Inflation Expectations and the Transmission of Monetary Policy

by
John M. Roberts
Board of Governors of the Federal Reserve System
Stop 80
Washington, D.C. 20551

jroberts@frb.gov

October 1998

Abstract: New Keynesian models with sticky prices and rational expectations have a difficult time explaining why reducing inflation usually requires a recession. An explanation for the costliness of reducing inflation is that inflation expectations are less than perfectly rational. To explore this possibility, I estimate the degree of nonrationality in two survey measures of inflation expectations. I find that the surveys reflect an intermediate degree of rationality: Expectations are nether perfectly rational nor as unsophisticated as simple autoregressive models would suggest. I also find that a structural New Keynesian model with expectations formation based on the survey results is able to match closely the empirical costs of reducing inflation.

JEL classification numbers: E31, E52

I am grateful to Larry Ball, Dean Croushore, Jeff Fuhrer, Robert King, Athanasios Orphanides, John Williams, and Michael Woodford, and seminar participants at the University of Virginia, Johns Hopkins University, the Federal Reserve System Committee on Macroeconomics, and the 1998 NBER Summer Institute for helpful comments. The views expressed in this paper are my own and are not necessarily shared by the Board of Governors of the Federal Reserve System or any other members of its staff.

New Keynesian economics strikes the following bargain: By accepting limited deviations from an otherwise standard new classical framework, the model can explain important empirical phenomena. In particular, by adding sticky prices to a model with otherwise standard microeconomic underpinnings (such as upward sloping supply curves, downward sloping demand curves, and, especially, rational expectations), New Keynesian economics can explain why monetary policy affects real variables.

Unfortunately, it turns out that there are empirical phenomena that the New Keynesian sticky-price framework with rational expectations can't explain. In particular, the model has difficulty accounting for the output losses that typically accompany a reduction in inflation. This point is easy to see by examination of the New Keynesian Phillips curve with sticky prices:

$$\Delta p(t) = E(t)\ \Delta p(t+1) + \gamma\ y(t) + \epsilon(t)\ , \qquad (1)$$

where $\Delta p(t)$ is inflation, $E(t)\ \Delta p(t+1)$ is the expectation in period t of inflation in the next period, and y is the deviation of output from its trend level. The problem for this model under rational expectations is that if an announcement by a central bank of a desire to reduce inflation can lead to a reduction in inflation expectations, then inflation can drop without the need for any drop in output.

Conventional estimates suggest that reducing inflation requires a loss of output. For example, Alan Blinder (1987, pp. 38-39) put the cost of a 1 percentage point permanent reduction in inflation at an increase in the unemployment rate of 2 percentage points for one year. And Larry Ball (1994) has documented that the employment and output losses associated with reducing inflation appear to be widespread across countries and over time.

Apparently, some modification of the New Keynesian sticky-price model is needed to explain why it is that inflation is so costly to reduce, and perhaps, in addition to the unattractive assumption of sticky prices, some other unattractive assumption may be necessary. In the past few years, two such unattractive assumptions have received attention. One, emphasized by Fuhrer Moore (1992, 1995), is to assume a deviation from standard microeconomics: Instead of assuming that labor supply constitutes a trade-off between employment and the *level* of the real wage, Fuhrer and Moore propose that labor supply represents a trade-off between employment and the *change* in the real wage.

An alternative option is to relax the assumption of perfectly rational expectations (Ball, 1991; Roberts, 1997). The unattractiveness of dropping this assumption is clear: Rational expectations has been one of the bedrock assumptions of macroeconomics for more than twenty years. Because it is the assumption about expectations that is most compatible with standard optimizing behavior, it should not be dropped lightly.

One piece of evidence supporting less-than-rational expectations is that surveys of inflation expectations for a year or more ahead violate the usual tests of rationality; this literature is reviewed in section 1 below. But economists are often skeptical that survey expectations tell us anything about actual inflation expectations, since respondents have little incentive to accurately report their expectations to a survey taker. So, survey expectations are at best a noisy measure of inflation expectations and at worst, worthless. However, in earlier work, reviewed in section 2 below, I have found that survey expectations prove to be quite "useful" in the estimation of New Keynesian aggregate supply models (Roberts, 1995, 1997). In particular, I've found that the estimates of the New Keynesian model are plausible conditional on survey expectations being an accurate measure of expectations, and that, given the choice, the data prefer the surveys as measures of expectations. Furthermore, I've also found that, conditional on the surveys being accurate measures of expectations, there is no need to resort to the nonstandard microeconomics of Fuhrer and Moore (Roberts, 1997). By contrast, Fuhrer and Moore found that, conditional on expectations being rational, they needed nonstandard microeconomics to fit U.S. data.

If less-than-rational expectations are to form part of a positive theory of the transmission of monetary policy, we need to move beyond econometric efficiency tests toward explicit formulations of how expectations deviate from rationality. I consider two specific alternative hypotheses about inflation expectations. One is that a fraction of the population has inflation expectations based on simple extrapolations of past inflation. The other is that inflation expectations adjust only gradually to the fully rational value. Neither of these alternatives is based on optimizing models. However, they may prove to be useful in characterizing the way in which expectations deviate from rationality. I introduce these models in section 2, discuss some alternative models in section 3, and derive empirical models in section 4.

I present estimates in section 5. I find that both models of less-than-rational expectations can account for the deviation of expectations from rationality reflected in two of the main surveys of inflation expectations. The point estimates suggest an intermediate degree of rationality: I can reject both perfect rationality and high levels of nonrationality. In terms of the model of partially adaptive expectations, the results can be interpreted as implying that perhaps 20 to 40 percent of the population has adaptive expectations. The results for the model of partial adjustment of expectations imply that expectations move 50 to 60 percent of the way toward the rational outcome per year. However, I also find that estimates of the partial adjustment parameter fail tests of stability over time, whereas estimates of the degree of adaptive expectations are more stable over time. These results suggest that the partly adaptive model may be the more robust.

In section 6, I address whether the degree of rationality implicit in the surveys can account for the costliness of reducing inflation. As a basis for comparison, I use a vector autoregression that includes inflation, a measure of output, and interest rates as my basis of comparison. I replace the equation for inflation in that model with a New Keynesian Phillips curve, assuming less-than-rational expectations, and then solve for the reduced form model of inflation consistent with that model. I find that with the intermediate degree of nonrationality I estimated from the survey data, the model does a good job of capturing the responses of the economy. In particular, it predicts a substantial output loss in response to a disinflationary monetary shock. By contrast, under the assumption of rational expectations, the model predicts very little cost to reducing inflation.

The results suggest that the degree of rationality represented by surveys of inflation expectations can explain the empirical costs of reducing inflation. These results bolster the case that less-than-rational expectations play a role in making inflation costly to reduce.

1. The literature on the rationality of surveys of inflation expectations

1.1 Survey expectations are not purely adaptive

There is a large literature on surveys of inflation expectations. One result from this literature is that survey measures of inflation expectations are not purely adaptive--that is, they reflect more information than that embedded in lagged inflation. Mullineaux (1980) and Gramlich (1983), for example, both found that the money supply helps explain movements in survey expectations, even controlling for lagged inflation.

1.2 But survey expectations aren't purely rational, either

Most researchers have concluded that survey expectations are not perfectly rational. For example, Bryan and Gavin (1986) and Batchelor and Dua (1989) find that the Livingston survey of economists' predictions is a biased forecast of inflation, while Pesaran (1987) finds that a survey of British manufacturers does not make efficient use of available information.[1] Some studies have suggested that for one measure of inflation expectations--the Michigan survey of household expectations--is an unbiased predictor of

---

[1] These results are for expectations of inflation one year ahead. Studies that have examined short-term expectations — of a month to a quarter ahead — have found that expectations are rational (Pearce, 1987; Keane and Runkle, 1990). Rationality one year ahead is obviously a higher hurdle to clear, but a failure of rationality over a one-year period is relevant in the New Keynesian staggered-contracts models, since the contracts may last a year or more.

inflation. However, it does not pass a stricter test of rationality, that a forecast make efficient use of all available information (Batchelor and Dua, 1989; Baghestani, 1992; Ball and Croushore, 1995; Roberts, 1997).

2. Sticky prices, sticky inflation, and "not-quite-rational expectations"

Fuhrer and Moore (1992, 1995) have emphasized the following model, originally proposed by Buiter and Jewitt (1981):

$$\Delta p(t) = [M(t)\ \Delta p(t+1) + \Delta p(t-1)]/2 + \gamma\ y(t) + \epsilon(t), \qquad (3)$$

where I use "M" to distinguish rational or "mathematical" expectations from other possible expectation-formation mechanisms. One way of thinking of equation 3 is that it is a model of "sticky inflation" rather than of sticky prices, since equation 3 can be derived by replacing prices with inflation in the sticky price model of equation 1. Costless disinflation is more difficult in this model, since lagged inflation can not move at the same time as inflation and expected inflation. Fuhrer and Moore (1995) found that this model was successful empirically: Under the assumption of rational expectations, they found that the sticky inflation model fit U.S. data well, whereas the standard New Keynesian sticky price in equation 1 did not.

