on the CPI by one minus the coefficient on the expected wage increases of other firms. Unfortunately, this estimate has a very high standard error, and we cannot rule out the possibility that the impact of an increase in expected CPI inflation on wage inflation would be one for one, but the point estimate is suggestive of our view.

11. Eckstein and Brinner (1972).

12. Blanchard (1999, pp. 153–54).

13. Blinder and others (1998, p. 98).

By price stability I mean a situation in which households and businesses in making their savings and investment decisions can safely ignore the possibility of sustained, generalized price increases or decreases.¹⁴

Second, even when people do pay attention to inflation, they may not use expectations in the way economists typically assume. If economic agents used a formal procedure to make wage and price decisions, they would first use available information to determine a desired real wage or price change. They would then add in the amount of inflation they expect between the time they are making the decision and some future time during the period over which they expect the price or wage to be in effect. But if they make the decisions intuitively, subjectively considering a number of factors simultaneously, including inflation, there is no reason to expect that the decision will give the appropriate weight to inflation. One decision heuristic, suggested to us by interviews with compensation professionals, is that information on inflation may simply be averaged along with other factors to arrive at a nominal wage or price increase. This would mean that an increase in inflation would lead to the setting of a higher wage or price, but the effect would be less than one for one. Thus, less than complete weighting of inflation is the second departure from full rationality that may influence the relationship between inflation and unemployment.

In fact, textbooks for compensation professionals warn against using the formal procedure that economists would imagine to be standard. For example, George Milkovich and Jerry Newman discourage their readers from granting automatic wage and salary increases, including increases for the cost of living.¹⁵ Such automatic grants, these authors say, reduce the funds available to reward employees for performance. Similar thoughts are expressed in the handbook of the influential Hay Group of compensation consultants, in which managers are advised to “avoid linking salary movement to changes in the cost of living, because this creates entitlement and reduces the amount of money available to differentiate for performance.”¹⁶

The third important departure from the hyperrational model comes from the way workers perceive inflation. Robert Shiller has documented very large differences between the intuitive models of inflation used by the lay public, most of whom are wage and salary recipients, and the mental

14. Greenspan (1988, p. 611).

15. Milkovich and Newman (1984).

16. Rock and Berger (1991, p. 556).

accounting of economists who study the effects of inflation scientifically.¹⁷ Wage and salary earners systematically underestimate the effects of inflation on the wages that their employers will want to pay them, even in questionnaires where the effects of inflation are quite explicit, so that it is highly unlikely that inflation is ignored. As a consequence, and especially at moderate rates of inflation when real wages are not perceptibly eroded, workers' job satisfaction may be enhanced by nominal wage increases even if they fail to fully reflect inflation.

There is considerable evidence for this sort of reaction on the part of workers. Economists see inflation as induced by changes in the money supply and thus as having a uniform effect on nominal wages and other prices, so that inflation causes no change in real income. In his questionnaire study, Shiller has shown that, in contrast, the public has no such expectations. For example, when asked "to imagine how things would be different if the United States had experienced higher inflation over the last five years,"¹⁸ only 31 percent of his noneconomist subjects believed that their own nominal income would have been higher than in the absence of inflation. When asked "to evaluate [a variety] of theories about [how] the effects of general inflation on wages and salary relates to your own experience and your own job," 60 percent of economists, but only 11 percent of the general public, elected that "competition among employers will cause my pay to be bid up. I could get outside offers from other employers, and so, to keep me my employer will have to raise my pay too." A popular answer for the general public (26 percent), in contrast to economists (4 percent), was the following: "the price increase will create extra profits for my employer who can now sell output for more; there will be no effect on my pay."¹⁹

These responses suggest that the public fails to understand inflation as a general-equilibrium phenomenon. They believe that inflation will make them poorer because it bids up the prices of the goods they consume, but they fail to appreciate fully, if at all, that inflation will also bid up the prices of other competing factors and other competing workers, thereby resulting in a rise in their own wages and salaries. Thus, according to Shiller, the "biggest gripe about inflation," expressed by 77 percent of the general

17. Shiller (1997).

18. Shiller (1997, p. 21).

19. Shiller (1997, pp. 31–32).

public (but only 12 percent of economists), was that inflation “hurts my real buying power. It makes me poorer.”²⁰

Economists should not be surprised that laypeople underestimate the effect of inflation on the demand for their own services. One of the most significant differences between trained economists and the lay public is economists’ greater appreciation of general equilibrium. The cognitive difficulty of general equilibrium has been indicated by the fact, noted by the Commission on Graduate Education, that even economics graduate students do not give the correct explanation for why barbers’ wages, in the technically stagnant hair-cutting industry, have risen over the past century.²¹ If economics graduate students fail to appreciate the effects on barbers’ opportunity costs from wage increases due to productivity change outside the hair-cutting industry, it would be a stretch to expect the lay public to see that, as inflation rises, the demand for their services (in nominal dollars) will similarly rise with it.

Findings by Eldar Shafir, Peter Diamond, and Tversky are consistent with those of Shiller. In one vignette, which they related to respondents, these authors draw a contrast between Ann, who receives a 2 percent nominal salary increase at zero inflation, and Barbara, who gets a 5 percent nominal salary increase at 4 percent inflation. Most respondents correctly answered that Ann would be better off economically,²² but they also said that Barbara would be happier and less likely to leave her job. This reaction to the vignette suggests that respondents have not *ignored* the inflation, as they would with editing—otherwise Ann would not be judged better off economically. But the other answer, favoring Barbara, suggests that they may also underestimate the effect that inflation will have on Barbara’s other alternatives, thus leading them to conclude that she will be happier and less likely to quit her job.

Unfortunately, these authors did not probe the reasons why their respondents believed Barbara should be happier than Ann, but they are responding as if the inflation has not increased Barbara’s alternatives by an equal amount. If the wages that she could get on the outside, as well as all of the prices that she would be paying, have increased by 4 percent, Barbara should be less happy than Ann and more likely to leave her job. Our model of inflation, however, suggests a good reason why Barbara should feel hap-

20. Shiller (1997, p. 29).

21. Krueger and others (1991, p. 1044).

22. Shafir, Diamond, and Tversky (1997).

pier than Ann does and be less likely to quit her job: she does not feel that her alternatives improve at the rate of inflation. Another question by Shiller suggests that the responses obtained to this vignette reflect the true opinion of the American public. He finds that about half of the U.S. general public—but only 8 percent of economists—think that they would feel more job satisfaction “if their pay went up . . . even if prices went up as much.”²³

Neither the vignette by Shafir and others nor Shiller’s question deals with the possibility, perhaps on the mind of the public, that the inflation is caused by a supply shock that decreases the real demand for workers rather than a money-neutral demand shock that leaves all demands unchanged in real terms. Of course, if that is really what is on the mind of the public, even when there is a persistent demand-induced increase in the rate of inflation, workers will still have greater job satisfaction with some small amount of inflation than with no inflation. This, then, is the third way in which we think that near rationality may impact the relation between inflation and unemployment. If higher job satisfaction at low rates of inflation leads to higher morale, less shirking, higher productivity, and less turnover, then firms face a different efficiency wage constraint at low rates of inflation than they face at either zero inflation or at high rates of inflation, when workers’ attitudes toward inflation may become more rational.

A Simple Model of Near-Rational Wage and Price Setting

We now present a simple formal model of the economy that incorporates the behavioral insights just described. In the model, some firms’ wage and price setters may ignore inflation, or firms may be aware of inflation but use it as only one of several factors in setting wages and prices, thus underweighting it relative to the behavior assumed in hyperrational models. And workers themselves may ignore or underweight inflation when considering their satisfaction at their current job, which in turn would affect their productivity. The net effect on unit labor costs of this behavior by workers may or may not be fully factored into firms’ wage setting. The implications of our model for the behavior of the macroeconomy

23. Shiller (1997, p. 37).

are not affected by this aspect of firms' behavior. However, we formally consider the case where firms do not correctly anticipate the effects on worker satisfaction and productivity, because this case permits a simple derivation of the profit shortfall a firm experiences from less than fully rational behavior.

The easiest place to begin the model is with its macroeconomic behavior. Income is determined by the quantity theory equation

$$(1) \quad \bar{p}Y = M,$$

where Y is real income, \bar{p} is the average price level in the economy, and M is the money supply. The usual constant of such quantity theory equations has been normalized to one by the choice of units.

The microeconomics of this economy begins with the boilerplate for models with monopolistically competitive firms. There are n firms in this economy. They divide up total aggregate demand, M/\bar{p} , according to the relative prices for their respective goods, so that the demand for the output of an individual firm is of the form

$$(2) \quad \frac{1}{n} \frac{M}{\bar{p}} \left(\frac{p}{\bar{p}} \right)^{-\beta},$$

where p is the price charged by a firm for its own product.

This takes us to the first innovation of the model, which occurs in the formulation of productivity and its effect on wages. All of these firms will pay an efficiency wage, which minimizes the unit labor cost of production. Productivity (as well as turnover costs) in each firm depends upon the morale of its workers. That morale, in turn, depends upon workers' conception of their outside opportunities, which has two major determinants. The first is the rate of unemployment, which determines how easy it would be for an individual worker to obtain another job. The higher the unemployment rate, the lower will be the opportunity cost of workers, and therefore the higher the morale inside the firm. The second determinant of morale is the workers' perception of the gap between their wage at their own firm and the wage outside the firm. That perception depends upon the wage being paid by the worker's current firm and his or her reference wage, which gives that worker's perception of the wages of other workers. Thus the productivity of the firm will depend also upon both the wage

it pays and the level of unemployment. For convenience we give productivity the following functional form:

$$(3) \quad P = -A + B \left(\frac{w}{w^R} \right)^\alpha + Cu,$$

where P denotes labor productivity, w is the wage paid by the firm, w^R is the reference wage of its workers, and u is the aggregate unemployment rate. α is chosen in the range $0 < \alpha < 1$.

Firms set both prices and wages one period ahead. In so doing they project the effects of inflation on the reference wages of their workers. These reference wages, of course, determine the level of wages that a firm should be paying. Totally rational firms will incorporate all of their expected inflation into the reference wage w^R . In contrast, near-rational firms—and similarly, fully rational firms whose workers underweight inflation in w^R —will incorporate only a fraction of inflation, a , into their projections of inflation. When a is zero, inflation is totally ignored. In the intermediate range, $0 < a < 1$, it is merely underestimated. Thus the reference wage for fully rational workers for the joint wage and price decisions of fully rational firms is

$$(4) \quad w_r^R = \bar{w}_{-1}(1 + \pi^e),$$

where \bar{w}_{-1} is the average wage paid to all workers in the previous period, and π^e is the expected rate of price inflation. The reference wage for the wage- and price-setting decision by near-rational firms, which are engaging in cognitive error, will analogously be

$$(5) \quad w_{nr}^R = \bar{w}_{-1}(1 + a\pi^e).$$

Equation 5 also describes the reference wage for the near-rational employees.

The profit-maximizing choice of the price for both the rational and the near-rational firms will take the following form. In both cases the prices will be a markup over expected unit labor costs,

$$(6) \quad p_j = m \frac{w_j}{P_j^e},$$

where j refers to both rational and near-rational firms, and P_j^e is the firms' expected productivity. The markup factor m will be $\beta/(\beta - 1)$. The rational and the near-rational firms have the same expectations about productivity.

These maximizing firms will, in turn, establish their wages as a multiple of their respective reference wages, which will differ for rational and for near-rational firms. The efficiency wage paid by each firm type will minimize its respective unit labor costs, w_j/P_j . Accordingly, each type of firm will choose, respectively,

$$(7) \quad w_j = \left[\frac{A - Cu}{B(1 - \alpha)} \right]^{1/\alpha} w_j^R.$$

Near-rational firms set wages that are different from those of fully rational firms, but the difference does not cumulate. The wages of near-rational firms are reset relative to their reference wage in each and every period. The reference wages for rational and near-rational firms are both multiples of the last period's average wage, and therefore both rise with steady inflation and always stand in the ratio $(1 + a\pi^e)/(1 + \pi^e)$. As a result, the difference between wages at the two types of firms will be small at low and moderate levels of inflation.

