Adaptive Least Squares Estimation of the Time-Varying Taylor Rule

J. Huston McCulloch

Ohio State University

June 6, 2007

Preliminary version prepared for
13th International Conference on Computation in Economics and Finance, Montreal,
June 14-16, 2007.

URL: www.econ.ohio-state.edu/jhm/papers/TaylorALS.pdf

Keywords: Adaptive Learning, Taylor Rule, Adaptive Least Squares, Expected Inflation, Unemployment Gap

JEL Codes:
E52 – Monetary Policy
C22 – Time Series Models

Abstract

Clarida, Galí and Gertler (CGG 2000), Orphanides and Williams (2005), Kim and Nelson (2006), and others have found time variation in the Fed's "Taylor Rule" interest rate policy response function. CGG arbitrarily break their period into two fixed-coefficient 20-year subperiods, however, rather than letting the data tell them when and if any shift in the coefficients occurred. They also proxy their expected inflation and output gap variables using fixed coefficients on several instrumental variables during each subperiod. This effectively models agents as making forecasts with coefficients derived from data that has not yet been observed.

The present paper uses the author's Adaptive Least Squares algorithm (McCulloch 2005a) to estimate this policy response function. First, it is used to simulate expected inflation and the expected unemployment gap using continually time-varying filter estimates that are based only on data that agents could have observed at the time the simulated forecast is made. Then, it is used to estimate the Fed's policy response function, using continually changing smoother estimates that employ both past and future data.

The Likelihood Ratio statistic rejects the hypothesis of constant coefficients in all three equations. Smoother estimates of the Taylor Rule indicate that the coefficient on expected inflation rose from barely 1.0 to 2.0 or higher during 1975-1980. It fell to 1.6 in the 1990's, but has been nearly 2.0 since 2003. The response to the unemployment gap has been negative throughout, and was strongest near 1980 and since 2002. The equilibrium inflation rate consistent with the Taylor Rule coefficients could be anywhere in the range of 1% to 6% since 2003, depending on the unobserved equilibrium real interest rate.

1. Introduction

Clarida, Galí and Gertler (CGG 2000), Orphanides and Williams (2005), Kim and Nelson (2006), and others have found time variation in the Fed's "Taylor Rule" interest rate policy response function. However, the widely cited CGG paper arbitrarily breaks the period it studies into two fixed-coefficient 20-year subperiods, rather than letting the data dictate when and if any shift in the coefficients occurred. They also proxy their expected inflation and output gap variables using fixed coefficients on several instrumental variables during each subperiod. This effectively models agents as making forecasts with coefficients derived from data that has not yet been observed.

The present paper uses the author's Adaptive Least Squares (ALS) algorithm (McCulloch 2005a, 2005b) to estimate this policy response function, in terms of expected inflation and the expected unemployment gap. Longbrake and McCulloch (2007) find that the conventional "output gap" employed by CGG and others may simply be a statistical illusion. The natural rate of unemployment, even if uncertain and time-varying, is a much more solid concept than potential output. This paper therefore uses the unemployment gap – the difference between the actual and natural rates of unemployment – in place of the customary output gap.

ALS is used to simulate expected inflation and the expected unemployment gap using continually time-varying *filter* estimates that are based only on data that agents could have observed at the time the simulated forecast is made. Then, it is used to retroactively estimate the Fed's policy response function, using continually changing *smoother* estimates that employ both past and future data.

Section 2 below briefly reviews the ALS model. Section 3 describes the data. Section 4 uses ALS to estimate a Vector Autoregression (VAR) of inflation on past inflation and unemployment. Section 5 models unemployment as a VAR of past inflation and unexpected inflation, as determined by the inflation model of Section 4. Section 6 then fits a Taylor Equation for the Federal funds rate, using the expected inflation series constructed in Section 4 together with the unemployment gap series implied by the model of Section 5. Section 7 concludes.