They interpreted their results as consistent with an underlying microeconomic model that, in addition to sticky prices, assumed a further deviation from conventional microeconomic foundations. They called their alternative specification "the relative real wage hypothesis"; it amounted to an assumption that when workers set their wages, they do not simply attempt to set the level of their nominal wages high relative to the expected price level over the life of their labor contract when employment is high; they instead are concerned about having a large *change* in their nominal wage relative to inflation when employment is high. Hence, the Fuhrer and Moore model "slips a derivative" relative to the conventional microeconomics.

One of the attractions of the New Keynesian Phillips curve is that, aside from sticky prices, its microeconomic foundations are conventional: That is, it can be derived assuming conventional supply and demand relationships, in which, for example, workers attempt to set nominal wages, which may be fixed for several periods, so that they will be high in real terms in periods when employment is expected to be high. Fuhrer and Moore's interpretation of their results requires abandoning this attractive aspect of the New Keynesian model.

However, an alternative interpretation of Fuhrer and Moore's aggregate results is that inflation expectations are less than perfectly rational. That is, suppose that:

$$E(t)\ \Delta p(t+1) = [M(t)\ \Delta p(t+1) + \Delta p(t-1)]/2. \qquad (4)$$

Equation 4 says that inflation expectations are not the same as perfectly rational expectations. Rather, they are only partly rational; they are also partly adaptive. A justification for this specification could be that half of the population has rational expectations and the other half has adaptive expectations. Algebraically, if equation 4 is substituted into equation 1, we obtain equation 3. So, under the assumption of rational expectations, the

To determine whether imperfect rationality or "sticky inflation" was the better explanation of U.S. inflation dynamics, I have tested whether Fuhrer and Moore's deviation from the standard sticky-price model was needed if we assumed that survey expectations were a good proxy for inflation expectations (Roberts, 1997). I found that conditional on the surveys being a good measure of expectations, the sticky-price model fit better than the sticky-inflation model, suggesting that less-than-rational expectations were preferred to sticky inflation.

In addition to partly adaptive expectations, another specification of expectations formation is that expectations are "stubborn," and adjust only gradually to their "rational" value. This specification suggests that expectations follow a partial adjustment process that can be written as:

$$S(t)\ \Delta p(t+1) = \rho\ S(t-1)\ \Delta p(t) + (1 - \rho)\ M(t)\ \Delta p(t+1)\ . \qquad (5)$$

In this model, inflation expectations are a weighted average of what they were last period and what they "should" be.

"Stubborn" expectations may be a more attractive specification than the partly adaptive model. One interpretation is that this specification could be thought of as a model of "habit persistence" in inflation expectations. Another interpretation is that a group that might adjust expectations in this way is professional forecasters. Professional forecasters might be hesitant to adjust their forecasts too rapidly in response to new information —perhaps out of concern for looking foolish for making a large change in their forecast. Alternatively, professional forecasters might not want to differ too much from the consensus of other forecasters. If that's the case, then they will be always looking backward at the published forecasters of others in formulating their own forecasts. Finally, the forecasts of other "agents," such as households, may be influenced by press reports of professional forecasts.

The implications of "stubborn" expectations for estimates such as Fuhrer and Moore's can be seen by substituting equation 5 into equation 1 and rearranging:

$$\Delta p(t) = \rho\ \Delta p(t\text{-}1) + (1 - \rho)\ M(t)\ \Delta p(t+1) + \gamma\ [y(t) - \rho\ y(t\text{-}1)] \\ + \epsilon(t) - \rho\ \epsilon(t\text{-}1)\ . \tag{6}$$

Clearly, the implications for the presence of lagged inflation are similar to equation 3.

3. Some related models

a. Credibility

Ball (1991, 1995) has argued that if central banks lack "credibility," inflation can be costly to reduce even if expectations are rational. Ball's argument requires that monetary policy behaves in such a way that people are legitimately confused as to what the central bank's objectives are. However, the results presented in section 6 and in Fuhrer and Moore (1995) show that using an estimated reaction function, the sticky-price model with rational expectations fails to predict large costs of reducing inflation. Since an estimated reaction function would capture any historical tendency for the central bank to send confusing signals, these simulation results can be interpreted as evidence against Ball's hypothesis.

Furthermore, Ball's (1994) evidence that costly disinflation is widespread suggests that misdirection is unlikely to be the whole story: While one central bank may have a history of misleading policy moves, it seems unlikely that they all would. And there is little evidence that efforts by central banks to bolster their credibility by announcing their objectives for monetary policy have done much good in reducing the costs of disinflation. Countries such as Canada, New Zealand, and Great Britain have announced plans to disinflate, usually several years ahead of time. The disinflation occurred according to plan — but along the way, these countries suffered significant recessions. While not ironclad proof, this evidence is also suggestive that credibility is not the whole story.

b. Structural breaks

Many authors have pointed out that rational learning following a structural break can lead to serially correlated forecast errors (Caskey, 1985; Lewis, 1989). If a sample is dominated by the aftermath of such a structural break, rationality tests of the kind I have used here will incorrectly suggest that rationality is rejected. Rational learning after structural changes can also be an explanation for costly disinflation. The serially correlated inflation forecast errors following the structural break will lead to serially correlated deviations of output from its trend level in a similar way as the serially correlated deviations of expectations from rationality implied by less than rational expectations.

If we could be sure that the macroeconomic environment has been stable over the sample period, our confidence in the results would be increased. One way to guard against the effects of a single structural break is to look at subsample results. If the results are similar in the two halves of the sample, our concern that a structural break was responsible for the rejection of rationality should be reduced. I explore this possibility in the empirical work of section 5.

As mentioned in the introduction, Ball (1994) found that costly disinflation was widespread both over time and across countries; in the United States, costly disinflation is almost as common as recession. If structural breaks are responsible for both costly disinflation and the rejection of the rationality of inflation expectations, such breaks must therefore have been quite common.

But frequent structural breaks present a formidable econometric challenge, since standard econometric techniques assume a stable environment. If the breaks are frequent enough, it will be difficult to find a subsample of any important length that is free of them. Hence, there is a sense in which the hypothesis of frequent structural breaks is nonrefutable. In any case, frequent structural breaks suggests an alternative interpretation of my results.[2]

4. Deriving the estimation framework

In section 2, I discussed some stylized models of how inflation expectations might be formed. In this section, I derive some implications of these models for empirical testing.

I should reiterate that these "models" are only empirical generalities and do not correspond to structural models derived from underlying economic behavior. While it would be desirable to have a structural model of the deviation of inflation expectations from rationality, it is nonetheless a useful first step to move from the broad statement that expectations are not rational to a more specific statement about how expectations deviate from rationality.

A generalization of the partly adaptive expectations hypothesis is:

$$S(t)[p(t+1) - p(t)] = \alpha(L)\, \Delta p(t-1) + [1 - \alpha(1)]\, M(t)[p(t+1) - p(t)]\ , \qquad (7)$$

where $M(t)[p(t+1) - p(t)]$ is the purely rational, or "mathematical," expectation over that period. The notation $\alpha(L)$ indicates a lag polynomial, while $\alpha(1)$ indicates the sum of the coefficients in the lag polynomial.

Equation 7 can be interpreted as indicating that inflation expectations are part-way between being perfectly rational and purely "adaptive." One could imagine that all individuals have such partly adaptive expectations, or that a fraction $1 - \alpha(1)$ of the population has rational expectations, while the remainder have adaptive expectations. Of course, if $\alpha(1) = 0$, then expectations are perfectly rational, whereas if $\alpha(1) > 0$, then expectations are less than rational.

---

2. Structural breaks are formally distinct from the policy regime shifts in, for example, Markov-type regime-shifting models. In a long-enough sample, linear econometric efficiency tests will be robust to the nonlinearities introduced by a Markov-type model. Of course, a small sample that is dominated by a single observation of a regime shift will be indistinguishable from a sample with a structural break.

Moving toward a specification that can be estimated, it is useful to rewrite equation 7 as:

$$M(t)[p(t+1) - p(t)] = (1/[1 - \alpha(1)])\{S(t)[p(t+1) - p(t)] - \alpha(L)\,\Delta p(t-1)\}. \quad (8)$$

We can exploit the properties of rational expectations by substituting in realized inflation for the expectation, which introduces an error term:

$$p(t+1) - p(t) = (1/[1 - \alpha(1)])\{S(t)[p(t+1) - p(t)] - \alpha(L)\,\Delta p(t-1)\} + \eta(t) \quad (9)$$

If M(t)[p(t+1) - p(t)] in equation 8 is truly rational, then the error term η will be uncorrelated any other information from period t or earlier. Equation 9 can be further rewritten to isolate the forecast error as the dependent variable:

$$p(t+1) - p(t) - S(t)[p(t+1) - p(t)] =$$
$$-(1/[1 - \alpha(1)])\{\alpha(L)\,\Delta p(t-1) - \alpha(1)\,S(t)[p(t+1) - p(t)]\} + \eta(t) \quad (10)$$

The stubborn expectations model can be written as:

$$S(t)[p(t+1) - p(t)] = \rho\, S(t-1)[p(t) - p(t-1)] + (1 - \rho)\, M(t)[p(t+1) - p(t)]\,. \quad (11)$$

Again, inflation expectations are a weighted average of what they were last period and what they "should" be. An interpretation of this specification is that individuals are stubborn, and so adjust gradually to the rational level. It could also be thought of as a model of "habit persistence" in inflation expectations.