The profits of each type of firm will be revenues net of labor costs. Given the demand function for firms' product in equation 2 and their labor productivity in equation 3, the profits for the two types of firms will be

$$(8) \quad \frac{1}{n} \frac{M}{\bar{p}} \left[p_j \left(\frac{p_j}{\bar{p}} \right)^{-\beta} - \left(\frac{p_j}{\bar{p}} \right)^{-\beta} \frac{w_j}{P_j} \right].$$

So far the model has described the case where the firm ignores or underweights inflation, and the case where the firm is rational but workers' reference wages are underindexed. Both situations will give us similar Phillips curves. In one case, near-rational firms will be switching to true rationality as their costs from near rationality mount with high inflation, whereas in the other case workers will eventually curb their misperceptions as inflation rises. But the two hypotheses are slightly different, and at this point we shall take the alternative route that analyzes the model where the near-rational firms fail to fully take account of inflation in forming w^R , but workers are always fully cognizant of inflation in determining their reference wage. This route permits an evaluation of the worst possible losses by near-rational firms from their failure to correctly perceive the effects of inflation.

Each of the terms p_j , w_j , and P_j is known relative to the value of the average wage \bar{w}_{-1} from equations 3 through 7, so it is possible to evaluate the relative profits of rational and near-rational firms. Using the profit function in equation 8, along with the assumption that both rational and near-rational firms have correct expectations about inflation, yields a formula for the relative profits of the two types of firms.²⁴ The relative increase in profits that a near-rational firm could make by becoming a rational firm is given by the following loss function,

$$(9) \quad L = \frac{r_r - r_{nr}}{r_r} = 1 - z^{1-\beta} \left[\beta - (\beta - 1) \frac{\alpha}{z^\alpha - 1 + \alpha} \right],$$

where r_r and r_{nr} are, respectively, the profits of rational and near-rational firms and z is the ratio $(1 + a\pi)/(1 + \pi)$. Equation 9 has three implications for this paper, which we shall explore in turn.

The first implication is that those who fail to maximize profits either by ignoring inflation ($a = 0$) or by taking it into account only partially ($0 < a < 1$) are near rational. When π is zero, the losses of such producers are zero, as can be seen by the fact that when π is zero, z is one. Thus, according to equation 9, the losses from being near rational when z is zero will also be zero. These losses will also continue to be small at levels of inflation near zero, since the derivative of equation 9 with respect to π is also zero when π is zero.

Second, equation 9 serves as the springboard for the completion of the model we will estimate below, which is based explicitly on the losses that are entailed from near-rational behavior. To complete the model, it is assumed that firm wage and price setters are willing to tolerate losses relative to their profits only up to a given threshold, ϵ , before they will switch to fully rational behavior. We assume that these thresholds are normally distributed with mean μ_ϵ and standard deviation σ_ϵ . The fraction of near-rational price setters, accordingly, will be

$$(10) \quad 1 - \Phi \left\{ \frac{1 - z^{1-\beta} \left[\beta - (\beta - 1) \frac{\alpha}{z^\alpha - 1 + \alpha} \right] - \mu_\epsilon}{\sigma_\epsilon} \right\},$$

24. A slightly more complicated formula will give the relative profits when π is different from π^e .

where Φ is the standard cumulative normal distribution, and μ_ε and σ_ε are, respectively, the mean and the standard deviation of the distribution of the thresholds ε .

Finally, equation 9 also yields benchmark estimates of the losses due to near-rational behavior. Table 1 shows the fraction of the profits of the fully rational firm sacrificed by the near-rational firm at different rates of inflation for two different values of a and two different values of both α and β .

To put the values in table 1 in perspective, consider the findings of Jonathan Leonard and of Steven Davis, John Haltiwanger, and Scott Schuh that the typical firm annually experiences shocks to demand that cause it to adjust its size up or down by roughly 10 percent.²⁵ Failing to adjust capacity to accommodate such a shock would cost a firm 10 percent of its profits. Thus it does not seem hard to believe that, for the typical firm, the issue of how to treat inflation in setting prices is far down the list of items demanding managerial attention, at least as long as inflation is under 5 percent.

Implications for the Long-Run and the Short-Run Phillips Curve

The model also allows easy derivation of both a short-run Phillips curve with given expectations of price inflation, and a long-run Phillips curve where actual and expected inflation must coincide. The short-run wage Phillips curve is obtained from wage-setting behavior and the equation for the average wage. The average wage in this economy will be

$$(11) \quad \bar{w} = \Phi w_r + (1 - \Phi) w_{nr}.$$

Using the wage-setting behavior of the rational and near-rational firms,

$$(12) \quad \bar{w} = \Phi \left[\frac{A - Cu}{B(1 - \alpha)} \right]^{1/\alpha} w_r^R + (1 - \Phi) \left[\frac{A - Cu}{B(1 - \alpha)} \right]^{1/\alpha} w_{nr}^R,$$

which can be rewritten as

$$(13) \quad \bar{w} = \Phi \left[\frac{A - Cu}{B(1 - \alpha)} \right]^{1/\alpha} w_{-1} (1 + \pi^e) + (1 - \Phi) \left[\frac{A - Cu}{B(1 - \alpha)} \right]^{1/\alpha} w_{-1} (1 + a\pi^e),$$

25. Leonard (1987); Davis, Haltiwanger, and Schuh (1996).

Table 1. Share of Profits of a Fully Rational Firm Lost by Near-Rational Behavior in the Treatment of Inflation

Percent

Inflation (percent)	Near-rational firms ignore inflation ($a = 0$)				Near-rational firms underweight inflation ($a = 0.7$)			
	Elasticity ^a $\beta = 3$		Elasticity $\beta = 5$		Elasticity $\beta = 3$		Elasticity $\beta = 5$	
	$\alpha = 0.1$	$\alpha = 0.75$	$\alpha = 0.1$	$\alpha = 0.75$	$\alpha = 0.1$	$\alpha = 0.75$	$\alpha = 0.1$	$\alpha = 0.75$
1	0.04	0.03	0.12	0.11	0.003	0.003	0.01	0.009
2	0.16	0.13	0.50	0.44	0.01	0.01	0.04	0.04
3	0.36	0.30	1.14	1.01	0.03	0.03	0.09	0.08
4	0.65	0.53	2.07	1.83	0.05	0.04	0.17	0.15
5	1.03	0.85	3.32	2.91	0.08	0.07	0.26	0.23

Source: Authors' calculations.

a. Elasticity of demand. The parameter α represents the curvature of the unit cost function.

using the definition of the reference wage. Dividing both sides by \bar{w}_{-1} and collecting terms yields the relation

$$(14) \quad (1 + \pi_w) = \left[\frac{A - Cu}{B(1 - \alpha)} \right]^{1/\alpha} [1 + \Phi\pi^e + (1 - \Phi)a\pi^e],$$

where π_w is the rate of wage inflation. Taking the logarithms of both sides of equation 14, and approximating $\ln(1 + \pi_w)$ by π_w , $\ln[1 + \Phi\pi^e + (1 - \Phi)a\pi^e]$ by $[\Phi + (1 - \Phi)a]\pi^e$, and $\ln[A - Cu]/[B(1 - \alpha)]^{1/\alpha}$ by its linear approximation, $d - eu$, yields the short-run wage Phillips curve

$$(15) \quad \pi_w = d - eu + (1 - f)\pi^e,$$

where $f = (1 - a)(1 - \Phi)$.

A price Phillips curve similar to equation 15 can also be derived from the model. The slight difference between the price Phillips curve implied by our model and the wage Phillips curve in equation 15 is the inclusion of a change-in-unemployment term in the price Phillips curve. This term enters because changes in the unemployment rate will cause changes in productivity and hence, through equation 6, in the price-wage markups.²⁶ We take this into account when we estimate the model by allowing lags on the unem-

26. The price Phillips curve will be of the form

$$\pi = c - eu + f\pi^e + (1 - f)h\Delta u^e,$$

where $h = -C/[B(1 - \alpha)]$, u is current unemployment, and Δu^e is the expected change in unemployment.

ployment rate. The steady-state Phillips curves with constant unemployment will be unaffected by varying markups caused by varying unemployment.

The short-run Phillips curve in equation 15 should come as no surprise. If all inflation had been included in the mental frames of the firms, which are the wage and price setters in this model, the coefficient f would be equal to zero. The near-rational firms, which constitute a fraction $(1 - \Phi)$ of all firms, ignore a fraction $(1 - a)$ of inflation. As a consequence, the Phillips curve in equation 15 mimics the usual inflation-augmented Phillips curve, but with a fraction $(1 - a)(1 - \Phi)$ of the expected inflation ignored. Thus the Phillips curve in equation 15 is not just an artifact of our illustrative model of price and wage setting. As long as a fraction of inflation is ignored or underweighted in near-rational wage and price setting, that fraction of inflation should fail to enter the inflation augmentation term. A whole spectrum of other models in which various combinations of firms and workers ignore or underweight inflation in their mental frames will yield similar results.

Using equation 15, the long-run Phillips curve, where actual and expected inflation are equal, will be

$$(16) \quad u^n - u = \frac{1}{e} f \pi.$$

where u^n is the natural rate of unemployment if all firms are rational. Its value in this model is d/e .

The long-run Phillips curve in equation 16 will be bowed inward and then forward bending. At zero inflation, π is zero, and therefore unemployment is at the natural rate. At very high inflation all firms will have given up being near rational. The losses from near-rational behavior will be sufficiently large that, by equation 10, Φ will be close to one—so that f , which is $(1 - \Phi)(1 - a)$, will be close to zero. Thus, at both very high and very low inflation, unemployment will be close to the natural rate, which is the level of unemployment that would occur if all firms were totally rational. At inflation above zero, unemployment will always be below the natural rate, since f will always be positive; however, at high rates of inflation the natural rate is an asymptote.

Figure 1 portrays the rate of unemployment that corresponds to different levels of inflation in the long run with benchmark parameters. We have assumed that near-rational firms completely ignore inflation ($a = 0$). We chose the parameters describing the distribution of Φ so that at least half of

Figure 1. A Hypothetical Long-Run Phillips Curve

Inflation (percent per year)



Source: Authors' calculations from calibration of the theoretical model.

all firms are always fully rational (thus μ_e is zero), and 95 percent of all firms are rational by the time inflation is at 5 percent a year. We also chose α at 0.75 and an elasticity of demand (β) of 3 (which implies a value for σ_e of 0.005, or 0.15 percent of normal profits), although, as we will discuss below, these assumptions hardly matter for the shape of figure 1.

The level of inflation that maximizes the product of f and π according to equation 16 will minimize unemployment. For the parameter values chosen to create figure 1, that inflation rate is 2.6 percent. At that rate of inflation the long-run equilibrium rate of unemployment is 1.7 percentage points lower than at either a rate of inflation of zero or a rate above 6 percent.²⁷

27. Interestingly, our choices of the values of the elasticity of demand (β) and the curvature of the productivity function (α) hardly matter for the shape of the curve in figure 1 or for the lowest sustainable rate of unemployment and its accompanying rate of inflation. Once we set the fraction of firms that are near rational at two points, we have described the curve for a given value of a . This result reflects a finding that will surface again later when we estimate the model, which is discussed in more detail in the next section. That is that the loss function is very nearly approximated by a constant times the square of inflation, so that the argument of the cumulative normal in our model can be very well approximated with two parameters.

Why does employment rise with inflation at low rates of inflation? In our model, inflation is not underestimated, but instead it is underweighted in the reference wage used for wage setting. This has the same consequences as underestimation. Near-rational firms either ignore or fail to fully project inflation, so they set lower wages, and therefore also set lower prices, relative to nominal demand, than they would if they were fully rational. Since the wages of fully rational firms are affected by the wages paid by near-rational firms, both types of firms pay lower wages than they would if all firms were fully rational. At these lower wages, employment will be higher. These higher levels would also occur in the slightly different version of the model in which workers underestimate the impact of inflation.²⁸

Empirical Evidence for Near-Rational Wage and Price Setting

In this section we discuss three related types of evidence for the importance of the type of behavior we describe. We begin by recounting the findings of Brainard and Perry's recent analysis of a Phillips curve model with time-varying parameters. We then do a simple exercise estimating Phillips curves on a split sample to see how the estimated coefficient of inflation differs between periods of high and low inflation. Finally, we estimate the model described in the previous section and present estimates of the unemployment-minimizing rate of inflation and the employment gains from being there rather than at the natural rate.