2. The Adaptive Least Squares Model.

The Adaptive Least Squares (ALS) model of McCulloch (2005a) considers a time-varying-parameter linear regression model of the form

$$\begin{aligned} y_t &= \mathbf{x}_t \boldsymbol{\beta}_t + \varepsilon_t, \quad \varepsilon_t \sim NID(0, \sigma_\varepsilon^2), \\ \boldsymbol{\beta}_t &= \boldsymbol{\beta}_{t-1} + \boldsymbol{\eta}_t, \quad \boldsymbol{\eta}_t \sim NID(\mathbf{0}_{k\times 1}, \mathbf{Q}_t), \end{aligned} \tag{1}$$

$t = 1, \ldots n$, in which $\mathbf{x}_t$ is a $1\times k$ row vector of ideally exogenous explanatory variables, $\boldsymbol{\beta}_t$ is a $k\times 1$ column vector of time-varying coefficients, and $\boldsymbol{\eta}_t$ is a $k\times 1$ column vector of transition errors that are independent of the observation errors ε_t. Let $\mathbf{y}_t$ be the $t\times 1$ vector of dependent variables observed up to and including time t. Ordinarily the first element of each $\mathbf{x}_t$ is unity, so that the first element of $\boldsymbol{\beta}_t$ is the intercept. $\mathbf{Q}_t$ is the $k \times k$ covariance matrix of the transition errors $\boldsymbol{\eta}_t$.

ALS makes the special assumption

$$\mathbf{Q}_t = \rho T_{t-1} \mathbf{P}_{t-1} . \tag{2}$$

where the *Effective Sample Size*, T_t, is determined, as in the famous Local Level Model of Muth (1960) and Kalman (1960), by

$$T_t = (1 + \rho T_{t-1})^{-1} T_{t-1} + 1 \tag{3}$$

with initialization

$$T_0 = 0 , \tag{4}$$

for a fixed signal/noise parameter ρ, to be estimated. The Effective Sample size at first grows almost as fast as the actual sample size t, but is ultimately bounded above by

$$T_\infty = \lim_{t \uparrow \infty} T_t = 1/2 + \sqrt{1/4 + 1/\rho} , \tag{5}$$

the unique positive root of the quadratic equation that defines the fixed points of (3).

Under this model, it can be shown that the Extended Kalman Filter implies that

$$\boldsymbol{\beta}_t \mid \mathbf{y}_t \sim N(\mathbf{b}_t, \mathbf{P}_t) \tag{6}$$

where

$$\mathbf{b}_t = \mathbf{W}_t^{-1} \mathbf{z}_t , \tag{7}$$

$$\mathbf{P}_t = \sigma_\varepsilon^2 \mathbf{W}_t^{-1} , \tag{8}$$

$$\mathbf{z}_t = (1 + \rho T_{t-1})^{-1} \mathbf{z}_{t-1} + \mathbf{x}_t ' y_t , \tag{9}$$

$$\mathbf{W}_t = (1 + \rho T_{t-1})^{-1} \mathbf{W}_{t-1} + \mathbf{x}_t ' \mathbf{x}_t . \tag{10}$$

This is equivalent to the Recursive Least Squares discussed by Sargent (1999) and Evans and Honkapohja, except that the gain $\gamma_t = 1/T_t$ varies optimally with t so as to nest the Local Level Model when $k = 1$, rather than being a constant. The gain is bounded below by $\gamma_\infty = 1/T_\infty$, however.

A slightly more complicated recursive formula (see McCulloch 2005a for details) gives the "smoother" estimates of the coefficient matrix, based on the entire sample of size n rather than just the first t observations as with the filter:

$$\boldsymbol{\beta}_t \mid \mathbf{y}_n \sim N(\mathbf{b}_t^S, \mathbf{P}_t^S) . \qquad ..(11)$$

The predictive error decomposition, based on the filter estimates, can be used to compute the Maximum Likelihood (ML) estimate of the hyperparameter ρ, and therefore of the sequence T_t and its limit T_∞. The hypothesis that there is no change in the coefficients, i.e. that $\rho = 0$, may then be tested by means of the Likelihood Ratio (LR) statistic. Since this null hypothesis is on the boundary of the permissible parameter space, it does not have the customary χ_1^2 distribution. Nevertheless, its critical values can be approximated by Monte Carlo methods. The exact critical values depend on the sample size and the number of regressors (k), but in general the 5% critical value is approximately 2.3.