Again thinking ahead to estimation, it is useful to rewrite the stubborn expectations model as:

$$p(t+1) - p(t) - S(t)[p(t+1) - p(t)] =$$
$$[\rho/(1-\rho)]\ \{S(t)[p(t+1) - p(t) - S(t-1)[p(t) - p(t-1)]\} + \eta(t)\,, \quad (12)$$

where $\eta(t)$ is an expectational error.

In estimation, there is a risk of spurious correlation between S(t)[p(t+1) - p(t)] on the left- and right-hand sides of the equation. I will avoid this problem will be to examine instrumental variables estimates.

5. Structural estimation

a. Data

I examine two surveys of inflation expectations that have been conducted for an extended period of time. One is the "Livingston" survey of economists' forecasts of inflation, which was started by a journalist of that name in the 1940s. The survey is conducted every six months. The panel of economists (currently fifty-five) is asked to report their forecast for the consumer price index. Because of publication lags, the usual lead time of the forecast is fourteen months. The survey is currently conducted by the Federal Reserve Bank of Philadelphia and the data are available from their web site.

The second survey I examine is part of the University of Michigan Survey Research Center's ongoing program of assessing household attitudes. A random sample of households — currently about five hundred — is asked what they expect the inflation rate will be for things they buy over the next twelve months. From May 1966 to the end of 1977, the survey was quarterly, and was taken in the middle month of each quarter; it has since been monthly. To have a longer series, I use quarterly data, using the observation for the middle month of the quarter for the more recent period.[3] The Michigan survey is available from the Survey Research Center on a proprietary basis.[4]

The basis of comparison for the forecasts is the consumer price index. The CPI is explicitly the measure of prices the Livingston respondents are trying to forecast, and it is conceptually close to the measure the Michigan survey is aiming at.

Structural estimation requires the use of instrumental variables, for several reasons. One is that the variables on the right-hand sides of both equations 10 and 12 include the current-period expectation of inflation, which also appears on the left-hand side of the equation. Another reason is measurement error, which may be an especially important problem for survey expectations, since they are based on small samples, and so may be very noisy indicators. Price data, however, are also likely to be measured with error. And

---

3. From 1948 to 1965, the Michigan survey also asked about inflation, but the question was qualitative: Do you expect prices to rise or to fall over the next year?

4. A third survey, the Survey of Professional Forecasters (also conducted by the Federal Reserve Bank of Philadelphia) is available for almost as long as the Michigan survey. However, the inflation question in this survey refers to the GDP deflator. Over the years, the GDP deflator has been periodically redefined. As a consequence, it is not possible to use current data to assess the forecast performance of the SPF.

in the case of the price data, measurement error may be serially correlated (see Roberts, 1994).

I use as instruments changes in the federal funds rate and two measures of economic activity, the unemployment rate and manufacturing capacity utilization. I do not use lagged inflation as an instrument because of the possibility of serially correlated measurement error mentioned above. For the interest rate, I use the federal funds rate. The fed funds rate has the advantage of being the interest rate most directly related to monetary policy, since it is the interest rate on overnight loans of reserves among banks. Hence, it can be thought of as providing a signal of monetary policy. I use the change in the funds rate to allow for nonstationarity in interest rates. Ball and Croushore (1995) and Roberts (1997) find that the change in the federal funds rate is strongly correlated with the inflation expectations survey forecast errors.

Capacity utilization and the unemployment rate have different relative strengths as measures of economic activity. Both are strongly cyclical and are highly correlated with other measures of activity such as gross domestic product. The chief advantage of capacity utilization is that it is unambiguously stationary. As a consequence, I use capacity utilization in the second part of this paper, when I examine vector autoregression models. However, capacity utilization has the disadvantage that it is revised. So it is possible information that was not available at the time households and economists formed their inflation expectations has been incorporated into the currently available data. By contrast, the unemployment rate is not revised.

I use four-quarter averages of the two measures of economic activity and the four-quarter change in the federal funds rate (which is equivalent to a four-quarter average of the one-quarter change); I use three nonoverlapping lags of each regressor; and the four-quarter averages are for periods t-1 to t-4. I use nonoverlapping four-quarter averages of the data to balance the need to include sufficient lagged information without introducing an excessive number of instruments. In small samples, excessive instruments can lead to biased results (intuitively, that's because overfitting will allow the OLS results to manifest themselves). However, preliminary estimation with annual data suggested that several lags were informative. Finally, I lag the information set one quarter to allow for publication and other high-frequency timing lags.

The serial correlation that results from overlapping forecasts introduces econometric complications for structural estimation: The serial correlation in the errors must be taken account of when making statistical inferences about the estimates. I therefore use the GMM procedure, allowing for up to eighth order serial correlation, with the Newey-West adjustment to the weighting matrix.

b. Results

I first discuss the estimates of the model using the unemployment rate as the measure of economic activity and then turn to the estimates using capacity utilization.

Table 1 shows the estimates of the two models of inflation expectations formation using the Livingston survey. By the criterion of the orthogonality of the residuals to the instruments, both models successfully explain the correlation of the instruments with the survey forecast errors. The estimated parameter of the partial adjustment model, shown in the top panel, is significantly different from zero at a high degree of confidence. It implies that expectations adjust about half way to the rational value each year.

The estimates for the partly adaptive expectations model assume that inflation expectations are partly based on a moving average of past inflation. I look at a four-quarter and an eight-quarter moving-average. With the four-quarter moving average, the underlying coefficient is statistically significant at the 6 percent level. The results imply that 36 percent of the population has adaptive expectations. With the eight-quarter moving average, however, the underlying coefficient is not significantly different from zero; the point estimate implies that just 21 percent of the population has adaptive expectations.

Table 2 shows results for the Michigan survey. Again, both structural models are able to account for the correlation of the instruments with the survey forecast error, the dependent variable. The underlying parameter for the partial adjustment model is once again precisely estimated, and implies that

Table 1
Structural Estimation, Models of Less-Than-Rational Expectations
Livingston Survey
Semiannual, 1957:2-1995:2

Instruments: three nonoverlapping lags of the four-quarter change in the federal funds rate, offset one quarter; and three nonoverlapping lags of the four-quarter average of the civilian unemployment rate, offset one quarter.

A. Partial adjustment of expectations

	$\rho/(1+\rho)$	ρ	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
	.867 (.294) [.003]	.46	.64	.25

B. Partly adaptive expectations

Partly adaptive portion of expectations	$-\alpha/(1-\alpha)$	α	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
Four-quarter moving average [p(t-1) - p(t-5)]	-.557 (.297) [.06]	.36	.66	.55
Eight-quarter moving average [p(t-1) - p(t-9)]/2	-.209 (.236) [.38]	.21	.65	.53

Numbers in parentheses are standard errors.
Numbers in brackets are p-values.

Table 2
Structural Estimation, Models of Less-Than-Rational Expectations
Michigan Survey
Quarterly, 1966:2-1995:4

Instruments: Same as table 1.

A. Partial adjustment of expectations

	$\rho/(1+\rho)$	ρ	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
	.643 (.151) [.0002]	.39	.40	.33

B. Partly adaptive expectations

Partly adaptive portion of expectations	$-\alpha/(1-\alpha)$	α	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
Four-quarter moving average [p(t-1) - p(t-5)]	-.289 (.192) [.13]	.22	.32	.36
Eight-quarter moving average [p(t-1) - p(t-9)]/2	-.342 (.126) [.007]	.25	.34	.54

Numbers in parentheses are standard errors.
Numbers in brackets are p-values.

expectations adjust about 60 percent of the way to the fully rational value each year.

With the Michigan survey, the partly adaptive expectations model yields more precise estimates when an eight-quarter moving average is the proxy for the adaptive portion of inflation expectations. In that case, the underlying coefficient is significant at the 1 percent level, and implies that 25 percent of the population has adaptive expectations. The percentage is similar when the four-quarter moving average is used as the adaptive expectations proxy, although the underlying parameter is less precisely estimated.

In tables 3 and 4, I use capacity utilization rather than the unemployment rate in the instrument set. In general, the results are similar, but they are more precisely estimated. In particular, the underlying parameters are now all significant at the 10 percent level and all but one is significant at the 5 percent level. In some cases, the significance level is higher because the standard error of the coefficient is smaller. But a higher absolute value of the parameters also contributes. Because it is not clear that all the information in the currently published capacity utilization data was actually in the real-time information set, however, the fact that some of the estimates suggest marginally less rationality should be viewed with some skepticism.