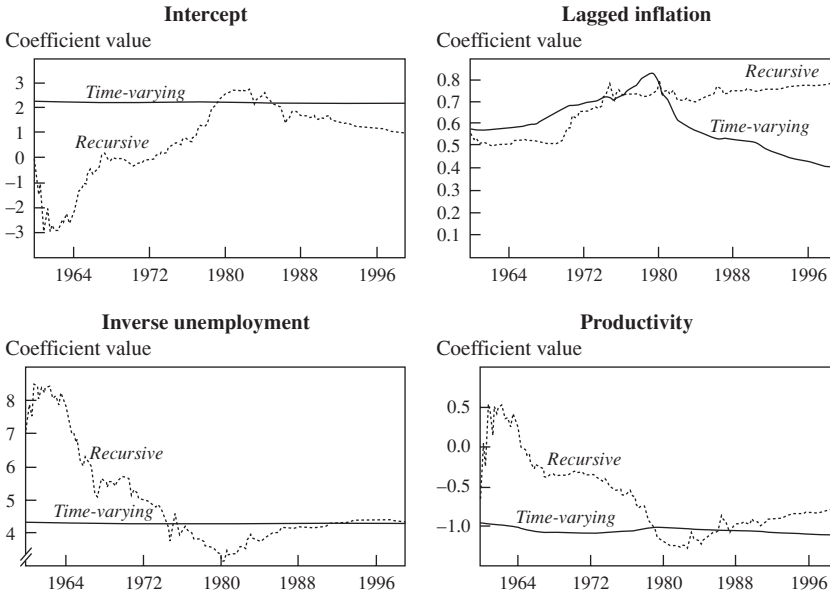
Time-Varying Parameters

In the Brainard and Perry paper that we described at the outset,²⁹ the authors were addressing how uncertainty affects policymaking. Their empirical work demonstrating one key source of uncertainty reveals precisely the departures from conventional NAIRU models that our model predicts. Previous work examining how the NAIRU varied over time assumed the NAIRU framework and allowed time variation only in the

28. We have resisted the temptation to call the unemployment-minimizing rate of inflation the optimal rate. In this model productivity also varies with the rate of inflation. Therefore, at the minimum unemployment rate, output is not at its maximum.

29. Brainard and Perry (2000).

Figure 2. Recursive Least-Squares Estimates and Time-Varying Kalman Filter Estimates of Price Phillips Curve Parameters, 1960–98



Source: Brainard and Perry (2000).

intercept of the equation.³⁰ Brainard and Perry applied a general Kalman filter estimation that permits all the key Phillips curve parameters to vary—lagged inflation and unemployment as well as the intercept—and lets the data choose the allocation of time variation among them. Figure 2 summarizes their results with CPI inflation as the dependent variable. The figure shows substantial time variation in the coefficient of the lagged inflation term, and virtual stability in the intercept and the coefficient of the inverse unemployment rate, which they measure by the unemployment rate of 25- to 54-year-old men to account for demographic changes over time. The coefficient on lagged inflation is low during periods of low inflation and approaches one only in the inflationary middle years of the sample period.

The virtual stability over time in the unemployment coefficient and the intercept in the Brainard-Perry time-varying estimates is also worth noting. Rather than attributing the episodes of sustained low unemployment-

30. For typical applications see Stock and Watson (1998) and Gordon (1998).

ment to declines in a NAIRU that is invariant to inflation, these results suggest a change in price- and wage-setting behavior that accompanied periods of low inflation. The model described above predicts coefficients on lagged inflation that change with the inflation regime while coefficients elsewhere are constant.

Brainard and Perry compare their Kalman filter estimates with recursive least-squares estimates, which are also shown in figure 2. These comparisons suggest why conventional estimation has seemed to support the NAIRU model since Franco Modigliani and Lucas Papademos introduced it in the inflationary mid-1970s.³¹ Before that time, lagged inflation in Phillips curves had been consistently estimated to have a coefficient well below one. But the large increase in inflation in the mid-1970s corresponded to a period of large variance in inflation, and fixed-coefficient estimation has been dominated by that episode ever since. If the coefficients in fact have varied over time, any procedure that assumes that they are fixed will yield misleading results. This includes the recursive estimates, which treat them as fixed in each interval over which they are estimated.

Periods of Low and High Inflation

The postwar U.S. economy has experienced extended episodes of both low and moderately high inflation that permit direct comparison of the NAIRU model with our own. Conventional NAIRU models use a modified Phillips curve in which lagged inflation is taken as a measure of adaptive inflationary expectations, and the coefficients on lagged inflation sum to one. By contrast, our model allows the possibility that the coefficient on expected inflation will be lower in extended periods of low inflation than in extended periods of high inflation. Absent estimation biases, we would expect the coefficient to approach one in a sufficiently inflationary environment. We first look at the empirical evidence using the conventional adaptive expectations framework. We then provide evidence using direct measures of inflationary expectations that address Sargent's criticism of the assumption that the coefficient on lagged inflation must equal one in an accelerationist model.³² Sargent argued that a coefficient of less than one on lagged inflation may reflect not incomplete projection of inflation, but rather forecasters' views that the process generating inflation does not have

31. Modigliani and Papademos (1975).

32. Sargent (1971).

a unit root. By using direct measures of inflationary expectations, we can rule out the possibility that our results reflect differences in how people form expectations rather than in how they use them.³³

In order to separately estimate wage and price Phillips curves for periods of low and high inflation, we sorted the quarters since the Korean War according to the average CPI inflation rate in the five-year period ending each quarter. We first classified quarters with average annualized inflation rates below 3 percent as low-inflation quarters, and those with average inflation rates above 4 percent as high-inflation quarters.³⁴ By this sorting, the low-inflation quarters run from 1954:1 through 1969:1 and from 1995:3 through 1999:4, the end of our sample period. The high-inflation quarters run from 1970:2 through 1986:1 and from 1990:4 through 1993:2. As it happens, there are seventy-seven quarters both in the high-inflation sample and in the low-inflation sample. The mean CPI inflation rates in the two samples are 6.3 percent and 2.0 percent, respectively. This separation was used in half of the wage and price inflation regressions. In the other half we limited the low-inflation sample to quarters with inflation rates below 2.5 percent, which brought the sample size down to sixty-two quarters and reduced the mean CPI inflation rate in the low-inflation sample to 1.9 percent.

Estimates with Adaptive Expectations

The quarterly Phillips curve equations we estimated were intended to span the specifications that analysts have used in conventional estimation of NAIRU models, except for the fact that we did not constrain the coefficients on lagged inflation. To this end, we tried a large number of data combinations and specifications on both wage and price Phillips curves, and ran each separately for the low- and high-inflation samples just described. In all cases the dependent variable was an annualized inflation rate in either wages or prices, and the explanatory variables were current or lagged values of unemployment, price inflation, and, for the wage equations, trend productivity growth. For price inflation we used the CPI, the GDP deflator, and the personal consumption expenditure (PCE) deflator

33. We are grateful to a seminar participant at the Bank of Canada for suggesting this approach.

34. By sorting our sample on the basis of long lags of the endogenous variable, we considerably reduce concern about sample selection on the basis of an endogenous variable.

and estimated price Phillips curves with each. Twelve values of lagged inflation were used as explanatory variables. For wage inflation we used the best series available for any time period, linking private wages and salaries as measured by the employment cost index (ECI) for 1980–99 to the adjusted hourly earnings index for the nonfarm economy for 1961–80 and to adjusted hourly earnings in manufacturing for 1954–61. Twelve lagged values of CPI inflation were used as explanatory variables. For unemployment we used the rate for all workers (the total rate), the rate for 25- to 54-year-old males, and Robert Shimer’s demographically adjusted series.³⁵ We used the current and three lagged values of unemployment and, alternatively, the current and eleven lagged values. For the wage Phillips curves we used two estimates of trend productivity growth, one being a series based on work by Robert Gordon and the other the series we constructed for our 1996 paper.³⁶ We ran regressions with the productivity coefficient either freely estimated or constrained to be one (for the wage inflation equations), and with just the current trend or with the current plus seven lagged values of the trend.³⁷

The key results are summarized in figure 3 for equations explaining wages and in figure 4 for equations explaining prices. The figures present the results of 144 and 72 specifications, respectively. Each point represents the sum of the coefficients on lagged inflation estimated for the low- and the high-inflation samples for one specification. If the sums of the coefficients were similar for the two samples, the points would cluster along the 45-degree line. If they were similar and near one, the points would cluster near the upper right corner. In fact, for both wages and prices, and over the wide range of specifications and data we used, the points cluster near one on the high-inflation axis, but on the low-inflation axis they range from around zero to around 0.5 for the wage equations. This is consistent with the predictions of our model. The range on the price equations is broader and less conclusive. The third of the observations at the highest end of the range are from equations using the PCE deflator. The mean values of the coefficients on the high- and low-inflation axes, respectively, are 0.25 and 0.82 for the wage equations and 0.60 and 0.95 for the price equations.


35. Shimer (1998).

36. Gordon (1998); Akerlof, Dickens, and Perry (1996).

37. All the equations also used the customary dummy variables for the guidepost period of the 1960s and the price controls period of the 1970s, and used the difference between inflation with and without oil prices in 1979–80 as an additional variable.

Figure 3. Wage Equations: Sum of Coefficients on Lagged Inflation in Low- and High-Inflation Samples, 1954–99

High-inflation samples




Source: Coefficient estimates from 144 regression specifications as described in the text.

Direct Measures of Inflationary Expectations

As in Brainard and Perry's paper, the results just described cast doubt on conventional estimates with the NAIRU model. However, both analyses treat expectations as adaptive and so cannot refute Sargent's criticism that rational expectations are formed differently and that the coefficient on properly measured expectations might be one. We now address this issue by using direct measures of expected inflation as explanatory variables in place of distributed lags of actual inflation rates, while maintaining our division of the sample into periods of high and low inflation. The other explanatory variables are the same as those used in the regressions behind figures 3 and 4. We used the two direct measures of expected rates of inflation that are available over our sample period: that from the University of Michigan's Survey of Consumers and that from the Federal Reserve's Livingston Surveys.

Figure 4. Price Equations: Sum of Coefficients on Lagged Inflation in Low- and High-Inflation Samples, 1954–99

High-inflation samples




Source: Coefficient estimates from 72 regression specifications as described in the text.

Figures 5 and 6 plot the estimated coefficients on expected inflation for the variously specified wage and price regressions, respectively. As with the results using adaptive expectations, the coefficients on expected inflation are substantially different in the low- and high-inflation periods. For 288 wage equations the low- and high-period means are 0.29 and 0.85, respectively. For 144 price equations the means are 0.25 and 1.00, respectively.

These results support our general hypothesis even more convincingly than do the results with adaptive expectations. Not only do they address the point that the relevant coefficient for natural rate theory is not necessarily the coefficient estimated with adaptive expectations, but the results are as clear about price inflation as they are about wage inflation.

One possible objection to the results presented here and in the next section is that the lower coefficients on inflationary expectations during peri-


Figure 5. Wage Equations: Coefficients on Price Expectations Variables in Low- and High-Inflation Samples, 1954–99



Source: Coefficient estimates from 288 regression specifications as described in the text.

ods of low inflation are an artifact of measurement error. For example, if the variance of measurement error is constant whereas the variance of true inflationary expectations is higher in times of high inflation, then the coefficient on expectations could be biased toward zero more in times of low inflation than in times of high inflation. We investigated this possibility. Although it is true that the variance of expectations is higher in periods of high inflation, it is also true that the sampling error in both the Survey of Consumers and the Livingston Survey is higher. In fact, the sampling error is so much higher that the computed bias is *higher* in the low-inflation periods, imparting a bias *against* our finding that the coefficient on expectations is lower in periods of low inflation.

Sampling error may not be the only source of error in the survey expectations. Neither survey may be asking the right people with the right weights. In an attempt to approximate how much error this problem might

Figure 6. Price Equations: Coefficients on Price Expectations Variables in Low- and High-Inflation Samples, 1954–99

Source: Coefficient estimates from 144 regression specifications as described in the text.

introduce, we computed the bias that would be caused if the measurement error variance in expectations were equal to the variance of the residual of a regression of one of our survey expectations on the other. Again we found that the measurement error variance grew faster than the conditional variance of the expectations, so that the bias caused thereby would work against our finding that the coefficient on expectations was lower when inflation was low.

Estimating the Model

Previously we showed how a Phillips curve–type relation can be derived from our theoretical model (equation 15). In this section we present estimates of the model and of the rate of inflation that permits the

lowest sustainable level of unemployment. We also show the gain in employment compared with operating at the natural rate. This section will first discuss the specification of the model we estimate, then our benchmark results, and finally, an analysis of their robustness.

Specifications

In theory, with a large enough sample, it would be possible to estimate the full model presented above. The elasticity of demand (β), the parameter for the curvature of the unit cost function (α), and the parameters of the distribution of rationality thresholds (μ and σ) all have different effects on the objective function. However, in practice it was impossible to estimate more than the mean of the distribution of rationality thresholds and one of the other parameters. The reason is that all three of them—the elasticity of demand, the curvature of the productivity function, and the standard deviation of the distribution of rationality thresholds—act in much the same way to determine the impact of past rates of inflation on the cumulative normal term (see equation 15).