3. Data

Data on the seasonally adjusted Personal Consumption Expenditures Chain-type price index P_t, the seasonally adjusted civilian unemployment rate U_t, and the Federal Funds rate R_t were obtained from the St. Louis Fed's FRED data base. The PCE deflator, series PCECTPI vintage 3/29/07, was available quarterly from 1947Q1 to 2006Q4, and annualized percentage quarterly inflation π_t computed from it as 400 $\log(P_t/P_{t-1})$.

Inflation prior to 1953Q1 is much noiser than that after this date, because of improvements in CPI data collection that occurred at that time and carried over to the PCE deflator. Consequently, inflation before 1953Q1 was not used. The series employed is shown in Figure 1 below.

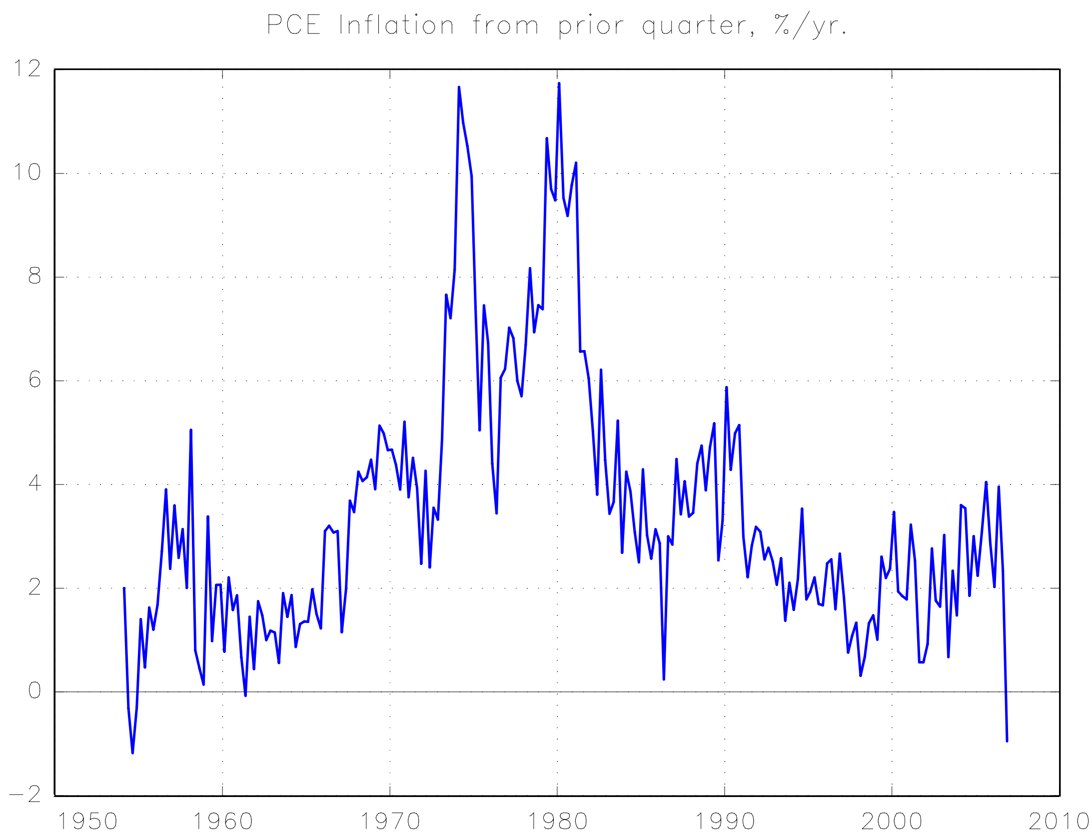

Figure 1

Unemployment, series UNRATE, was obtained monthly, 1948.01 – 2007.03, and quarterly averages of the monthly figures computed from these for 1948Q1 to 2007Q1. In order to align with the inflation data, only 1953Q1 – 2006Q4, shown in Figure 2 below, was employed.