The estimates suggest that the households in the Michigan survey have expectations that are closer to rational than the economists in the Livingston survey. However, the equations were not estimated over the same period. In results not shown here, I estimated the Livingston survey equation over the same sample as was used for the Michigan survey estimates; the results were similar to those in tables 1 and 3, and so suggest that the sample period is not the explanation for the differences in the results. This result, and its implication that households are more rational than the professional economists represented in the Michigan survey, has been noted before (Bryan and Gavin, 1986).

c. Comparison with earlier estimates

Fuhrer (1997) also examined the issue of the degree to which inflation expectations were perfectly rational. Unlike the present approach, however,

Table 3
Alternative Instrument Set
Structural Estimation, Models of Less-Than-Rational Expectations
Livingston Survey
Semiannual, 1957:2-1995:2

Instruments: Same as table 1, ***except*** manufacturing capacity utilization replaces the civilian unemployment rate

A. Partial adjustment of expectations

	$\rho/(1+\rho)$	ρ	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
	.918 (.280) [.001]	.48	.45	.26

B. Partly adaptive expectations

Partly adaptive portion of expectations	$-\alpha/(1-\alpha)$	α	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
Four-quarter moving average [p(t-1) - p(t-5)]	-.623 (.299) [.04]	.38	.52	.56
Eight-quarter moving average [p(t-1) - p(t-9)]/2	-.502 (.208) [.02]	.33	.57	.50

Numbers in parentheses are standard errors.
Numbers in brackets are p-values.

Table 4
Alternative Instrument Set
Structural Estimation, Models of Less-Than-Rational Expectations
Michigan Survey
Quarterly, 1966:2-1995:4

Instruments: Same as table 1, ***except*** manufacturing capacity utilization replaces the civilian unemployment rate.

A. Partial adjustment of expectations

	$\rho/(1+\rho)$	ρ	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
	.630 (.142) [.0001]	.39	.44	.48

B. Partly adaptive expectations

Partly adaptive portion of expectations	$-\alpha/(1-\alpha)$	α	Are residuals orthogonal to instruments? (P-value)	First-stage R-squared
Four-quarter moving average [p(t-1) - p(t-5)]	-.371 (.207) [.07]	.27	.31	.39
Eight-quarter moving average [p(t-1) - p(t-9)]/2	-.435 (.135) [.001]	.30	.38	.63

Numbers in parentheses are standard errors.
Numbers in brackets are p-values.

Fuhrer did not make use of information on survey measures of inflation expectations but instead embedded a specification similar to the partly adaptive model in a structural model with sticky prices, and estimated α. Fuhrer found that he could not reject the hypothesis of completely adaptive behavior — that is, that $\alpha = 1$. However, his estimates were not very precise, and his preferred point estimates were around 0.8.

With the specification I use, I do not obtain a direct point estimate of α, and I cannot directly examine the hypothesis that $\alpha = 1$. However, I can report the estimate of α implied by the upper 5 percent confidence bound of the underlying coefficient. The largest estimate of α I obtain is 0.38, in table 4. The upper 5 percent confidence bound associated with the underlying coefficient implies an α of 0.55. Hence, the results of this paper are precise enough to allow me to reject the high degree of nonrationality in Fuhrer's preferred estimate.

The overall conclusion, therefore, is of an "intermediate" degree of rationality, with perhaps 20 to 40 percent of the population using simple, backward-looking expectations. Alternatively, the partial-adjustment estimates suggest that expectations adjust between 50 and 60 percent of the way to rationality per year. Again, that's less than perfectly rational but also less than complete ignorance of future expectations.

d. Subsample stability

As mentioned in section 3, it is possible that what appears to be a deviation of expectations formation from rationality is in fact the result of rational learning after a structural break. One type of evidence that may help gauge the importance of this hypothesis is the subsample stability of the estimates: If the estimates of the model are stable over the sample period, that is encouraging that the estimates are structural and not the result of adjustments after a single structural break that dominated the sample.

To address this issue, I look for structural breaks in the key parameter in each of the models of expectations formation. I examine only results using the Livingston survey, because the sample for the Livingston survey is longer and so makes it more likely that estimates will be precise in both parts of the sample. I begin by looking for the structural break in the middle of the sample, between the first and second semesters of 1976. The two halves of the sample were dominated by distinct monetary institutions, with the first half dominated by the Bretton Woods regime, which ended in 1973. To guard against the possibility that my results are affected by the transition following the Bretton Woods regime, I also compare a first sample that ends in 1973:H1 with a second sample that begins in 1980:H1.[5]

Table 5 shows the results splitting the sample in half. It is immediately apparent that the partial adjustment model of expectations formation is not structurally stable: Using either set of instruments, the hypothesis that the estimated parameter is the same in both halves of the sample is rejected at the 1 percent confidence level. The lack of parameter stability suggests that the partial adjustment model is not a good structural representation of expectations formation. Furthermore, the estimated coefficient is large and negative in the first half of sample, suggesting that the partial adjustment model of expectations formation does not characterize the first half of the sample.

The partly adaptive expectations model fares considerably better, however. The hypothesis that the estimated coefficient is the same in the two halves of the sample is not rejected at conventional confidence levels for either set of instruments. Furthermore, the implied structural parameters are similar in the two halves of the sample.

Splitting the sample in half may allow each of the subsamples to include part of the transitional period after the breakdown of Bretton Woods, and the rational learning hypothesis suggests that it is transitions after such regime changes that may bias the results. To guard against that possibility, the

5. Another option would be to let the data choose the mostly likely breakpoint, using the techniques described in Andrews (1993) and Andrews and Ploberger (1994). However, those techniques are formally only appropriate under ordinary least squares. In particular, the small sample properties of analogous techniques under instrumental variables are not known. In any case, the 1973 and 1980 break points discussed in the text have prior economic motivation, making the Andrews-type techniques unnecessary.

Table 5

Tests of Structural Stability: Splitting Sample in Half
Using Livingston Survey

A. Unemployment rate is instrument for economic activity.

	Expectations:Partial adjustment		Partly adaptive	
	1957:2-1976:1	1976:2-1995:2	1957:2-1976:1	1976:2-1995:2
Estimated coefficient	-1.79 (.88)	1.367 (.402)	-.571 (.170)	-.933 (.159)
Structural coefficient	—	$\rho=.58$	$\alpha=.36$	$\alpha=.48$
Orthogonality of instruments	.62	.82	.58	.90
Is coefficient the same?[a]	.007		.14	

Note: Instruments are the same as those used in table 1.

B. Capacity utilization is instrument for economic activity.

	Expectations:Partial adjustment		Partly adaptive	
	1957:2-1976:1	1976:2-1995:2	1957:2-1976:1	1976:2-1995:2
Estimated coefficient	-1.32 (.82)	1.550 (.357)	-.657 (.192)	-1.147 (.141)
Structural coefficient	—	$\rho=.61$	$\alpha=.40$	$\alpha=.53$
Orthogonality of instruments	.55	.76	.55	.57
Is coefficient the same?[a]	.003		.27	

Note: Instruments are the same as those used in table 2.

a. Probability level of hypothesis that the raw coefficients were different in the two samples, using the dummy-variable technique.

estimates in table 6 limit the earlier portion of the sample to the period before 1973:H2 and the later part of the sample to the period after 1979:H2. Looking first at the results for partly adaptive expectations, the estimated coefficient is statistically significant in both portions of the sample, suggesting that expectations formation was less than perfectly rational in both periods. Furthermore, the point estimates are now closer together than when the sample was split in half — indeed, when the unemployment rate is used as the instrument for economic activity, the estimated structural parameters in the two samples are nearly identical.

The results for the partial adjustment model are more subtle. When capacity utilization is the instrument for economic activity, the partial adjustment parameter is estimated to be positive in both subsamples, the point estimates are close, and the degree of rejection of the hypothesis that they are the same is much smaller than in table 5. When the unemployment rate is the measure of economic activity in the instrument set, the estimated coefficient in the first part of the sample is once again negative. However, the estimate is sufficiently imprecise that the hypothesis of equality is no longer rejected at conventional confidence levels.

Overall, the results on structural stability cast some doubt on whether partial adjustment is a robust model of expectations formation. But the model of partly adaptive expectations held up well: It suggested a similar degree of nonrationality in two different monetary-policy environments, the Bretton Woods period and the 1980s and 1990s. Moreover, the model results were robust to including the transitional period in between.

6. Empirical implications of less-than-rational expectations

With estimates of the degree of nonrationality in hand, I next turn to the question of how well these estimates of less-than-rational expectations do in helping the New Keynesian model fit the data. The approach I take is to couple the sticky price model with less-than-rational expectations with reduced-form estimates of equations for output and monetary policy. Since the models of

Table 6

Tests of Structural Stability: Pre-1973:H2 and Post-1979:H2
Using Livingston Survey

A. Unemployment rate is instrument for economic activity.

Expectations:	Partial adjustment		Partly adaptive	
	1957:2-1973:1	1980:1-1995:2	1957:2-1973:1	1980:1-1995:2
Estimated coefficient	-1.17 (1.39)	1.605 (.444)	-.698 (.222)	-.734 (.139)
Structural coefficient	—	ρ=.62	α=.41	α=.42
Orthogonality of instruments	.63	.71	.61	.69
Is coefficient the same?[a]	.19		.44	

Note: Instruments are the same as those used in table 1.