This lack of identification in practice can be understood if we consider a Taylor series approximation to the argument of the cumulative normal in equation 15, expanded around a value of zero inflation. There is no reason to expect that the argument will be exactly zero at zero inflation, so the constant term will likely be present. As noted, the first derivative of the firm's loss function with respect to inflation is zero at zero inflation and very small at most rates of inflation less than 10 percent. Thus the first-order term of the Taylor series expansion of the argument of the cumulative normal will also be zero. Second- and higher-order terms will be present, but our analysis of the loss function suggests that, with inflation between zero and 10 percent, and with reasonable values for the elasticity of demand, the curvature of the unit cost function, and the standard deviation of the distribution of rationality thresholds, the third-order and higher terms are unimportant. An approximation of the loss function of the form $E\pi^2$, where the constant E was chosen so that the approximation was exactly equal to the loss at 5 percent inflation, was never off by more than 3 percent of the loss. One parameter is all that is necessary to capture the effects of all three parameters from the model (α , β , and σ) on the derivative of the argument of the cumulative normal with respect to inflation.

We thus estimate a Phillips curve of the form

$$(17) \quad \pi_t = d + \Phi(D + E\pi_{L,t}^e)\pi_t^e - eu_t + gX_t + \epsilon_t,$$

where π is the rate of inflation, Φ is the cumulative standard normal density function, π^e is inflationary expectations, u is a term capturing the effects of current and lagged unemployment on inflation, X is a vector of dummy variables for oil shocks and price controls, ϵ is the error term, and $d, D, E, e,$ and g are parameters to be estimated.³⁸

The term π_L represents the effects of past inflation on the likelihood that people will act rationally toward inflation. Our theory tells us nothing about the way in which inflation should matter other than the sign of E , so we proxy π_L with several different, parsimonious specifications. The first is a geometrically declining, weighted moving average of past values of inflation:

$$(18) \quad \pi_L = (1 - \delta)\pi_{L-1} + \delta\pi_{L-1},$$

where δ is a parameter to be estimated.

Alternatively we estimate π_L as

$$(19) \quad \pi_L = \frac{\sum_{\substack{i=1 \\ i\lambda < 1}}^{24} (1 - i\lambda)\pi_{-i}}{\sum_{\substack{i=1 \\ i\lambda < 1}}^{24} (1 - i\lambda)},$$

where the parameter λ is estimated. Our final two specifications for π_L treat it as a four-year moving average of past inflation with equal weights, or with the relative weights of quarters from each year estimated (three additional parameters).

It is standard practice in Phillips curve estimation to proxy inflationary expectations with lagged values of inflation. In many specifications discussed below we follow that tradition. When we do, we use either a twelve-quarter unrestricted lag or one of the methods used to construct π_L to construct π^e . However, we also want to rule out the possibility that

38. This specification ignores the parameter a from the theoretical model. In theory that parameter could be estimated, but we do not take the theoretical model that literally. Instead we imagine that there is a continuum of reactions to increasing inflation, with people putting more and more weight on it until their behavior resembles that of the rational economic actor in the standard model. The model we estimate here can be thought of as a model where a fraction $(1 - \Phi)$ of people are ignoring inflation. Or the Φ function can be thought of as approximating a more general function that reflects how much weight the average person is putting on inflation in making economic decisions.

changes in the coefficient on π^e might reflect changes in the process by which expectations are formed rather than how they are used. Thus we also use direct survey measures of inflationary expectations for π^e in some specifications.

Our different specifications include several different measures of unemployment and different numbers of lags. The unemployment term, u , is constructed using one of three data series. The first is the aggregate U.S. unemployment rate from the Current Population Survey of the Bureau of Labor Statistics. Because this variable may be influenced by changing demographics, we have also considered two alternative measures: the unemployment rate for prime-age males and Shimer's demographically corrected series.³⁹ We also vary the number of unemployment lags from zero to eleven quarters.

For the dependent variable we variously use four different measures of inflation: the annualized percentage change in the urban consumer price index, the GDP deflator, the PCE deflator, and the index of wage and salary compensation constructed by Brainard and Perry.⁴⁰ When we use the percentage change in the compensation index as the dependent variable, we subtract off a measure of trend productivity growth. The three specifications of this trend are a measure based on Gordon, the measure we constructed for our 1996 paper, and a sixteen-quarter moving average.⁴¹

The form of the Phillips curve here is similar in some respects to the one in our 1996 paper that modeled the implications of downward nominal wage rigidity. Therefore we also examine whether we can successfully estimate a Phillips curve that embodies the insights from that model as well as the current one. Below we estimate a number of specifications that augment equation 17 with the term for nominal rigidity from that previous paper.⁴² When we nest that model, we must also estimate its key parameter—the standard deviation of desired wage changes—along with the other parameters from the current model.⁴³

39. Shimer (1998).

40. Brainard and Perry (2000).

41. Gordon (1998); Akerlof, Dickens, and Perry (1996).

42. The inclusion of the term for nominal rigidity could be motivated if we included firm profitability or firm-specific labor market considerations in the productivity function. That would produce heterogeneity in desired wage setting, with firms constrained by the floor of no nominal wage decrease forced to pay a higher wage as in the model in our previous paper.

43. See Akerlof, Dickens, and Perry (1996, appendix A) for its specification. We leave out the term for the change in profits, which could not be robustly estimated.

The model was estimated with quarterly U.S. data from the first quarter of 1954 through the last quarter of 1999, although we vary the end date in some specifications to check the extent to which our results depend on the experience of the 1990s. Data sources and the specification of the dummy variables for price controls and oil shocks can be found in table A1 in the appendix.⁴⁴ All the parameters of the model were estimated simultaneously by nonlinear least squares.

Results

Table 2 presents results for four different estimates with five types of variation: in the dependent variable, in the method of constructing π^e and π_L , in the unemployment measure and its lags, in the sample period, and in the inclusion of the term for nominal rigidity.

Our first focus of attention is the estimated value of the cumulative normal multiplying inflationary expectations when inflation is zero. In the theoretical model this corresponds to the fraction of firms behaving in a fully rational fashion at zero inflation. The model predicts that this fraction will be less than unity, and that as inflation increases above zero, the fraction of rational firms will rise. Both of these predictions yield tests of the model.

The NAIRU specification for the Phillips curve is nested in our model and can be obtained if the value of D is sufficiently high. For example, if D were 2 or higher, the coefficient on inflationary expectations would never fall below 0.97, and there would be little room for changing experience with inflation to affect the coefficient on inflationary expectations. All of the four estimated values of D imply coefficients on expected inflation of 0.5 or less at zero inflation. The lowest implies a coefficient of 0.19. In all four cases a value of D that would imply a coefficient of 0.9 or greater (1.28) can be rejected at conventional levels of significance.

The instantaneous effect of increasing inflation above zero can be computed as one minus the cumulative normal evaluated at D , divided by the sum of the coefficients on unemployment and its lags. Those values are about -1.2 or larger (in absolute value) in the specifications presented here. Thus, to a first-order approximation, raising inflation from zero to 1 percent will cause a reduction in unemployment of 1.2 percentage points or more.

44. We use dummy variables rather than an import price or energy price measure because we believe that these were atypical events that had atypical effects on the economy.

Table 2. Estimated Parameters for the Near-Rational Phillips Curve^a

<i>Independent variable and characteristics</i>	<i>Dependent variable</i>			
	<i>CPI</i>	<i>GDP deflator</i>	<i>PCE deflator</i>	<i>Compensation index minus productivity growth</i>
Constant	0.038 (0.008)	0.028 (0.008)	0.024 (0.012)	0.018 (0.003)
Unemployment	-0.51 (0.12)	-0.45 (0.12)	-0.43 (0.18)	-0.42 (0.07)
<i>D</i> (constant in coefficient on expectations)	-0.47 (0.36)	-0.88 (0.45)	0.02 (0.60)	-0.33 (0.23)
<i>E</i> (coefficient of π_t^2 in coefficient on expectations)	488 (152)	2,824 (1,120)	1,172 (708)	1,309 (368)
Standard deviation of desired wage change from term for nominal rigidity	Term not included	Term not included	0.022 (0.014)	Term not included
Method for constructing π_t	Geometric	16-quarter moving average (equal weights)	Geometric	Linear
Method for constructing π_t	Surveys of consumers	12 unrestricted lags	Geometric	Livingston
Unemployment measure and number of lags	Total; 0 lags	Total; 11 lags	Shimer; 7 lags	Male; 3 lags
Sample period	1954:1–1989:4	1954:1–1999:4	1954:1–1999:4	1954:1–1999:4
Natural rate of unemployment (percent)	7.5	6.4	5.6	4.4
Inflation at lowest sustainable unemployment rate (percent)	3.4	1.6	2.4	2.0
Lowest sustainable unemployment rate (percent)	4.5	4.4	4.7	2.3
Durbin-Watson statistic	1.4	2.0	1.9	1.1
<i>R</i> ²	0.824	0.698	0.717	0.729

Source: Authors' calculations.

a. Standard errors are in parentheses.

The term that most distinguishes our model from that of the textbooks is the coefficient of the square of lagged inflation in the cumulative normal multiplying inflationary expectations (E). If E is zero, the coefficient on expectations will not vary with past rates of inflation. Our theory says it should, and that is what we find in each of the specifications we have estimated. In all four specifications presented above, E is large, and more than 1.65 times its estimated standard error. Going from zero to 5 percent inflation would increase the argument of the cumulative normal by 1.2 to 7.0, depending on the specification. Except in specifications with CPI inflation as the dependent variable, the coefficient on inflationary expectations is above 0.95 by the time inflation has reached 4 percent. For the CPI specification the coefficient is 0.6 at 4 percent inflation and rises above 0.95 at about 6.5 percent.

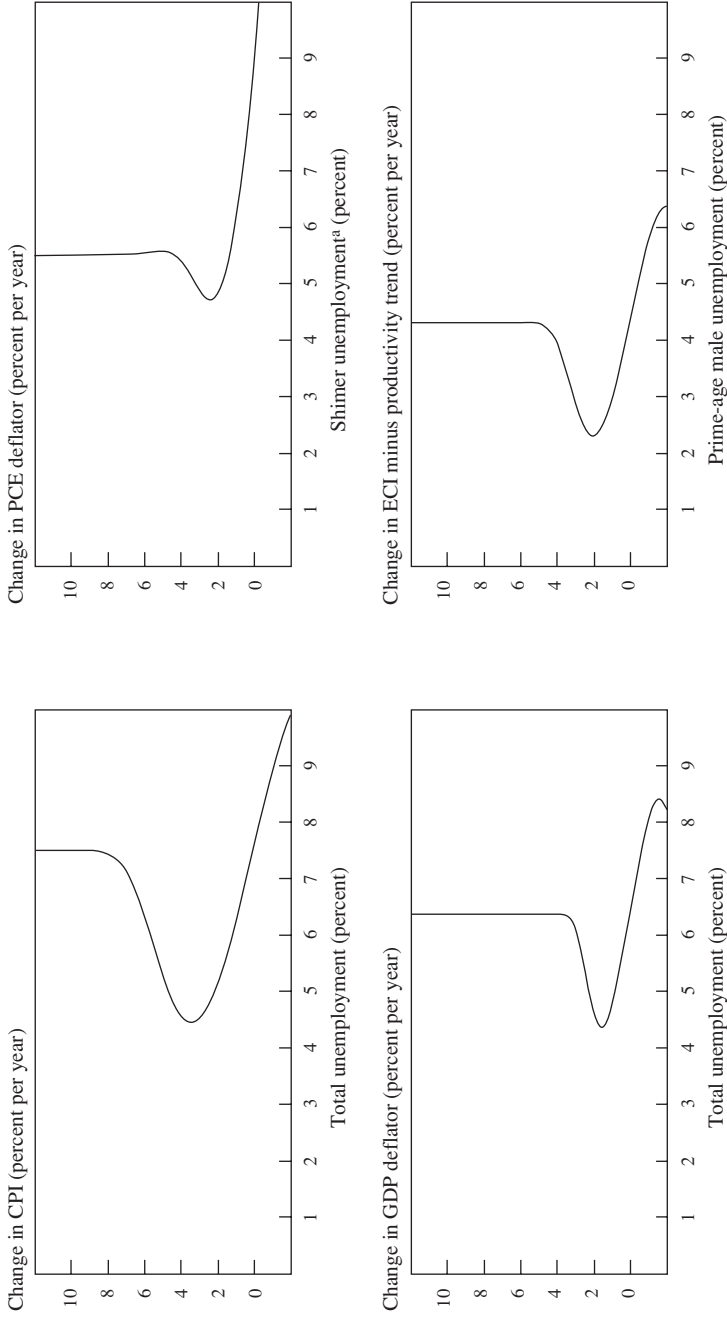
Besides allowing us to estimate the effect of inflation on the use of inflationary expectations, estimating our model also allows us to calculate the lowest sustainable rate of unemployment and the accompanying rate of inflation. We have computed these rates of inflation for each of the four models in table 2 from the estimated parameters numerically. We have also computed the natural rate in each model and the lowest sustainable rate of unemployment. Inflation at that unemployment rate ranges from 1.6 to 3.4 percent. The difference between the natural rate and the lowest sustainable rate of unemployment ranges from 1.5 to 3.1 percentage points. Figure 7 shows the long-run relationship between inflation and unemployment implied by each of the four specifications estimated in table 2.