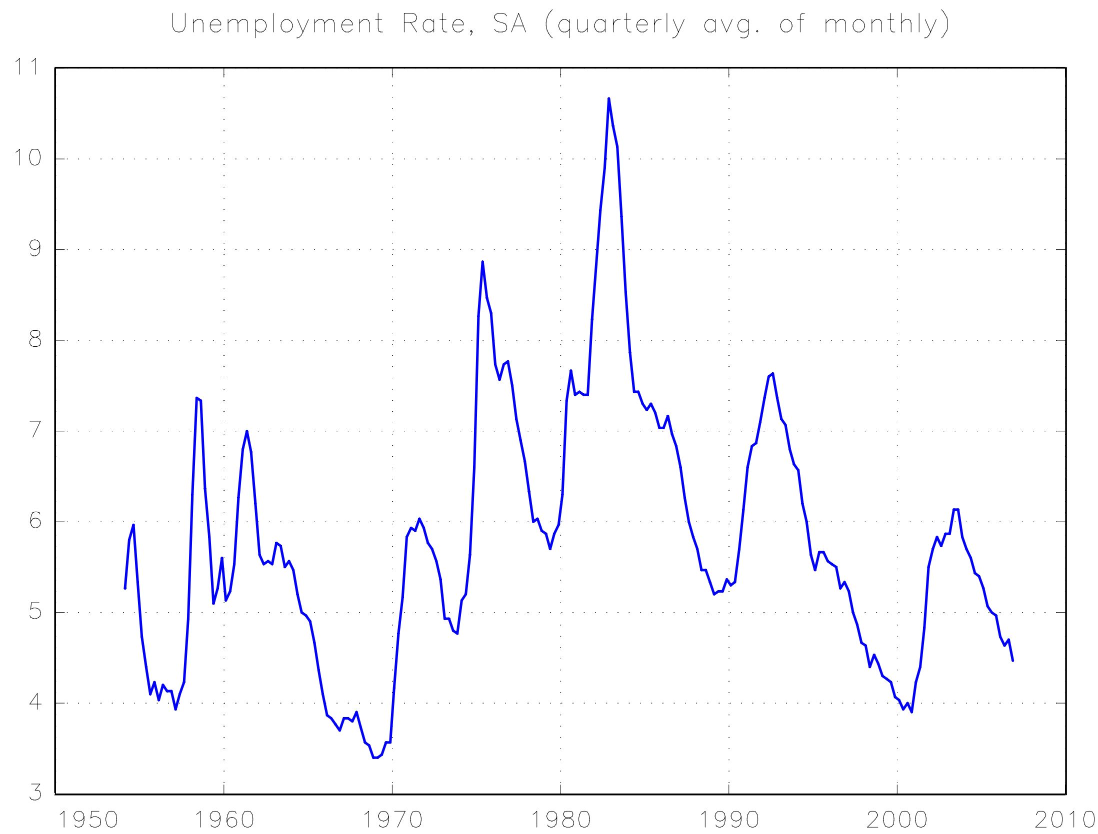

Figure 2

The effective Federal Funds rate, series FEDFUNDS, was obtained as monthly averages of daily figures, 1954.07 – 2007.03. Each quarter, from 1954Q3 to 2007Q1, was represented by the last month of the quarter, so as to reflect Fed policy at a "point" in time when most of the data for the quarter in question was already realized (even if not yet published). The portion of the series employed, from 1960Q1 to 2006Q3, is depicted in Figure 13 below.

4. The Inflation Equation.

Inflation was regressed on a constant and up to 4 lags of itself and unemployment, using ALS. In order to accommodate up to 4 lags of inflation without using data before 1953Q1, the regressions were run on data from 1954Q1 to 2006Q4.

In order to keep the model relatively parsimonious, a general-to-specific model selection procedure was used, in which all 4 lags of both variables were first used. The maximum absolute t-value for the smoother estimates of the longest lag of each variable was computed, and if either value was less than 1.96, the lag with the smaller value was omitted. The model was then reestimated with one fewer lag coefficient, and the procedure repeated until both of the longest surviving lag coefficients had maximum absolute t-values in excess of 1.96. The selected model had 3 lags of inflation itself and 4 lags of unemployment.[1] Its estimates are given below:

$n = 212$
$T_\infty = 38.6$ (quarters)
$LR = 4.55$

The limiting effective sample size of 38.6 quarters may be thought of as indicating that the parameters turn over approximately every 38.6/4 = 9.6 years, so that there were in effect about 5 or 6 different regimes observed. The LR statistic for the hypothesis of no change in the parameters ($\rho = 0$) well exceeds the approximate 5% critical value of 2.3.