B. Capacity utilization is instrument for economic activity.

Expectations:	Partial adjustment		Partly adaptive	
	1957:2-1973:1	1980:1-1995:2	1957:2-1973:1	1980:1-1995:2
Estimated coefficient	2.39 (1.34)	1.496 (.397)	-.954 (.290)	-.703 (.146)
Structural coefficient	ρ=.71	ρ=.60	α=.49	α=.41
Orthogonality of instruments	.69	.73	.64	.69
Is coefficient the same?[a]	.06		.20	

Note: Instruments are the same as those used in table 2.

a. Probability level of hypothesis that the raw coefficients were different in the two samples, using the dummy-variable technique.

imperfect rationality I examine nonetheless involve an important degree of forward-looking behavior, I use rational expectations solution techniques to obtain a reduced-form equation for inflation. I then compare the impulse responses from the model with the solved reduced-form equation with those from the original VAR.

I do this analysis with semiannual data, in contrast to the more usual practice of using quarterly data. Semi-annual data have a number of advantages. First, as I discuss in Roberts (1994), high-frequency movements in price data may be dominated by measurement error. Second, the version of the sticky-price model I use is Taylor's staggered contracts model (Taylor, 1979, 1980). Using semiannual data allows me to use a sensible version of that model — one with an annual contract length — without introducing the additional complication of multiple contract lengths. (Of course, an additional advantage of semi-annual data is that it "lowers the bar" for the model, since it does not need to fit quite as many movements in the time series data.)

As I've shown elsewhere (Roberts, 1995), Taylor's model implies a formulation that has the broad features of the standard New Keynesian Phillips curve (equation 1). However, the Taylor model captures an additional element of inflation dynamics, since it takes account of the fact that most wages are set for fixed periods of time.

The version of the staggered contracts model I use is:

$$p(t) = [x(t) + x(t-1)]/2 + \epsilon(t) , \tag{13}$$

$$x(t) = E(t)\,[p(t) + p(t+1)]/2 + \gamma\, y(t) + \nu(t), \tag{14}$$

where p is the log of the price level and x is the log of the "contract" wage. The contract wage is assumed to stay in effect for one year, or two semesters. Thus, firms are assumed to set prices as a fixed markup over the average wages of their workers, half of whom set their wages in each period. And wages are set with an eye to trading off the real wage over the life of the contract against the level of economic activity.

I consider the two mechanisms of expectations formation discussed above, partly adaptive and partial adjustment. I also consider the limiting case of perfectly rational expectations. In all cases, inflation expectations have an important "rational" component. In order to solve the model and obtain an impulse response that I can compare to that from a VAR, I use the "AIM" procedure developed by Anderson and Moore (1985), which solves linear rational expectations models and obtains a reduced-form solution.

Aside from the degree of rationality, the only other parameter in the inflation model is the γ parameter, which determines the slope of the Phillips curve. I use the value of 0.025 for γ, which is from Roberts (1997), table 5, for results referring to capacity utilization.[6] I also assume that all of the error in the inflation process comes from the "price" shock, ϵ, and none from the " contract wage" shock, ν. That assumption amounts to assuming that shocks to inflation are dominated by prices that are not a simple markup over wages —in particular, that description would fit crude oil and other commodity prices.

The VAR includes four lags of: the semester-average manufacturing capacity utilization; the change in the inflation rate, where inflation is the percent change over the semester in the consumer price index; and the real federal funds rate, which I measure by subtracting the current-period inflation rate from the semester-average of the federal funds rate.[7] I use the change in inflation in the model since the evidence suggests that inflation was non-stationary over the estimation period (1957 to 1997).

As is always the case in a VAR analysis, it is necessary to choose a contemporaneous ordering of the variables. I've assumed that capacity utilization is contemporaneously exogenous, that inflation is affected by capacity utilization but not the interest rate, and that the interest rate is affected

---

6. In Roberts (1997), the coefficient on capacity utilization in an equation where the dependent variable is annualized semiannual inflation is 0.1. Since the parameter in a non-annualized version of the equation is 2 γ, the result, correctly adjusted, implies γ is 0.025.

7. Ideally, I would use model-consistent inflation expectations to form the real interest rate. However, because the federal funds rate is literally an overnight interest rate, it is not clear over what horizon the expectations ought to apply; contemporaneous inflation seems as good as any other option.

by both variables. This ordering has previously been discussed at length by Christiano, Eichenbaum, and Evans (1994). Briefly, its justification is that the Federal Reserve, in setting monetary policy, has a good idea of the current state of the economy. So, inflation and output can affect monetary policy contemporaneously. However, monetary policy is slow in affecting output and inflation, so the contemporaneous relationship is one way only. Finally, the shock to inflation is assumed not to affect output in the current period.

The solid lines in figure 1 show the results for the impulse responses from the VAR. Recall that an impulse response shows the forecasts of each of the variables to a one-standard-deviation perturbation of each of the shocks. The results are entirely conventional: A shock that raises the federal funds rate leads to a large and statistically significant reduction in output. Inflation gradually falls, and there is a statistically significant permanent reduction in inflation.[8] Eventually, the fed funds rate also falls, as the long-run affect on the real interest rate is zero. The output loss that occurs during the period in which inflation achieves its permanent reduction can be used to calculate a sacrifice ratio. The VAR results imply that reducing inflation by one percentage point at an annual rate through a surprise increase in the federal funds rate entails a reduction in capacity utilization below what it otherwise would have been of 4.3 percentage points for one year. That's comparable to the 2 percentage point sacrifice ratio for the unemployment rate cited in the introduction to the paper, since capacity utilization is more cyclical than the unemployment rate.

The output shock also leads to a persistent increase in output. Output returns to normal after about two years, and actually dips below its original level for about two years, although the effect is only marginally significant. As we might expect in response to an aggregate demand shock, inflation rises in

8. Some readers may find it surprising that inflation falls steadily in response to the interest-rate shock. In particular, it does not exhibit the so-called price puzzle, where inflation rises, sometimes significantly, in response to a monetary shock (Christiano, Eichenbaum, and Evans, 1994). These results suggest that the price puzzle may not be a robust feature of the data.

Figure 1

Comparison of Impulse Responses: VAR and Model with Partly Adaptive Expectations

25 percent adaptive expectations; adaptive expectations are a four-semester moving average.