The values of the coefficient of inflationary expectations implied by our parameter estimates are plotted in figure 8 for each of our four specifications. In all cases the coefficient values vary considerably over the sample. In all four specifications the coefficient on inflation reaches a maximum value of one for at least a year at some point during the sample period in the early to mid-1970s. The four specifications differ in the exact timing of the increase in the 1970s, in how the 1950s and the 1990s are treated, and in the date of the end of the period of a coefficient of one on inflation.

These figures can be compared with the time path Brainard and Perry estimated for the coefficient on inflation.⁴⁵ Our estimates imply consider-

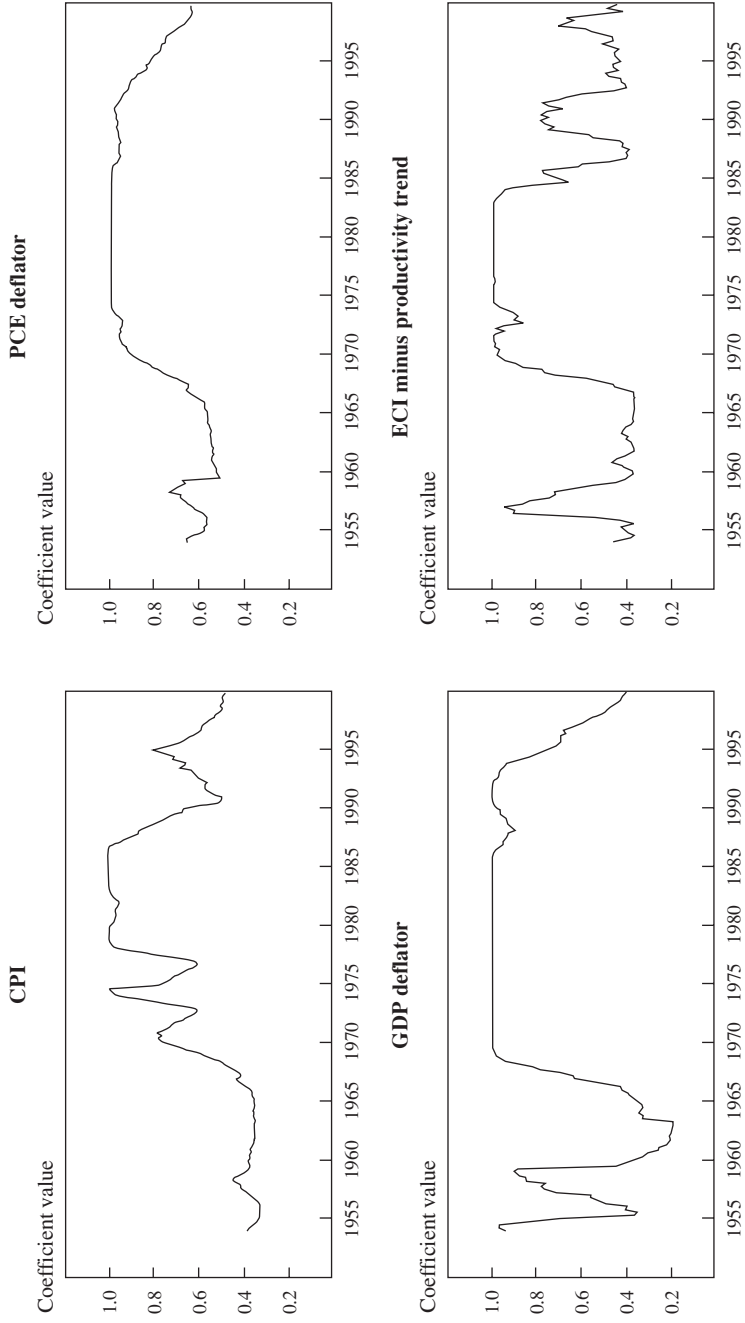
45. Brainard and Perry (2000).

Figure 7. Long-Run Inflation-Unemployment Relation Estimated with Various Price Indexes



Source: Authors' calculations from the estimated parameters reported in table 2.
 a. Shimer's (1998) demographically adjusted unemployment rate.

Figure 8. Coefficient on Inflation in Models with Various Price Indexes, 1954–99



Source: Authors' calculations from the estimated parameters reported in table 2.

ably more abrupt changes and more persistence. They also imply more variation. However, it must be remembered that the method Brainard and Perry used to estimate their values for the coefficient on inflation imposes smoothness on the changes. When we smooth our estimates (results not shown), they begin to resemble the time path that Brainard and Perry found, with one major difference. The Brainard-Perry estimates peak earlier and fall off more abruptly than our smoothed estimates.

We have varied the specifications presented above to anticipate possible objections to our results. The specification with the CPI as the dependent variable, which is estimated using only data through 1989:4, shows that our results do not depend on the experience of the 1990s, which may be atypical. Since nonlinear estimation is difficult when many parameters are being estimated, we have generally used very parsimonious specifications for the lags on past price inflation when constructing inflationary expectations. One might object that this parsimony forces the coefficient on inflation to do the work that a richer lag structure would do. The specification where the GDP deflator is the dependent variable answers this by matching the richest possible lag structure for price expectations (twelve-quarter unrestricted) with the most parsimonious specification of the term in the coefficient of expectations (a sixteen-quarter equally weighted moving average with no free parameters). Likewise, in most specifications, the inclusion of lagged unemployment, or our term for nominal rigidity, or both does not change our fundamental results.

Our Durbin-Watson statistics for the two specifications using survey expectations show considerable serial correlation. We have not attempted to correct for this problem because we lack a credible instrument for price expectations, which are endogenous with respect to the error in the Phillips curve. We are unhappy with this drawback of the analysis, but estimates of our model that we have tried using simulated data suggest that the bias from ignoring the serial correlation in the parameters we care about is minor.

Robustness of the Results

As noted above, there are many aspects of the specification that are not dictated by the theory. Our approach to this problem has been to estimate a wide array of different specifications to determine whether our primary results are sensitive to changes in the specification.

Because both the estimation of the model and the numerical analysis of the results currently require human intervention, we have not been able

to mechanize the process of sensitivity testing. Thus we have not been able to do an exhaustive specification search. Instead we estimated 218 different specifications, many of which were run to test specific concerns. However, most were chosen randomly.⁴⁶ Our survey of the results of these specifications yields the following generalizations.

The most important is that nearly all the point estimates imply that significant gains in employment are possible by increasing inflation from zero to a rate above 1.5 percent. This can be seen in figure 9, which plots for each specification the rate of inflation at the lowest sustainable unemployment rate and the reduction in unemployment that one obtains from increasing inflation from zero to that rate. In only twelve specifications was the estimated gain less than 1 percentage point, and in only one was it negligible. This specification was a wage Phillips curve with a rich lag structure for price expectations in which the inflation term in the coefficient on expectations was constrained to be an equally weighted, sixteen-quarter moving average of past inflation. Allowing a richer specification for the impact of inflation on the use of expectations eliminates this result. Of the other eleven specifications where the estimated impact is less than 1 percent, all are at least a half a percentage point. Most of the specifications are wage equations, and none use the PCE deflator as the dependent variable. Only one uses survey expectations. In no case are the parameters of the inflation coefficient very precisely estimated, so that values more typical of other specifications cannot be ruled out.

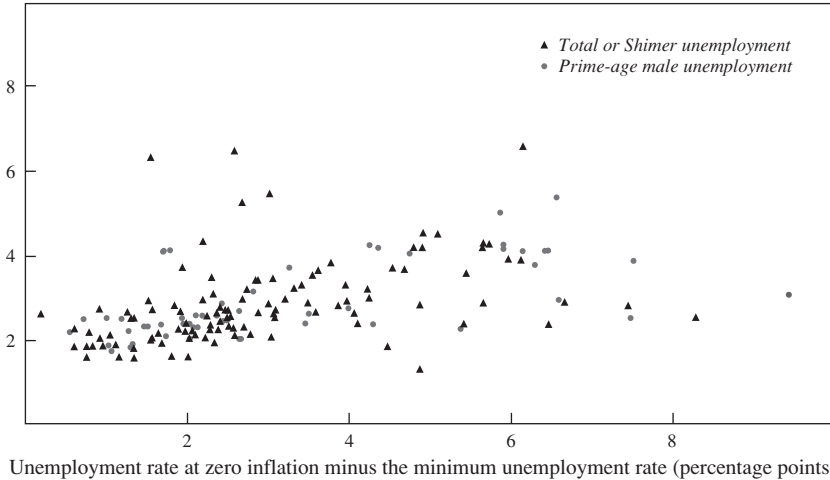
A second generalization is that it is not possible to robustly distinguish the relative importance of the effects of nominal rigidity and those of near rationality. The majority of specifications that included our term for the effects of nominal rigidity give results like those for the PCE in table 2. These do suggest a role for both nominal rigidity and near rationality. However, in many specifications that include both effects, the effect of past inflation on the coefficient of expectations is not measured precisely, being about the same size as its estimated standard error.⁴⁷ In other cases, the

46. We set a goal of 200 specifications, met that goal, and then estimated a few more to check specific concerns that arose in the process of evaluating the original 200 specifications. In randomly choosing specifications we allowed all options with equal probability, except that we found that the twelve-quarter unrestricted lag on inflation for the price expectations term was always computationally burdensome, so we did not include those specifications in those that were randomly chosen.

47. In contrast, when the term for nominal rigidity was not included, the coefficient on the square of past inflation was nearly always at least 1.7 times its estimated standard error or more.

Figure 9. Robustness Analysis for the Nonlinear Model

Rate of inflation at the lowest sustainable unemployment rate (percent per year)



Source: Coefficient estimates from 218 regression specifications as described in the text.

optimization routine was trying to drive estimates of the standard deviation of desired wage changes to zero. In six specifications not represented in figure 9, we obtained converged estimates for the parameters, but the estimated values for σ were sufficiently large that there was no single rate of inflation at which the unemployment rate was minimized. It simply fell to the natural rate asymptotically as in the models estimated for our 1996 paper.⁴⁸

We encountered few problems with applying nonlinear estimation. We did look for and find a few cases where there were multiple local minimums, but these reflected minor differences in the lag structures that were not substantive. Of the 218 specifications we estimated, we were unable to obtain converged values for about 30. This might be a serious concern because, under the hypothesis of fully rational behavior, the model's parameters are not identified. It might be that the nonlinear estimation program is trying to drive the constant term in the coefficient on inflationary expectations to infinity in order to drive the coefficient on expectations to

48. Akerlof, Dickens, and Perry (1996).

one.⁴⁹ However, this is not what was happening in any of the cases of convergence problems that we encountered.⁵⁰

Overall, the results from estimating our model support the theory we have laid out. They suggest that macroeconomic policy should aim for a rate of inflation in the range of 1.5 to 4 percent. Either higher or lower rates seem likely to result in lower output and employment.

Conclusion

This paper offers an alternative to natural rate models of unemployment. Natural rate models provide a wonderful economics “just-so” story, based on the idea that firms and workers take full account of expected inflation in setting current wages and prices. This behavior produces a unique long-run unemployment rate, the natural rate, that is consistent with any steady rate of inflation, and a short-run Phillips curve in which unemployment above or below the natural rate causes inflation to decelerate or accelerate.

Our model of the macroeconomy rests on different behavioral underpinnings, which are supported by a range of related evidence, including the psychological literature on decisionmaking and perception, direct survey evidence on how people react to inflation, and the advice of compensation professionals. We propose that when inflation is low it is not especially salient, and wage and price setting will respond less than proportionally to expected inflation. At sufficiently high rates of inflation, by contrast, anticipating inflation becomes important, and wage and price setting responds fully to expected inflation. This behavioral difference between our model and the natural rate model has significant implications both for estimating the relation between inflation and real activity in the macroeconomy and for informing the conduct of macroeconomic policy.

49. When we generated simulated data with a standard Phillips curve model and attempted to estimate our model on it, this is what happened.

50. Instead we had one of three other problems. Either the program was trying to drive σ to zero. Or the program was driving the constant term in the coefficient of expectations to *negative* infinity, and the coefficient on the square of past inflation to infinity, to eliminate coefficient values between one and the lower floor. Or, in some very rich specifications, the first derivatives of a group of unrelated parameters became so close to collinear that it was impossible to invert the approximation to the Hessian used in the maximization routine.