[1] In the selected model, the fourth lag of U_t had a maximal t value that only barely exceeded 2, while the third lag was insignificant throughout, so that 2 lags of U_t would have given similar results. The first and second lags of U_t and all 3 lags of π_t were highly significant, however.

Figure 3 below shows the filter coefficients on the three included lags of inflation in the inflation equation, along with their sum.[2] Because the scaled precision matrix $\mathbf{W}_t$ is singular until the sample size t equals the number of coefficients estimated k, the filter is not even calculable until k = 6 quarters have elapsed. Even then, the effective sample size starts off quite small, resulting in very low precision and high variance to the filter estimates. The erratic behavior of the filter coefficient estimates during the first couple of years is therefore to be expected. By 1958 they appear to have stabilized, however.

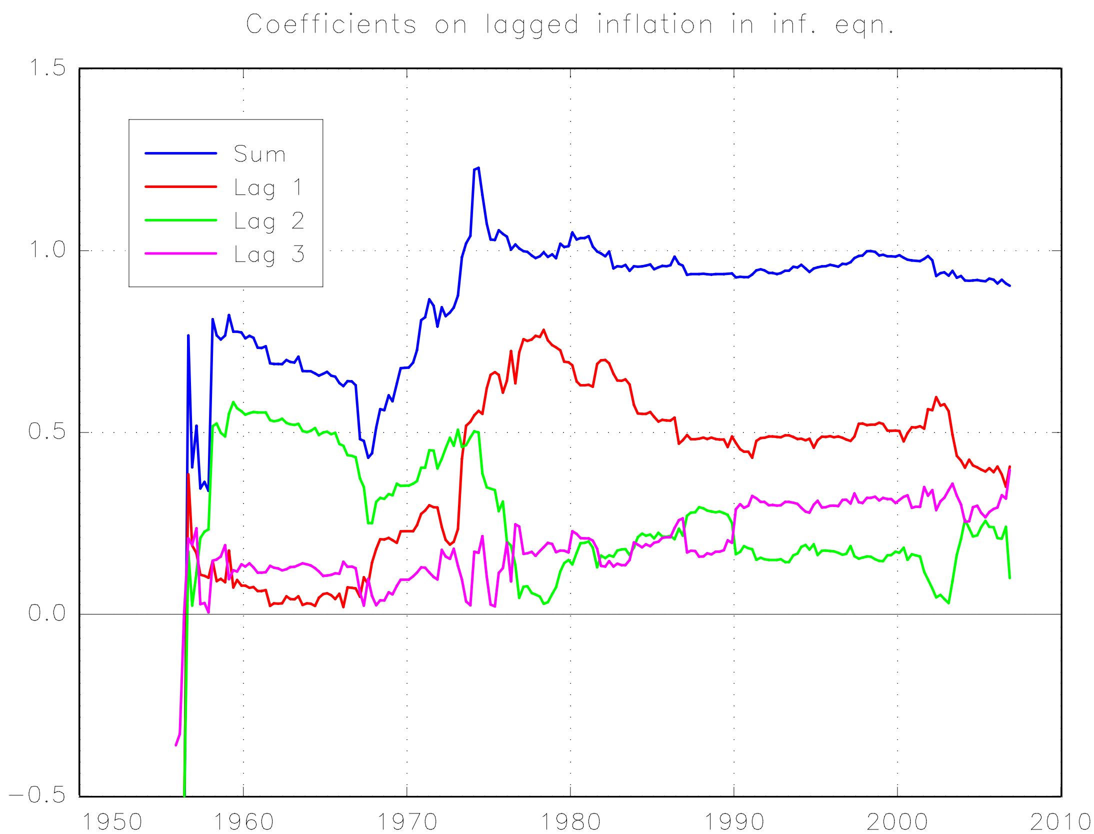

Figure 3

In a general 2-variable VAR, the stability of each variable will depend, in general, on all the coefficients of both equations. However, as long as unemployment is itself stationary (about a time-varying natural rate), the stationarity of inflation (about a time-varying long-run inflation rate) will depend only on the inflation coefficients in the inflation equation, and in particular on their sum. We will assume that this is the case.