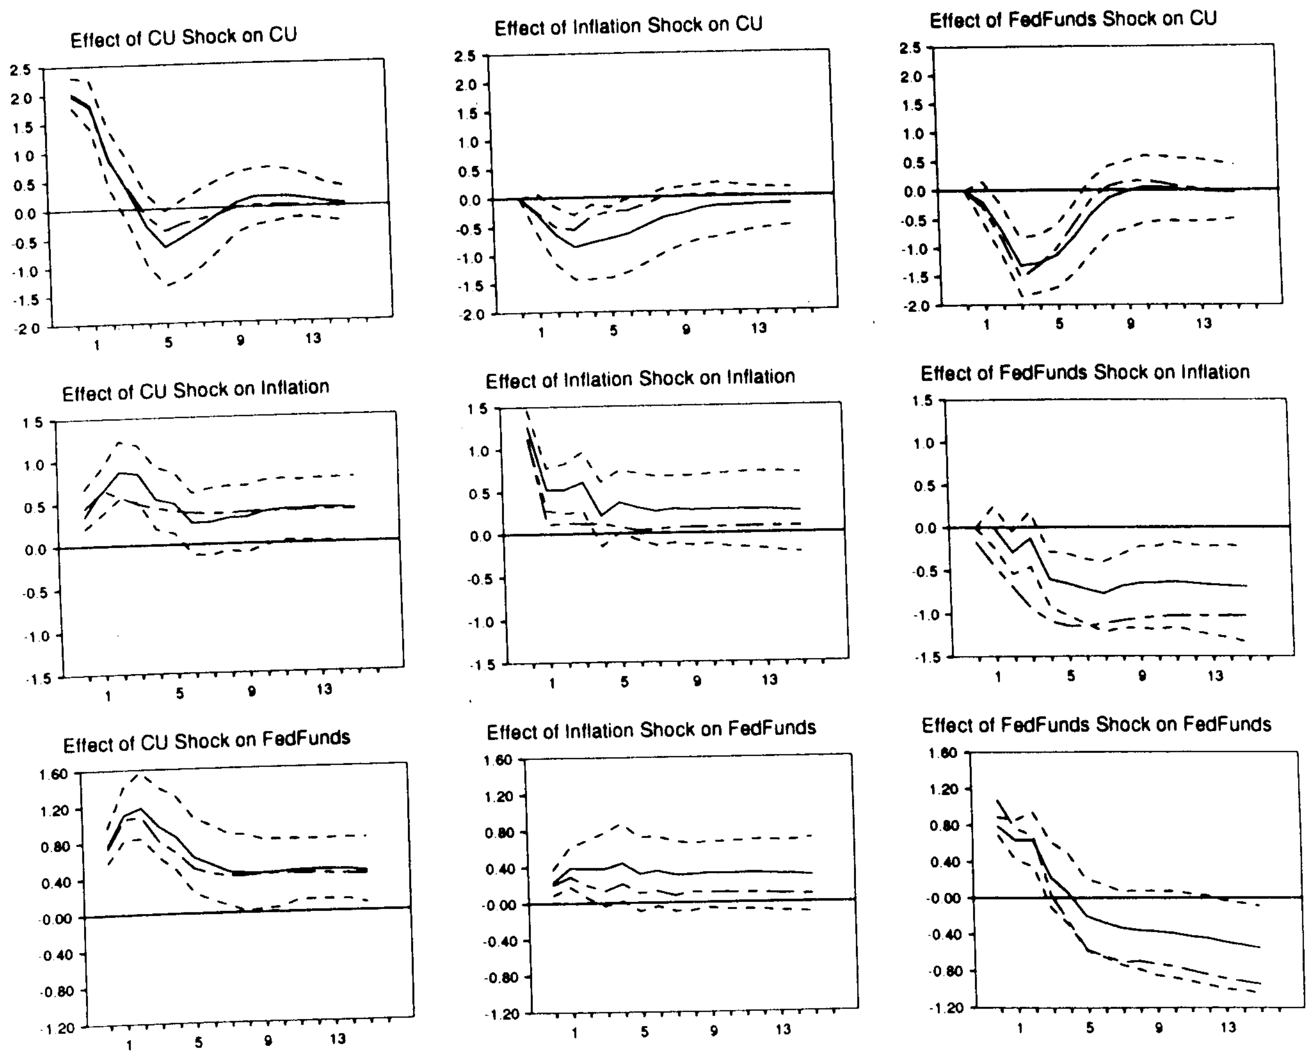

Solid lines are unconstrained VAR impulse responses.
Dashed lines are 5 and 95 percent confidence bands.
Mixed dashed lines are model impulse responses.

response to the high output. Inflation ends up higher than initially, but the effect is only marginally significant at the 5 percent level. The interest rate rises in response to the output shock, and by more than inflation, implying an increase in real rates. The increase in rates may represent a tightening in monetary policy, and so may explain why output drops below its initial level.

The inflation shock raises inflation initially. The effect wears off in large part after two years. Inflation ends up higher than it was initially, but the effect is not significant. The shock leads to a statistically significant reduction in output. The effect is long-lived, and only drops below statistical significance after four years. Interest rates move very little, and generally less than inflation.

The mixed dashed lines show simulation results assuming adaptive expectations. Using the unemployment rate as an instrument, the most precise estimate of this model was from the specification for the Michigan survey and a two-year moving average of past inflation. The coefficient estimate in that case had the interpretation that 25 percent of the population had adaptive expectations.

The results in figure 1 suggest that the model does a good job of capturing major features of the data, especially considering that no parameters of the inflation model were estimated in order to fit these data. In particular, the model predicts that there will be a large output loss in response to a monetary policy shock that permanently reduces inflation, and the path of output is similar to that in the VAR. Inflation, however, falls by rather more than in the VAR. For the first six semesters after the shock, inflation is significantly lower than in the VAR. With the output response similar to the VAR but the inflation response larger, the sacrifice ratio is only 2.5, compared with 4.3 for the VAR. Finally, consistent with the lower inflation in the model's response, interest rates soon fall below those of the VAR response.

The patterns for the output shock are also good. Inflation rises in response to the output shock almost as much as in the VAR, although there is a noticeable "hump" pattern to the inflation response that the model fails to capture.

The model does perhaps least well in capturing the effects of the inflation shock. The effects on inflation itself fall off much more rapidly in the model simulation than in the VAR. The shock's effects on output are also muted relative to the VAR. As in the VAR, there is little effect on interest rates.

Figure 2 repeats the exercise assuming perfectly rational expectations. The results are very different. In particular, the monetary policy shock leads to disinflation with very low costs: Inflation comes down to about its long-run level after half a year, and output moves very little from trend. The sacrifice ratio is only 0.7, a small fraction of the estimate from the unconstrained VAR. There is less difference with the VAR in response to the output shock than with the interest rate shock, although there, too, the differences for inflation are statistically significant. These results indicate that even assuming only 25 percent of the population has adaptive expectations has important effects on the properties of the model. These results are similar to those of Fuhrer and Moore (1995), although their interpretation was different: Recall from section 2 that they referred to the modification of the sticky-price model that I call "partly adaptive expectations" as "real staggered contracts."

If we are willing to accept the results using capacity utilization as an instrument, then up to one-third of the population has two-year adaptive expectations (table 4). In this case (not shown), the overall pattern of results is similar to that of figure 1. However, the reduction in inflation in response to an interest rate shock is a bit less than in the earlier case, which keeps the impulse response within the lower 5 percent confidence bound in all but one period. Also, the output loss becomes a bit greater than in the VAR. Both of these changes tend to raise the sacrifice ratio, to 3.6, not far from the unrestricted estimate of 4.3.

For the Livingston survey, the estimates of the adaptive expectations parameter were more precise when a four-quarter moving average of inflation was used. In this case, the fraction of the population that had adaptive expectations was 36 percent (when unemployment was used as an instrument). The impulse responses from the model with this assumption are shown in

Figure 2

Impulse Responses from Model Assuming Rational Expectations

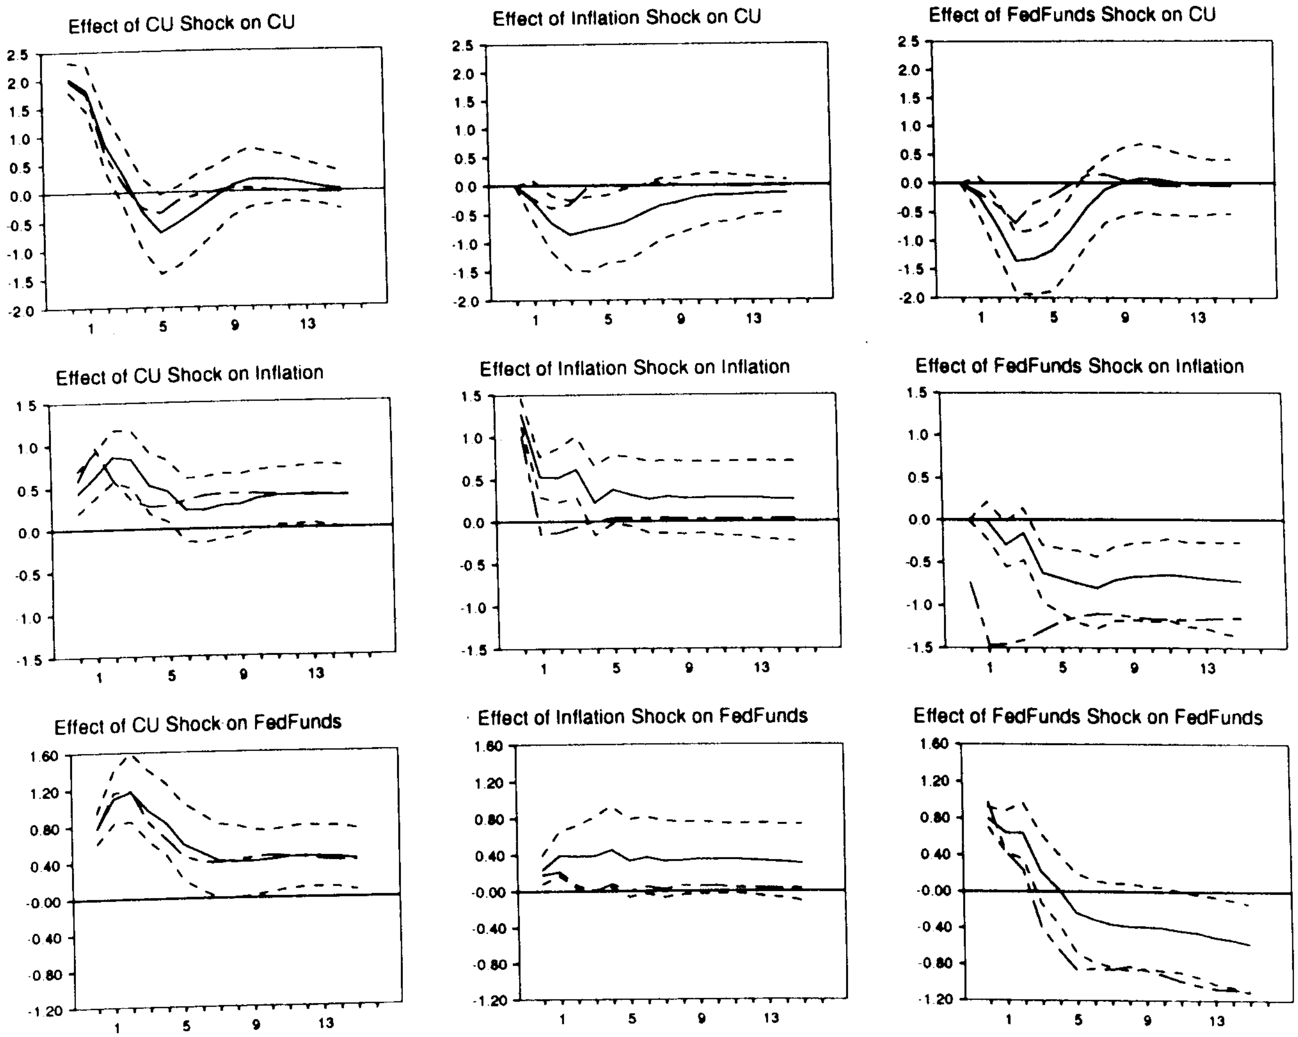

Solid lines are unconstrained VAR impulse responses.
Dashed lines are 5 and 95 percent confidence bands.
Mixed dashed lines are model impulse responses.

figure 3. The results for the response to the federal funds rate shock are qualitatively similar to those in figure 1. In the present case, the effects on output are somewhat smaller than in the VAR, but they are within the 5 percent confidence bound. The inflation response is somewhat smaller than in figure 1; it breaches the 5 percent band in several periods early on, but its long-run response is close to the unrestricted response. Although the responses to the interest rate shock are closer to the VAR impulse responses than in figure 1, the sacrifice ratio is actually a bit smaller than before, at 2.3.