Our model is supported by the evidence. In a conventional Phillips curve, estimates of coefficients on expected inflation, whether measured by lagged inflation or by direct surveys of expectations, are greater when inflation is high than when it is low. Estimates of our model provide further support. Rather than a natural rate of unemployment that is invariant to the rate of inflation, our model traces out a range of equilibrium unemployment rates associated with different ongoing inflation rates. The lowest sustainable unemployment rate is the minimum of this range. The natural unemployment rate is a special case: it is the equilibrium unemployment rate at high inflation rates (and, ignoring downward wage rigidity, at zero inflation). The natural rate is noticeably higher than the lowest sustainable unemployment rate. The rate of inflation that accompanies the lowest sustainable unemployment rate is low, perhaps not far from current values, but not zero. Operating with inflation either higher or lower leads to a higher rate of unemployment in the long run.

The distinctive feature of our model is especially important for estimation. In recent years, as low inflation rates have come to be the norm, NAIRUs estimated from the empirical counterpart of the natural rate model have proved to be misleading guides to policymakers and economic analysts. In the mid-1990s these models typically projected 6 percent as the lowest sustainable unemployment rate, yet real output has grown at a 4 percent annual rate since then, and the unemployment rate has fallen as low as 3.9 percent. The NAIRUs estimated for the early 1960s, the previous period of moderate inflation, also appear unrealistic. When adapted for estimation, the model we have developed should provide more useful estimates of the attainable levels of employment to serve as a guide for stabilization policy and as an anchor to longer-run projections.

Not only does our model fit the data better than NAIRU models, it is also more cogent theoretically. NAIRU models serve well as what Irving Fisher might call “the first approximation.” They are derived from the assumption that all people behave according to what economists call economic rationality, or else their deviations from that behavior are totally random. This paper relies, as a first approximation, on exactly such economic thinking. But Fisher also made “a Second [and even a Third] Approximation.”⁵¹ With aggregate Phillips curves, such further approximations involve departures from perfectly rational decisionmaking. The

51. Fisher (1930).

evidence available on the subject suggests that, in setting wages and prices, the lay public do not use the same model of the economy as economists. Given the complexity of their decisions and, for the most part, their lack of training as economists, it would, indeed, be surprising if they did. It is thus highly unlikely that the welter of *interdependent*, intuitively based decisions of a real economy will produce a coefficient of inflationary expectations on wage and price inflation that is always exactly one. This paper has offered a theory for such a departure as price and wage setters under-adjust for inflation when it is not very salient and when the cost of such behavior is low. This theory yields a lowest sustainable rate of unemployment and an accompanying rate of inflation. It also fits the facts.

APPENDIX A

Table A1. Data Sources and Variable Construction

<i>Data series</i>	<i>Source and description</i>
Consumer price index (CPI)	Bureau of Labor Statistics (BLS) official data on the seasonally adjusted average of the CPI-U, 1982–84 = 100. An experimental CPI calculated using geometric means, CPI-U-XG, was used for the period 1990:1 to 1998:4. The CPI-U-XG was seasonally adjusted using quarterly dummies. Data were extracted from the BLS website's Selective Access page, www.bls.gov/sahome.html . We used quarterly averages of the change in the logarithms of monthly data. (The CPI-U-XG data were not used for some of the specifications in the robustness check.)
Gross domestic product (GDP) deflator	Bureau of Economic Analysis (BEA) official data reflecting the October 1999 benchmark revision for data back to 1959:2. We used unrevised data before 1959:2 by necessity. Data were extracted from the BEA website, www.bea.doc.gov . We used quarterly averages of the change in the logarithms of monthly data.
Personal consumption expenditure (PCE) deflator	BEA official data reflecting the October 1999 benchmark revision for data back to 1959:2. We used unrevised data before 1959:2 by necessity. Data were extracted from the BEA website. We used quarterly averages of the change in the logarithms of monthly data.
Wage inflation	This series comprises the best available BLS series for each time period: employment cost index (ECI) wages and salaries for 1980:2 through 1999:4, the adjusted hourly earnings index for the nonfarm economy for 1964:2 through 1980:1, and adjusted hourly earnings in manufacturing for 1954:1 through 1964:1. We use the change in the logarithm of quarterly data.
ECI wages and salaries	BLS official data on seasonally adjusted, private industry wages and salaries of all workers, extracted from the BLS website's Selective Access page.
Adjusted hourly earnings index	BLS official data on seasonally adjusted, current dollar average hourly earnings (1977 weights) in the nonfarm economy. Data were extracted from the BLS website's Selective Access page.
Adjusted hourly earnings in manufacturing	BLS official data on seasonally adjusted, current dollar average hourly earnings (1977 weights) in the manufacturing sector. Data were extracted from an unpublished BLS series.

(continued)

Table A1. (continued)

Data series	Source and description
Gordon productivity trend	We constructed this variable using estimates of trend productivity over several time periods provided by Robert Gordon, Northwestern University. We smoothed the transitions so that the variable is equal to 0.0263 from the start of the sample through 1971:2, after which it changes over nine quarters to a value of 0.0124, which is maintained through 1986:2, after which it drops to 0.0094 over eight quarters. From 1988:2 through 1994:4 it is constant and then rises over seven quarters to 0.0153, where it remains until the end of our sample. The exact series is available from the authors by request.
Akerlof-Dickens-Perry productivity trend	This variable is equal to 0.0230 from the beginning of our sample through 1973:4, then declines over nine quarters to a value of 0.0100, where it remains through the rest of the sample. The values for the quarters 1974:1 through 1975:4 are 0.0216, 0.0201, 0.0187, 0.0172, 0.0158, 0.0144, 0.0129, and 0.0115.
Surveys of Consumers Inflation Expectations series	The University of Michigan's Survey Research Center, which conducts the Surveys of Consumers, asks its respondents how much they expect prices to go up or down on average during the next twelve months. We use the median response to the survey in the last month of the previous quarter as a measure of inflationary expectations. Studies of this series suggest that it provides the most accurate forecast of inflation one quarter ahead. A complete survey description can be found at http://athena.sca.isr.umich.edu/scripts/contents.asp .
Livingston Survey of Inflation Expectations	This survey was originally conducted by Joseph Livingston, a journalist with the <i>Philadelphia Inquirer</i> , and has been continued by the Philadelphia Federal Reserve Bank since his death in 1990. The survey is conducted every January and June and asks respondents, who are professional economists, their forecast for the level of the CPI six months hence. We use the forecast of the inflation rate made in January for the second and third quarters and the forecast made in June for the fourth and following first quarters. See http://www.phil.frb.org/files/liv/document.html for a complete description of the survey.
Unemployment rate	BLS official data on the total unemployment rate among the nation's labor force. Data were extracted from the BLS website's Selective Access page at http://www.bls.gov/webapps/legacy/cpsatab9.htm .

(continued)

Table A1. (continued)

<i>Data series</i>	<i>Source and description</i>
Unemployment rate for males aged 25–54	BLS official data of the unemployment rate among men aged 25–54 in the labor force. Data were extracted from the BLS website’s Selective Access page at http://www.bls.gov/webapps/legacy/cpsatab9.htm .
Shimer demographically adjusted unemployment rate	Robert Shimer’s demographically adjusted unemployment rate. For full details see Shimer (1998). Values for 1998:3 through 1999:4 are too recent to be in his series, so we took the arithmetic average of the difference between the Shimer series and the total unemployment rate from 1954 to 1998 and added that to the total unemployment rate to arrive at values for the Shimer series.
Guidepost dummy	A dummy variable set equal to one from 1964:1 through 1966:2. This was used only in the split-sample regressions.
Wage and price controls Phase One dummy	A dummy variable set equal to one from 1971:3 through 1972:4.
Wage and price controls Phase Two and 1973 supply shocks dummy	A dummy variable set equal to one from 1973:1 through 1974:1.
Wage and price controls removal dummy	A dummy variable set equal to one from 1974:2 through 1975:1.
1979 oil price shock dummy	A dummy variable set equal to one from 1979:1 to 1980:1 and for 1981:1.
Price wedge	The change in the total CPI-U minus the change in the CPI-U less food and energy. Used as an alternative to the above dummies in some of the split-sample regressions to account for fluctuations in food and energy prices. Data for both series were extracted from the BLS website’s Selective Access page. This variable was used only in the split-sample regressions.

Comments and Discussion

Truman F. Bewley: I would like to interpret what this paper says in terms of my own experience interviewing employers about wage and employment policies. The paper's main assertion is that the perceived external rate of wage inflation has less impact on a firm's wage policy when the rate of inflation is low than when it is high. I will suggest that there may be an explanation of this phenomenon different from that given by the authors. The phenomenon itself is, no doubt, real, given their persuasive empirical results. However, the weak feedback of inflation on wages at low rates of wage inflation may have an explanation involving the mechanics of wage setting and may have little to do with near-rational behavior.

In order to present my explanation, I have to describe some things I observed. The authors make the Keynesian assumption that morale depends on the relationship between a firm's wages and a reference wage, which is roughly the average of wages of other firms competing in the same labor market. I do not think this assumption is correct. Most employers that I have talked to say they think their workers have only a very vague idea about pay rates outside their own firm. External pay rates therefore have little impact on morale. This is especially true in nonunion companies. Instead workers focus almost entirely on internal comparisons: the wages they use as references are their own past pay and the pay of co-workers at the same workplace. Enough employees know each other's pay that comparisons internal to a firm are an important compensation issue, at least for workers who have long-term, full-time jobs. If pay rates are internally inequitable, employees complain and morale falls apart, causing management major problems. Past pay is an obvious reference, since everybody remembers his or her own compensation history.

The importance of this reference wage makes it difficult to cut nominal pay.

Because of the significance of internal pay equity, firms normally have a strict internal pay structure that relates pay to tenure on the job, to skill, and to many other things, such as experience and job level. Internal structures vary in form from firm to firm. They are part of a firm's culture and must be followed rigorously if a firm is to remain healthy. The structure may be described explicitly in a booklet, especially if the firm is large. This document gives instructions to the managers who actually determine the pay of individual workers.

The relationship between the pay of different firms is the result of labor market pressures. That is, the overall level of the pay structure is determined by the ability to recruit and retain workers. Firms usually adjust the whole level of their structure, moving pay rates at different levels in concert.

It is very risky for a firm to cut the nominal pay of individual workers, unless the firm is in danger of closing and its work force clearly understands this danger. It may be possible, however, to reduce the nominal level of the pay structure without cutting the pay of any worker. There is an important distinction between the level of a pay structure and the level of individual workers' pay, a matter to which I will return. It may also be risky to allow the purchasing power of workers' pay to fall, although reductions in the real value of pay have less serious consequences than do nominal pay cuts.

The choice of pay structure level is a strategic business decision discussed at high levels of management. In making this selection, business leaders make use of labor market surveys and look at their own experience in recruiting and retaining workers. They also take account of the fact that, if they raise the level of pay today, it will remain high in the future.

I hear a lot about this last point now. I am currently surveying pricing practices, and so I talk to businesspeople almost every day. Some say that they are not now increasing pay to ease labor shortages, because they know they will not be able to reverse the increases during the next economic downturn.

It is important that, in setting pay levels, businesspeople look at wage surveys. They buy surveys, make their own surveys, try to project rates of wage inflation, and look at inflation as measured by the consumer price index as an indicator of what other firms will do with regard to pay. They also talk to competitors in the labor market and in the product market.

They discuss what they themselves are doing with their pay structures, without giving away information that would help other firms compete with them in product markets. Employers need this information to be sure they will pay enough to recruit and retain labor.

Because employers discuss wage setting among themselves, there may not be a direct and mechanical feedback from projected wage increases to wage policy, as the paper assumes. Through mutual consultation, businesspeople can, I believe, coordinate a slowdown of wage increases. I do not mean conscious coordination and am not implying collusion, but they may create an atmosphere that decreases the size of the pay increases perceived as needed. Such an ambience would be beyond the control of any one firm, but would nevertheless be achieved. It is not clear how much coordination occurs in this way, but businesses do keep each other informed as to what they are doing with pay levels.

I now turn to the distinction between the pay of individual workers and the level of a pay structure. The former tends to increase over time, because workers rise through the internal pay structure as they acquire seniority, gain skills and experience, and become more valuable to the company. Workers in unskilled jobs may hit a pay ceiling after two or three years. But there are many skilled workers in the U.S. economy, and in many firms these well-paid workers account for a large share of labor costs. These workers may continue to rise in the pay structure throughout their career.