Prior to 1970, the sum of the lag coefficients was distinctly less than unity, indicating a stationary process for inflation, as argued at the time by Sargent (1971). By

[2] As in OLS, ALS autoregressive coefficients will tend to be biased, and in particular their sum will tend to be biased downwards, particularly with small sample sizes as the sum approaches unity. In ALS, it is in fact the effective sample size T_t, rather than the full sample size n that is relevant. Adjusting for this baias along the lines of e.g. Andrews and Chen (1994) goes beyond the scope of this paper.

1974, however, the sum of the lag coefficients had reached and even briefly exceeded unity. This phenomenon was noted already in the late 70s by Klein (1978). The three non-zero lags, whose relative size has moved around considerably, indicate that the simple Local Level Model of Muth (1960), which formalizes the Adaptive Expectations model of Cagan (1956) is insufficiently general to capture inflationary dynamics.

Figure 4 below shows the sum of the filter estimates of the coefficients on lagged unemployment in the inflation equation. The individual coefficients, which often have large magnitude and either sign, but are often locally insignificantly different from 0, are not shown. Although the net effect of U on inflation is small, it is consistently negative, and was larger before the early 1970s than afterwards.

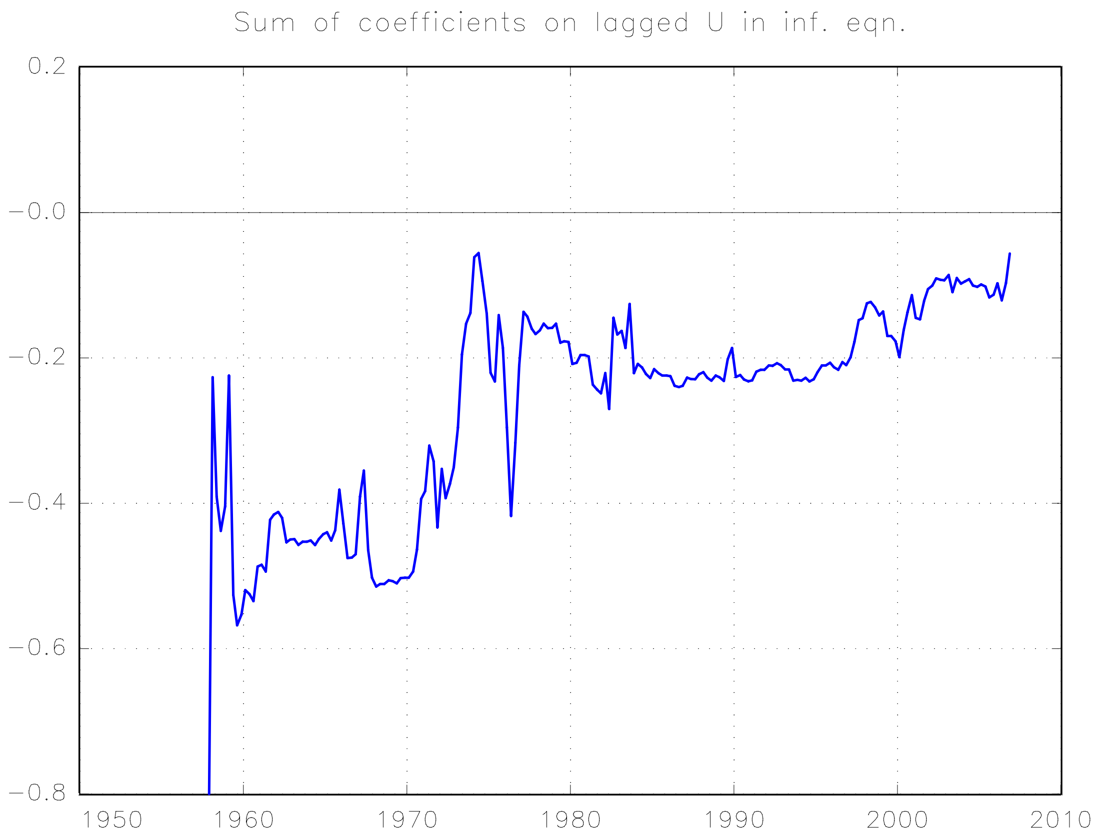

Figure 4

Figure 5 below shows actual inflation π_t and simulated expected inflation $\hat{\mathrm{E}}_t\,\pi_{t+1}$ computed from the filter estimates of the parameters. Figure 6 gives unexpected inflation, $\pi_t - \hat{\mathrm{E}}_{t-1}\,\pi_t$.[3]