Where the model using a one-year inflation lag is a notable improvement is in its ability to capture the effects of the inflation shock. The model predicts much greater short-run persistence in inflation than the model using two years of lagged inflation. The responses of output and the interest rate are also much closer to the actual pattern.

In figure 4, I switch to the partial adjustment model of expectations formation, assuming that expectations move 61 percent per year, or 38 percent per semester, to the rational value, consistent with the results using the Michigan survey. For the fed funds shock, the model simulations match the VAR impulse responses almost exactly. Nonetheless, the sacrifice ratio is 5.8, about a third larger than the value from the VAR. The model also matches the effects of the output shock almost exactly. However, the partial adjustment model does a poor job of accounting for the effects of the inflation shock: The model implies virtually no short-run persistence in inflation, and virtually no response of the other two variables. If I had used the partial adjustment speed estimated from the Livingston survey, which suggested only 54 percent annual adjustment, the sacrifice ratio would have been even larger, at 7.9.

Figure 3

Impulse Responses from Model with 36 Percent Adaptive Expectations
Adaptive expectations are a two-semester moving average.

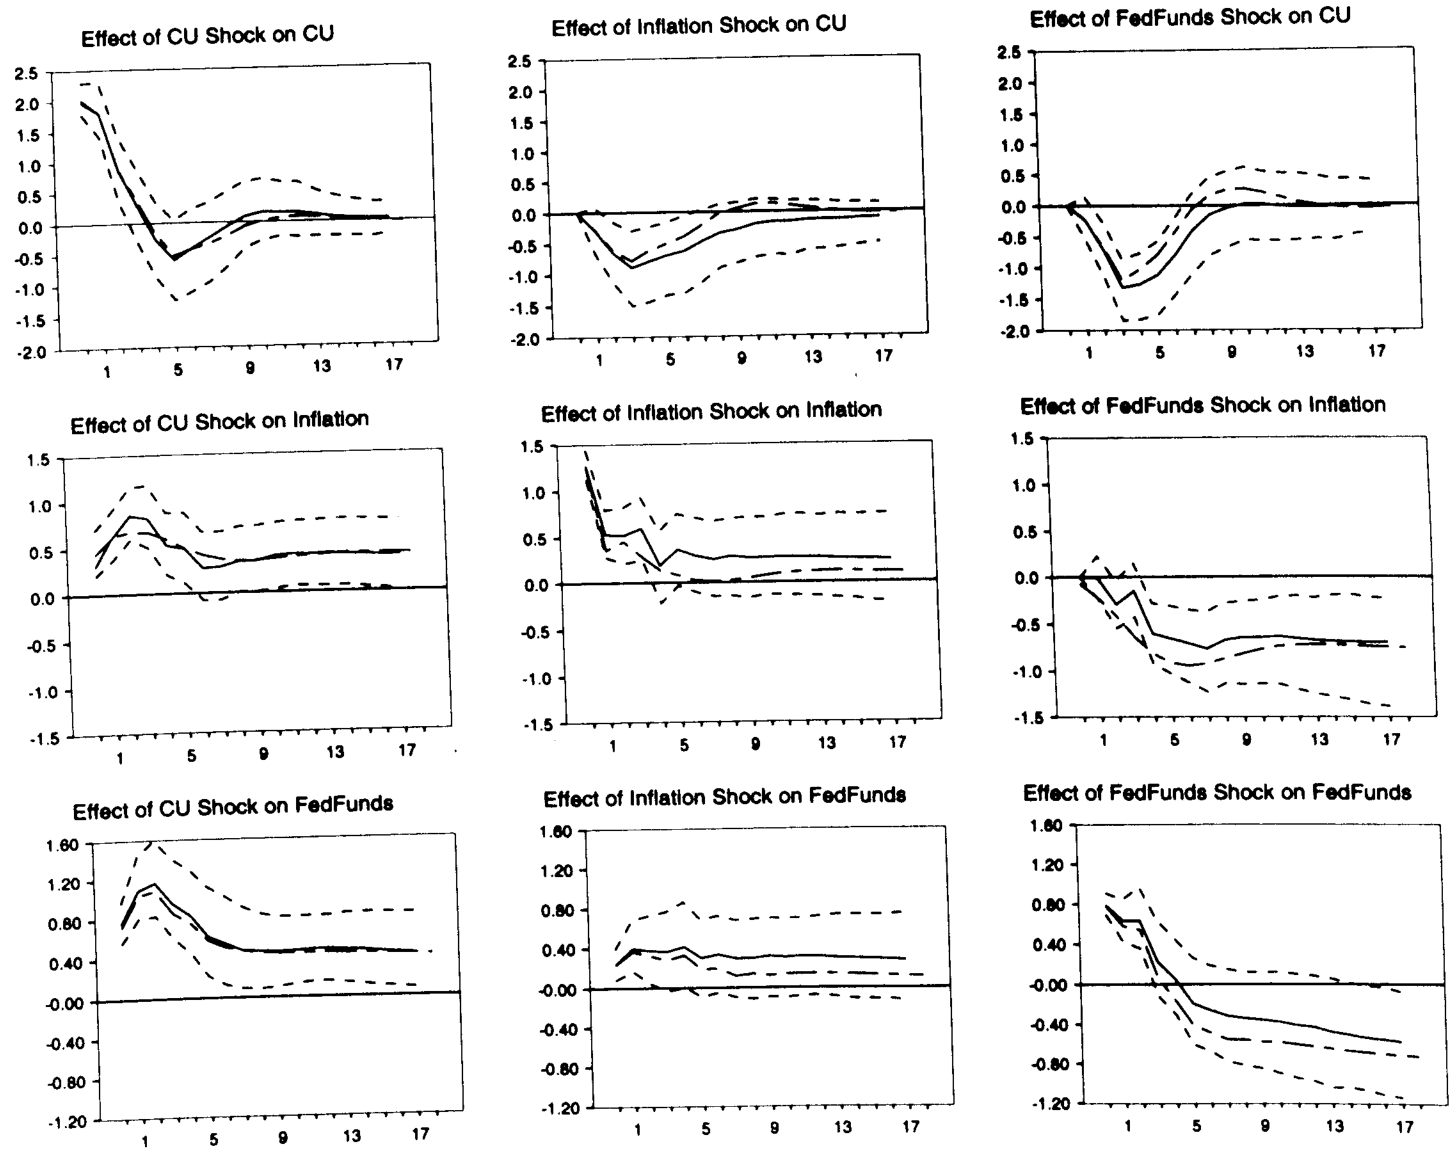

Solid lines are unconstrained VAR impulse responses.
Dashed lines are 5 and 95 percent confidence bands.
Mixed dashed lines are model impulse responses.

Figure 4

Impulse Responses from Model with Partial Adjustment of Expectations

Expectations adjust 38 percent per semester.