Because most people are rising in the pay structure, they receive pay increases every year even if the structure stays constant. If there is only a little inflation in the cost of living, these increases might exceed that rate of inflation. A company can even decrease its pay structure while giving everybody real and nominal increases.

A firm's average cost of labor and its average and marginal cost of output are determined by the level of the pay structure. Hence, if a firm gives an average wage increase of 5 percent, it does not follow that its average cost of labor rises by 5 percent. It may even decrease. The impact on the average cost of labor depends on the steepness of the company's internal pay structure, on how fast employees rise within it, and on labor turnover. The increase in the average cost of labor tends to be less than the pay increase of an average worker because of what are called turnover savings. These savings result from the fact that the average pay of workers entering a firm is less than the average pay of those leaving.

Department stores and supermarkets depend on turnover savings to control labor costs. Stores offer a schedule of pay increases to floor staff to provide work incentives, and these increases are paid for, to some extent, by turnover savings. One of the disadvantages of economic downturns is that they cause turnover to decrease, making existing staff steadily more expensive.

Firms have considerable latitude in choosing the level of their pay structure, for that choice is not a clear-cut decision but rather a judgment call. A firm can allow its pay structure to lag behind that of other firms. It may choose to do so if it is in trouble financially, if it considers that it can tolerate higher labor turnover or lower-quality employees, or if it believes that its labor recruiters should work harder to find new workers. Social mood may also affect the level of the pay structure, and I feel that this mood has changed now from what it was during the recession of the early 1990s. Employers are talking tougher about pay now than they were then. Whereas in the 1990s they talked about keeping up with other firms, now they talk about avoiding excessive labor costs. This mood shift may help explain why we have recently had such low inflation.

Imagine that a firm is holding its pay structure constant. A lot of the firms whose managers I interviewed were doing just that. There may be reasons for holding it constant despite recessionary conditions. A pay structure's level is ultimately tied to starting pay rates, and a firm that has a reputation as a premier employer may think that potential employees are aware of changes in starting rates. Such companies imagine that declines in these rates would look bad. Other firms do not care. For example, fast-food restaurants, supermarkets, and department stores allow starting pay and hence the pay structure for floor staff to float up and down freely with labor market conditions.

Assuming that the pay structure remains constant, let us think about the feedback from price inflation to the wage—not the feedback from wage inflation, which is the subject of the paper. I am changing the topic somewhat. If price inflation is small, the firm may not be under any pressure to raise wages to offset it, because most employees automatically receive raises large enough to compensate for it.

It might be imagined that, on the contrary, wage increases should equal normal increases plus the rate of price inflation, no matter how small that rate is, because inflation reduces the value of the normal increases and hence reduces their impact on incentives. The incentive effects are, after

all, an important reason for having the raises that are incorporated in the internal structure. The reduction in the real value of the increases does not seem, however, to have much impact on motivation, as long as the real value of the increases remains positive. Employees think about their increases relative to those of co-workers, so that the main impact on incentives is achieved simply by rising within the internal pay structure.

In summary, it could be that companies do not have to react at all to low price inflation by increasing wages, because people are, in any case, receiving annual increases that offset the inflation. Employees may react if inflation is so high that living standards decline. But if living standards do not decline, people do not suffer from the inflation and believe that the normal increases reward their efforts adequately.

Another reason for increasing the pay structure would be competition for labor, and a firm may decide not to compete. I have talked to a number of managers lately about this topic in the context of the effects of capacity constraints on prices. When I ask about capacity, a typical response is, "Our capacity limitation is labor. We can't operate a third shift, because we can't get the workers."

"So why don't you raise wages to attract more labor?"

"Well, sure, we could raise wages, but it wouldn't do much good. There is not enough labor out there."

That is one argument. Another is, "We don't want to get stuck with high pay."

This is the kind of strategic thinking that goes into determining the level of a pay structure. Although in today's low-inflation environment the primary consideration is not the general rate of price inflation, the situation would probably change if price inflation were high enough to exceed the rate of pay increase that most people enjoy as a result of rising through the pay structure. In that case, people's living standards would decline, and when living standards fall, people cannot afford the things they have grown used to, and they become angry with their company. Business executives take such reactions seriously. When wages are increased to compensate for price inflation, the pay structure rises, and with it the firm's marginal costs of output.

Thus, labor turnover and the internal pay hierarchy create a distinction between the increase in the average cost of labor to a firm and the increase in the average pay of individual workers. And the choice of the level of a firm's pay structure is dominated by competition for labor, except when

the rate of price inflation is so high that firms have to increase their structures just to protect employee living standards. Living standards tend not to be an important factor at low rates of inflation, because the living standards of many or most workers are protected by the raises that they would receive even if there were no inflation. These points give another possible explanation for the paper's observation that the sensitivity of wages to inflation increases with the rate of inflation.

Alan S. Blinder: This is a throwback paper. I say that without disparagement, however. After all, the Phillips curve was published in the year of my bar mitzvah. George Akerlof, William Dickens, and George Perry (henceforth ADP) offer us a fascinating and skillfully executed paper, and the fact that it is decidedly old-fashioned does not mean that it is wrong. The authors' objective is to bring money illusion back to the Phillips curve, from whence it was banished decades ago. To put their latest effort into historical perspective, I ask you to remember The Lone Ranger and "return with me now to those thrilling days of yesteryear." Money illusion rides again.

The shape of the Phillips curve. The original Phillips curve, as presented by A. W. Phillips and Richard Lipsey,¹ took the form $\dot{w} = f(U) + a\pi$. If we give the inflation term on the right-hand side an expectational interpretation and append a fixed markup, this gives rise to the usual price-price Phillips curve, which is sloping in both the short and the long run.

Subsequently, Milton Friedman and Edmund Phelps pointed out that such equations embody money illusion if $a < 1$: each 1 percent increase in inflation raises \dot{w} by only a percent and hence reduces real wage gains.² This, they pointed out, is fundamentally illogical. The parameter a must be exactly 1.0, which makes the long-run Phillips curve vertical at the natural rate.

In their paper, ADP embody money illusion in the Phillips curve when inflation is low, but not when it is high; this leads to the curve shown in their figure 1. Its odd shape comes from the following reasoning.

First, remember that we are looking at a long-run Phillips curve here, so we are tracing out alternative steady states. Start at zero inflation. Since there can be no money illusion when inflation is zero, unemployment is at the natural rate. At modest inflation rates, money illusion kicks in, mak-

1. Phillips (1958); Lipsey (1960).

2. Friedman (1968); Phelps (1967).

ing labor in the near-rational sector cheaper and thereby (as Keynes might have put it) raising the money supply in wage units. In the negatively sloped portion of the curve, higher inflation leads to more money illusion and therefore to lower unemployment, just as in an old-fashioned Phillips curve. But at even higher inflation rates, more and more workers and firms shake off their money illusion, and the economy relinquishes these employment gains. At high enough inflation rates (roughly above 6 percent in the authors' illustrative figure), all money illusion is gone and the Phillips curve is vertical for the usual reasons.

Now compare this Phillips curve with the one the authors offered in their earlier Brookings Paper (henceforth ADP-I).³ The two carry quite different empirical and policy implications at low inflation rates. ADP-I warned us that low inflation could be catastrophic because of the zero floor on nominal wage increases. Their Phillips curve veered sharply to the right, toward extremely high unemployment rates, as inflation fell to low levels. ADP-II argues, in stark contrast, that some low, positive inflation rate will lead to minimal unemployment—that is, unemployment below the natural rate and, indeed, stunningly below it in some specifications.

At least superficially, U.S. data since 1996 seem far more consistent with ADP-II than with ADP-I, which, I suppose, is why ADP-II was written. The authors are entirely correct to point out that the mechanisms of the two papers—nominal wage floors in ADP-I and money illusion in ADP-II—are not inconsistent. In fact, they are first cousins: one can meld them into a single model. But the contrast between the Phillips curves in ADP-I and ADP-II raises an important practical question for the Federal Reserve: should it welcome, or shun, inflation in the range of, say, 2 percent?

To pursue this question, I will concentrate on an important property of the Phillips curve that ADP emphasize in this paper. The leftmost point in their figure 1 represents the lowest rate of unemployment that is *sustainable* and the inflation rate that accompanies it. I would like to raise two questions about this point: Is it optimal? And is it sustainable?

The answer to the first question is, obviously not. The ADP model allows for no explicit costs of inflation. Were any such costs of inflation added to the model, the optimal point would be one with lower inflation and higher unemployment than the leftmost point in figure 1.

3. Akerlof, Dickens, and Perry (1996).

And there is a second point. Since we have no hard-core “Chicago school” critic on the Brookings Panel, let me assume that role for a moment. We Chicago types hasten to remind the authors that money illusion is an error: reacting to nominal magnitudes as if they were real magnitudes is a mistake that leads to bad decisions. ADP do not offer any utility analysis of labor supply. But if they did, it would presumably show that utility declines because people work too much given the low level of real wages. (In fairness, this utility loss will be small.)

Who's got the money illusion? What about sustainability? Here, once again, I find it useful to go back to “those thrilling days of yesteryear” and remember some old-time Keynesian economics.

We old-timers used to tell our young charges that the distinguishing theoretical feature of Keynesian economics was *nominal wage rigidity*. That assumption, in turn, combined two distinct notions: first, that the labor market is characterized by downward wage rigidity rather than market clearing, and second, that nominal, not real, wages are rigid (because of money illusion). I used to pose the following question to my students: Which assumption produces the characteristic Keynesian policy result that raising aggregate demand will boost employment? The answer is, the second. If it is real wages that are rigid, the model generates classical unemployment that cannot be cured by expanding demand. On the other hand, if the labor market clears, but with money illusion in the labor supply function, higher aggregate demand raises the price level, shifts the labor supply curve outward, and boosts employment. This conclusion is so old that it may seem new to young economists, many of whom have never been exposed to such archaic musings.

Why did Keynes assume that workers suffered from money illusion? Although he never lived to read Kahneman and Tversky, Keynes's “evidence” was not so different from ADP's. He wrote:

Now ordinary experience tells us, beyond doubt, that a situation where labour stipulates (within limits) for a money-wage rather than a real wage, so far from being a mere possibility, is the normal case. . . . It is sometimes said that it would be illogical for labour to resist a reduction on money-wages but not to resist a reduction of real wages. . . . But, whether logical or illogical, experience shows that this is in fact how labour behaves.⁴

4. Keynes (1936, p. 9).

I have tortured you through this historical detour for two reasons. The first is to raise a question: Who has the money illusion? ADP tell their readers it does not matter: as long as one side of the wage bargain has money illusion, it all comes out the same. They write, “A whole spectrum of other models in which various combinations of firms and workers ignore or underweight inflation in their mental frames will yield similar results.” Their specific model takes up the case in which firms, not workers, are “near rational.” But if labor demand is a decreasing function of the real wage, it seems to me that firms plagued by money illusion would *underdeflate* the money wage, and hence behave as if the real wage were *higher* than it actually is. If the labor demand schedule is downward sloping,⁵ such firms will hire *less* labor than in the neoclassical equilibrium, not more.⁶

Second, we old-fashioned Keynesians used to argue that money illusion, being a clear departure from rationality, should disappear over time. Thus the money illusion that characterizes the short run should not persist into the long run. In Phillips curve terms, this leads to a familiar conclusion: the Phillips curve is negatively sloped in the short run, but vertical in the long run. ADP, by contrast, offer figure 1 as their long-run Phillips curve.

These two views carry quite different policy implications. In the ADP view, money illusion will last forever as long as inflation is modest. Thus the leftmost point in their figure 1, with its super-low unemployment rate, can persist indefinitely. In the post-Friedman-Phelps Keynesian view, however, money illusion will not last forever.

ADP use a variety of psychological evidence (including a quote from me—dirty pool!) to argue instead that higher inflation, not the passage of time, dissipates money illusion. I have not the slightest objection to their psychologically based notion that higher inflation rates are more salient, are less likely to be edited out of the worry lists of harried decisionmakers, and so forth. But doesn't time matter, too? As long as the economy remains in the range where money illusion is operational, real wages in the near-rational sector remain lower than in the rational sector.

5. The ADP model contains no condition analogous to the neoclassical equilibrium: real wage = marginal product of labor.

6. An exception arises if labor supply is totally inelastic. In that case, employment is unaffected.

Eventually, workers should catch on to the fact that their neighbors are earning more, even if only a little more.

The evidence for money illusion. Of course, my pretending to be a Chicago economist was just playacting. In fact, I am persuadable—indeed, pretty much persuaded—that money illusion is a fact of life. But I do want to enter a couple of objections to the evidence that ADP offer in its defense.