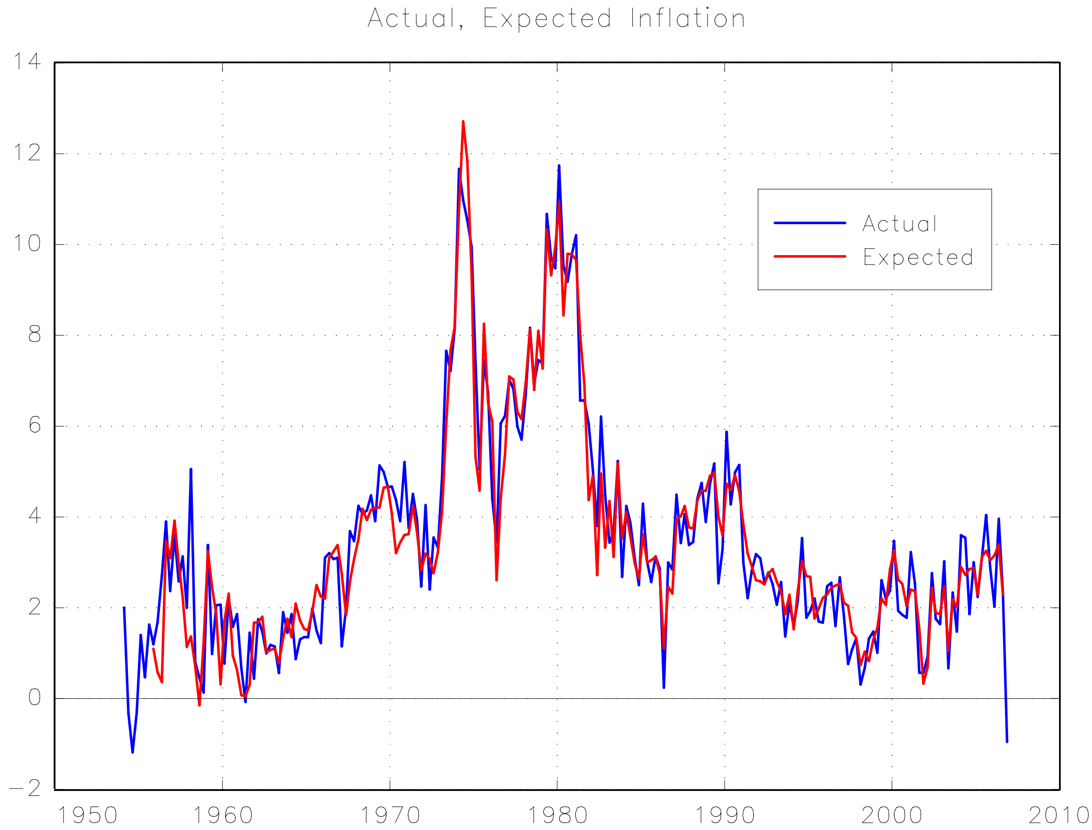

Figure 5

[3] There is clear visual evidence in Figure 6 of volatility clustering. It is anticipated that future versions of this paper will allow for this in this and the other two equations estimated here with a GARCH-ALS model, along the lines discussed in McCulloch (2006a). This may modify somewhat the standard error estimates and test statistics in the present paper.