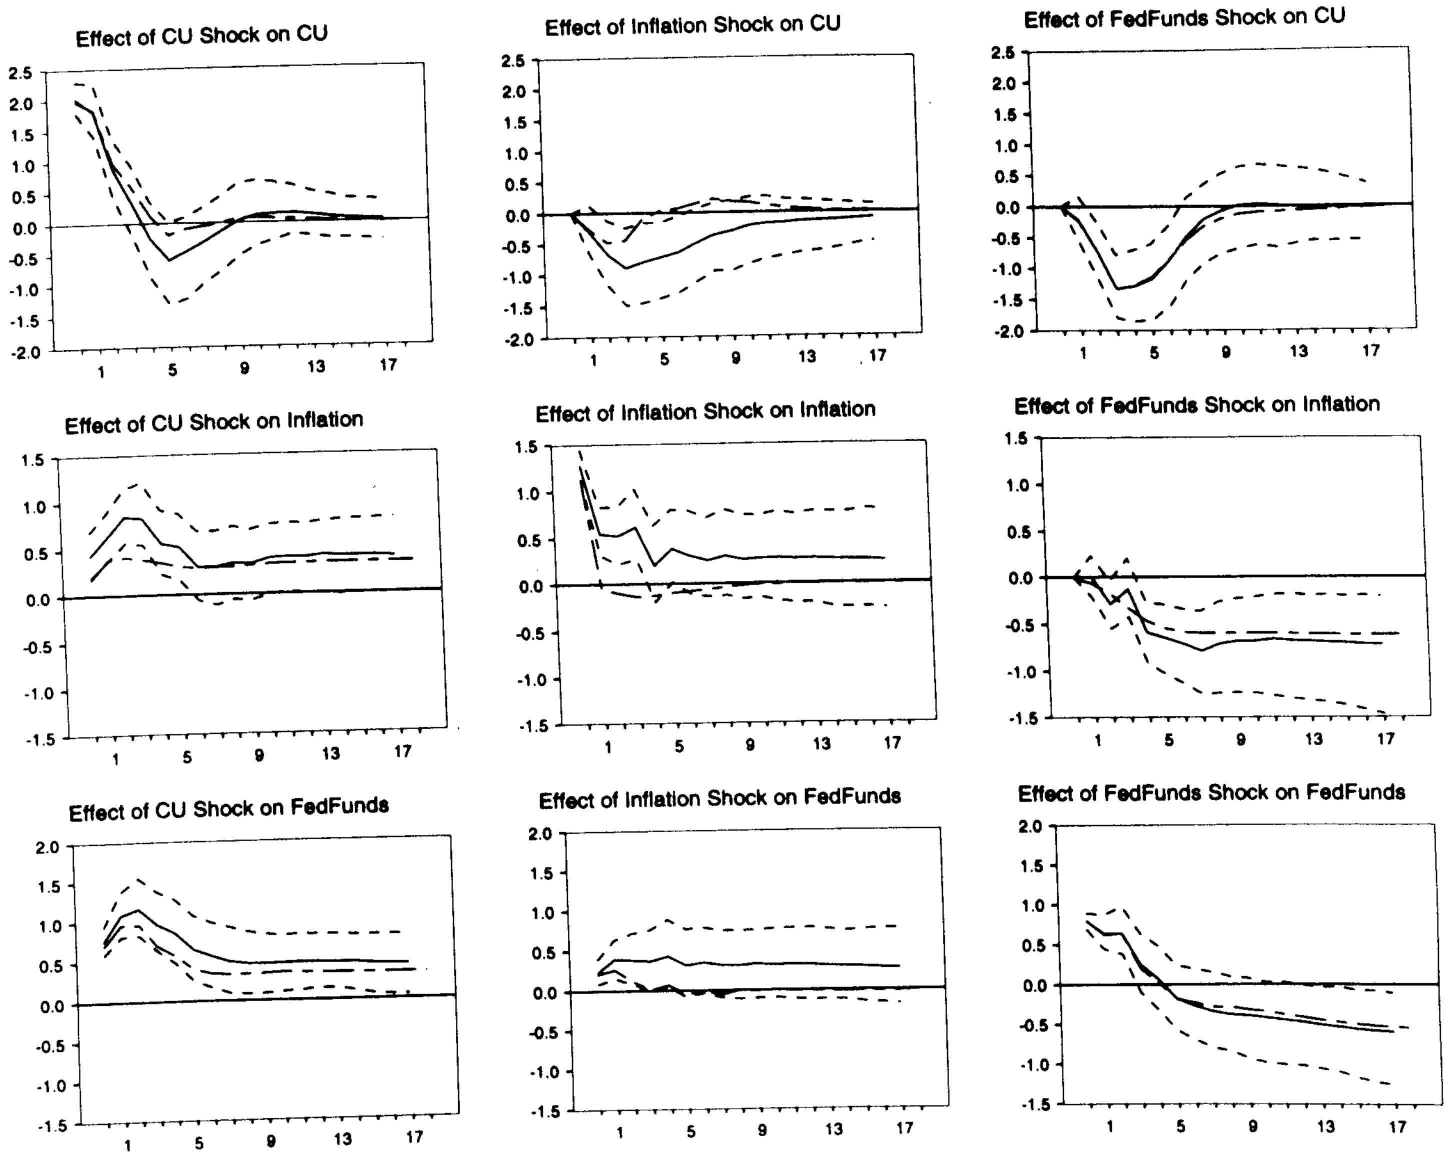

Solid lines are unconstrained VAR impulse responses.
Dashed lines are 5 and 95 percent confidence bands.
Mixed dashed lines are model impulse responses.

7. Conclusions

The main results of the paper are two:

- Survey measures of inflation expectations represent an intermediate degree of rationality. The results in this paper are precise enough to reject both perfect rationality and very high levels of nonrationality.
- Simulation results that assume the intermediate degree of rationality suggested by the surveys are able to capture the costs of reducing inflation in historical U.S. data.

Together, these results bolster the case that deviations of inflation expectations from rationality play an important role in explaining why attempts by central banks to reduce inflation have historically required costs in terms of lost output and employment.

References

Anderson, Gary S., and George Moore, 1985, A linear algebraic procedure for solving linear perfect foresight models, **Economics Letters** 17, 247-52.

Andrews, Donald, 1993, Tests for parameter instability and structural change with unknown change point, **Econometrica** 61, 821-56.

______________, and Werner Ploberger, 1994, Optimal tests when a nuisance parameter is present only under the alternative, **Econometrica** 62, 1383-1414.

Baghestani, Hamid, 1992, Survey evidence on the Muthian rationality of the inflation forecasts of U.S. consumers, **Oxford Bulletin of Economics and Statistics** 54, 173-86.

Ball, Lawrence J., 1995, Disinflation with imperfect credibility, Journal of **Monetary Economics** 35, 5-24.

______________, 1994, What determines the sacrifice ratio? in Gregory Mankiw, ed., **Monetary Policy**, MIT Press.

______________, 1991, The genesis of inflation and the costs of disinflation, **Journal of Money, Credit, and Banking** 23, 439-52.

______________, and Dean Croushore, 1995, Expectations and the effects of monetary policy, Working Paper no. 95-22, Federal Reserve Bank of Philadelphia, October.

Batchelor, Roy A., and Pami Dua, 1989, Household versus economist forecasts of inflation: A reassessment, **Journal of Money, Credit, and Banking** 21, 252-7.

Blinder, Alan, 1987, **Hard heads, soft hearts** (Addison-Wesley, Reading, Massachusetts).

Bryan, Michael F., and William T. Gavin, 1986, Models of inflation expectations formation: A comparison of household and economist forecasts, **Journal of Money, Credit, and Banking** 18, 539-44.

Buiter, Willem H., and Ian Jewitt, 1981, Staggered wage setting with real wage relativities: Variations on a theme of Taylor, in Willem H. Buiter, **Macroeconomic theory and stabilization policy** (University of Michigan Press, Ann Arbor, 1989), 183-99.

Calvo, Guillermo A., 1983, Staggered contracts in a utility-maximizing framework, **Journal of Monetary Economics** 12, 383-398.

Caskey, John, 1985, Modeling the formation of price expectations: A Bayesian approach, **American Economic Review** 75, 768-776.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans, 1994, Identification and the effects of monetary policy, Working Paper no. WP-94-7, Federal Reserve Bank of Chicago, May.

Fuhrer, Jeffrey C., 1997, The (un)importance of forward-looking behavior in price specifications, **Journal of Money, Credit, and Banking** 29, 338-50.

Fuhrer, Jeffrey C., and George R. Moore, 1995, Inflation persistence, **Quarterly Journal of Economics** 110, 127-60.

______, 1992, Monetary policy rules and the indicator properties of asset prices, **Journal of Monetary Economics** 29, 303-36.

Gramlich, Edward M., 1983, Models of inflation expectations formation, **Journal of Money, Credit, and Banking** 15, 155-73.

Keane, Michael P., and David E. Runkle, 1990, Testing the rationality of price forecasts: New evidence from panel data, **American Economic Review** 80, 714-35.

Lewis, Karen K., 1989, Changing beliefs and systematic rational forecast errors with evidence from foreign exchange, **American Economic Review** 79, 621-36.

Mullineaux, Donald J., 1980, Inflation expectations and money growth in the United States, **American Economic Review** 70, 149-61.

Pearce, Douglas K., 1987, Short-term inflation expectations: Evidence from a monthly survey, **Journal of Money, Credit, and Banking** 19, 388-95.

Pesaran, M. Hashem, 1987, **The limits to rational expectations** (Basil Blackwell, Oxford).

Rich, Robert W., 1989, Testing the rationality of inflation forecasts from survey data: Another look at the SRC expected price change data, **Review of Economics and Statistics** 71, 682-6.

Roberts, John M., 1997, Is inflation sticky? **Journal of Monetary Economics** 39, 173-96.

____________, 1995, New Keynesian economics and the Phillips curve, **Journal of Money, Credit, and Banking** 27, 975-84.

______________, 1994, Sticky prices and measurement error, Photocopy, Federal Reserve Board (September).

Taylor, John B., 1979, Staggered contracts in a macro model, **American Economic Review** 69, 108-113.

_______, 1980, Aggregate dynamics and staggered contracts, **Journal of Political Economy** 88, 1-23.