They suggest a “decision heuristic” whereby inflation is “averaged along with other factors to arrive at a nominal wage or price increase.” Maybe I am trapped inside a rationality box, but this does not make sense to me. What “other factors” are in this average? Productivity improvements? Demand pressures? As determinants of nominal wage increments, shouldn’t these factors be *added*, not averaged? It seems to me that inflation is a factor unlike any other. Because it is fundamentally about adjusting units of measurement, it is something one *adds on* to the basic determinants of real wages. Besides, if we are averaging factors that contribute to higher wages, shouldn’t the other factors, like productivity, also be attenuated?

Second, the authors observe that compensation professionals caution firms not to grant automatic cost-of-living adjustments. That may be true, but it is not evidence for money illusion. Suppose the firm raises its salary pool by inflation plus a small increment for productivity gains. That would be an illusion-free outcome. Still, the firm would not want to give every worker the same percentage raise, because that would destroy its ability to reward superior performance—that is, to adjust *relative* wages.

Finally, let me suggest a potential cognitive error that cuts in the other direction. The Bureau of Labor Statistics adjusts various prices for quality improvements, and for good reasons. But ordinary consumers, being less sophisticated than Jack Triplett or Robert Gordon, may see higher prices as higher prices. So, for example, most people think automobile prices have soared over the last two decades. But the BLS disagrees: after accounting for large quality adjustments, new car prices have risen much slower than the overall rate of inflation. And everyone knows that computers have fallen in price. But how many people think they now cost only 6 percent of what they did in 1981? Yet that is what Dale Jorgenson and Kevin Stiroh imply in their paper in this issue. As a final example, is there anyone who doesn’t believe that medical care prices have skyrocketed, even though some excellent recent research suggests that many quality-

adjusted prices have actually fallen? Mere mortals who have this kind of money illusion will overweight inflation, not underweight it.

The last word. I hasten to add that I find much to admire in this paper, especially the authors' thorough and creative empirical work, which includes myriad robustness checks. I was especially pleased to find the Sargent critique addressed by the use of directly observed expectational data. Such an approach is no panacea, but it is certainly a step in the right direction.

Finally, even if one harbors doubts about the details of the ADP model, as I do, the theory calls attention to an interesting hypothesis that is well worth exploring in its own right, namely, that the coefficient on expected inflation in the Phillips curve depends on past inflation. The authors could have skipped all the theory, appealed to the psychological evidence, and jumped directly to that testable hypothesis. To test it, however, one needs a way to constrain the coefficient to range between zero and one. Using a probability distribution function is a clever way to accomplish that. And so far, the tests look pretty favorable. Now that, *kemo sabe*, is something the Fed really should know.

General discussion: Benjamin Friedman expanded on Alan Blinder's comment about whether the point that maximized the sustainable employment rate should be considered the optimal point at which to operate the economy. He pointed out that, even if one accepts the standard macroeconomic argument that zero inflation minimizes various distortion costs and that the losses from these costs are quadratic, then given the higher employment that a positive inflation rate makes possible in the authors' model, the optimal inflation rate is above zero even after accounting for the associated distortions. Friedman also noted that other recent empirical work had cast doubt on the idea of a vertical long-run Phillips curve, and he saw the paper as both confirming those results and providing a theoretical underpinning for them. He observed that the paper is consistent with the recent literature that finds, for the United States and other OECD countries, that many aspects of the real economy are not invariant to the inflation rate.

Edmund Phelps interpreted the paper as showing that declining inflation accounted for the decline in U.S. unemployment. He noted that he and

his coauthors, in their study in this volume, had found that real variables, such as real interest rates and changing productivity trends, were important in explaining the divergence in unemployment rates across OECD countries in the 1990s. Once these factors were accounted for, inflation was not important. He wondered whether the model of this paper would hold up if real variables were given a chance to explain the data. William Dickens responded that the Brainard-Perry results suggest that it is the change in the coefficient on inflation that is driving changes in the sustainable rate of unemployment, not changes in the intercept, as would be the case if the factors Phelps cited were responsible for the improvement.

Robert Gordon recalled James Stock's observation that the real puzzle in today's economy might be low unemployment rather than low inflation. This perspective draws attention to the fact that some alternative measures of resource utilization, such as capacity utilization in manufacturing, are not tight. It also draws attention to a number of structural factors relating to the labor market, such as those examined by Lawrence Katz and Alan Krueger in the Spring 1999 issue of *Brookings Papers*, which the present authors did not consider. Responding to Phelps's and Gordon's observations that real variables and labor market changes had not been taken into account, George Perry noted that productivity was in the equations and that demographic changes were allowed for by using alternative unemployment measures.

William Nordhaus suggested that the issue of the costs of inflation went beyond tax distortions and shoe leather costs, and involved how people misperceive inflation. He recalled a point made by Henry Wallich at a time when the division of labor in most households fell along then-traditional lines. Wallich had observed that inflation was sexist: the husband liked it because his wages were rising rapidly, but the wife disliked it because the prices she had to pay at the grocery store were rising rapidly. The persistent irrationality in the conflicting views of economic agents as producers and the same economic agents as consumers belongs in any analysis of the cost of inflation. Nordhaus also suggested that the authors test whether inflationary expectations in fact have responded to experienced inflation in the nonlinear way suggested by the paper. Such a response would imply that the elasticity of inflationary expectations to inflation was low at low inflation rates, then rose above one as inflation rose, and settled at one at sufficiently high inflation rates. While recognizing that the model underpinning this paper had the response to expectations, rather than the expect-

tations themselves, varying with inflation, he reasoned that the alternative model was worth examining.

Olivier Blanchard observed that both the changing prevalence of cost-of-living clauses in wage contracts and the characteristics of such clauses supported the authors' model of how inflation entered into decisions. Such clauses typically kick in only above a certain threshold, with no adjustment made at low inflation rates, and the fraction of contracts that are indexed varies directly with the inflation rate. During the high-inflation period, two explanations were offered for this changing fraction. One was that certain fixed costs associated with understanding the point of indexation had to be overcome, and the other was that incentives for indexation are high only when inflation is high. The fact that the proportion of indexed contracts declined after inflation fell supports the second explanation, which is consistent with the model.

Despite being skeptical of the econometrics, Christopher Sims found the model and its implications plausible. His own work led him to believe that agents react in nonlinear fashion to shocks, with small shocks eliciting weak reactions and large shocks eliciting strong ones. He had modeled the idea that the reactions of individuals to stochastic signals, like prices or wages, involve a finite rate of information flow, using this as a constraint on dynamic optimizing behavior. This implies that individuals react sluggishly to market signals, with their reactions delayed and smoothed, or made at discrete intervals, introducing idiosyncratic error. The extent of such processing depends on how important the signal is.

Sims suggested that this work provides a further theoretical justification for the kind of results the authors obtain, but in addition provides a rationale for a broader pattern of results that emerges in the literature on vector autoregressive analyses. Its findings contradict the notion that some variables are inherently costly to move. Shocks typically affect own variables quickly, whereas cross-variable relationships show slow, smooth reactions. In closing, Sims suggested that this work implies a modification of the Lucas critique: perhaps what matters is not whether a policy action is predictable, but how important it is. This would make it possible to have apparently permanent sluggish responses when policy is stable and shocks are small.

References

- Akerlof, George A., William T. Dickens, and George L. Perry. 1996. "The Macroeconomics of Low Inflation." *BPEA*, 1:1996, 1–76.
- Akerlof, George A., and Janet L. Yellen. 1985. "A Near-Rational Model of the Business Cycle, with Wage and Price Inertia." *Quarterly Journal of Economics* 100 (Supplement): 823–38.
- Ball, Laurence, N. Gregory Mankiw, and David Romer. 1988. "The New Keynesian Economics and the Output-Inflation Tradeoff." *BPEA*, 1:1988, 1–82.
- Blanchard, Olivier J. 1999. *Macroeconomics: Second Edition*. Upper Saddle River, N.J.: Prentice Hall.
- Blinder, Alan S., and others. 1998. *Asking About Prices: A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.
- Brainard, William C., and George L. Perry. 2000. "Making Policy in a Changing World." In *Economic Events, Ideas, and Policies: The 1960s and After*, edited by James Tobin and George Perry. Brookings (forthcoming).
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. MIT Press.
- Eckstein, Otto, and Roger Brinner. 1972. "The Inflation Process in the United States." Joint Economic Committee of the Congress of the United States. Government Printing Office.
- Fair, Ray C. 2000. "Testing the NAIRU Model for the United States." *Review of Economics and Statistics* 82(1): 64–71.
- Fisher, Irving. 1930. *The Theory of Interest*. New York: Macmillan.
- Friedman, Milton. 1968. "The Role of Monetary Policy." *American Economic Review* 58(1): 1–17.
- Gleitman, Henry. 1996. *Basic Psychology*. Norton.
- Gordon, Robert J. 1998. "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *BPEA*, 2:1998, 297–346.
- Gray, Jo Anna. 1978. "On Indexation and Contract Length." *Journal of Political Economy* 86(1): 1–18.
- Greenspan, Alan. 1988. Statement before the Committee on Banking, Housing and Urban Affairs of the U.S. Senate, July 13. *Federal Reserve Bulletin* (September): 607–13.
- Hendricks, Wallace E., and Lawrence M. Kahn. 1985. *Wage Indexation in the United States: Cola or Uncola®?* Ballinger.
- Holland, A. Steven. 1995. "Inflation and Wage Indexation in the Postwar United States." *Review of Economics and Statistics* 77(1): 172–76.
- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2): 263–92.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest, and Money*. Harcourt, Brace.

- King, Robert G., and Mark W. Watson. 1994. "The Post-War U.S. Phillips Curve: A Revisionist Econometric History." *Carnegie-Rochester Conference Series on Public Policy* 41: 157–219.
- Krueger, Anne O., and others. 1991. "Report of the Commission on Graduate Education in Economics." *Journal of Economic Literature* 29(3): 1035–53.
- Kunreuther, Howard. 1978. *Disaster Insurance Protection: Public Policy Lessons*. Wiley.
- Leonard, Jonathan S. 1987. "In the Wrong Place at the Wrong Time: The Extent of Frictional Unemployment." In *Unemployment and the Structure of Labor Markets*, edited by Kevin J. Lang and Jonathan S. Leonard. New York: Basil Blackwell.
- Levine, David I. 1993. "Fairness, Markets, and Ability to Pay: Evidence from Compensation Executives." *American Economic Review* 83(5): 1241–59.
- Lipsey, Richard. 1960. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1862–1957: A Further Analysis." *Economica* 27(February): 1–31.
- Lucas, Robert E., Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4(April): 103–24.
- Mankiw, N. Gregory. 1985. "Small Menu Costs and Large Business Cycles: A Macroeconomic Model of Monopoly." *Quarterly Journal of Economics* 100(2): 529–37.
- Milkovich, George T., and Jerry M. Newman. 1984. *Compensation*. Plano, Tex.: Business Publications.
- Modigliani, Franco, and Lucas Papademos. 1975. "Targets for Monetary Policy in the Coming Year." *BPEA*, 1:1975, 141–65.
- Nisbett, Richard, and Lee Ross. 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*. Prentice-Hall.
- Phelps, Edmund. 1967. "Phillips Curves, Expectations of Inflation and Optimal Unemployment Over Time." *Economica* 34(August): 254–81.
- . 1968. "Money-Wage Dynamics and Labor-Market Equilibrium." *Journal of Political Economy* 76(4, Part 2): 678–711.
- Phillips, A. W. 1958. "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957." *Economica* 25(November): 283–99.
- Rock, Milton L., and Lance A. Berger. 1991. *The Compensation Handbook: A State-of-the-Art Guide to Compensation Strategy and Design*, 3rd ed. McGraw-Hill.
- Sargent, Thomas J. 1971. "A Note on the 'Accelerationist' Controversy." *Journal of Money, Credit, and Banking* 3(3): 721–25.
- . 1973. "Rational Expectations, the Real Rate of Interest, and the Natural Rate of Unemployment." *BPEA*, 2:1973, 429–72.
- Shafir, Eldar, Peter Diamond, and Amos Tversky. 1997. "Money Illusion." *Quarterly Journal of Economics* 112(2): 341–74.

- Shiller, Robert J. 1997. "Why Do People Dislike Inflation?" In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer. University of Chicago Press.
- Shimer, Robert J. 1998. "Why Is the U.S. Unemployment Rate So Much Lower?" In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Julio Rotemberg. Cambridge, Mass.: National Bureau of Economic Research.
- Stock, James H., and Mark W. Watson. 1998. "Median Unbiased Estimation of Coefficient Variance in a Time-Varying Parameter Model." *Journal of the American Statistical Association* 93(March): 349–58.