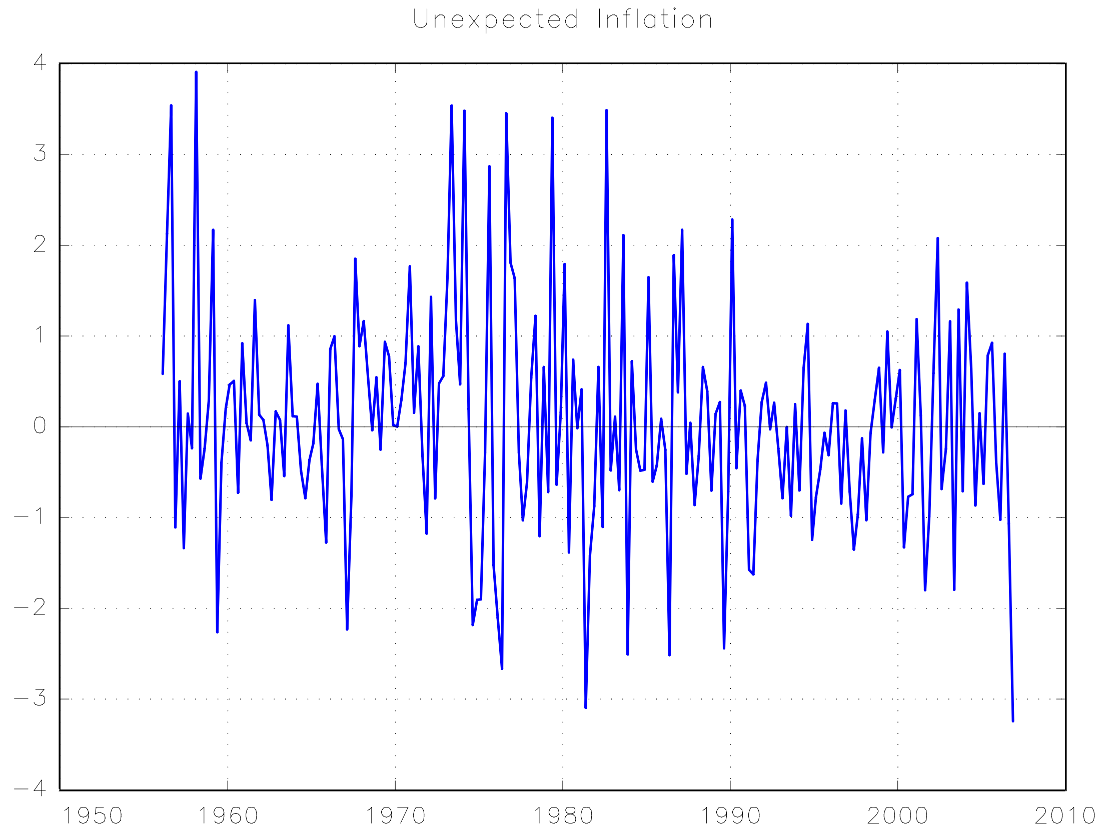

Figure 6

5. The Unemployment Equation

Since inflation is a non-stationary process, an unconstrained 2-equation VAR on inflation and unemployment would, with probability 1, imply that unemployment is also non-stationary. In order to permit unemployment to be stationary and therefore to have a well-defined natural rate, we instead model unemployment as a VAR in a constant and up to 4 lags of itself and of simulated unexpected inflation, as constructed from the model of Section 4 above. This system is in effect a 2-equation structural VAR , with the natural rate restriction imposed, that unemployment depends inflation only to the extent it is unanticipated.

Since 4 lags of the unexpected inflation variable are not available until 1957Q2, the unemployment equation begins with that date. The general-to-specific model selection procedure like that used for the inflation equation permitted a reduction of the equation for unemployment to 2 lags of itself, plus two lags of unexpected inflation. The ML estimates of the signal/noise parameter imply

$n = 199$

$T_\infty = 34.7$ (quarters)

$LR = 3.76$

The asymptotic effective sample size (34.7 quarters = 8.7 years) is similar to that obtained for the inflation equation. The LR statistic again well exceeds the 2.3 approximate 5% critical value, permitting rejection of parameter constancy.

By far the most important variables are the lags of inflation itself, shown in Figure 7 below. The first lag is consistently greater than unity, while the second lag is consistently negative, indicating the change in unemployment has considerable inertia to it. Since unexpected inflation is, by assumption, stationary with mean 0, the stationarity of unemployment depends on its own lag coefficients. These sum consistently to less than unity, which permits the construction of a natural rate estimate as the constant term, divided by 1 minus the sum of the unemployment lag coefficients.

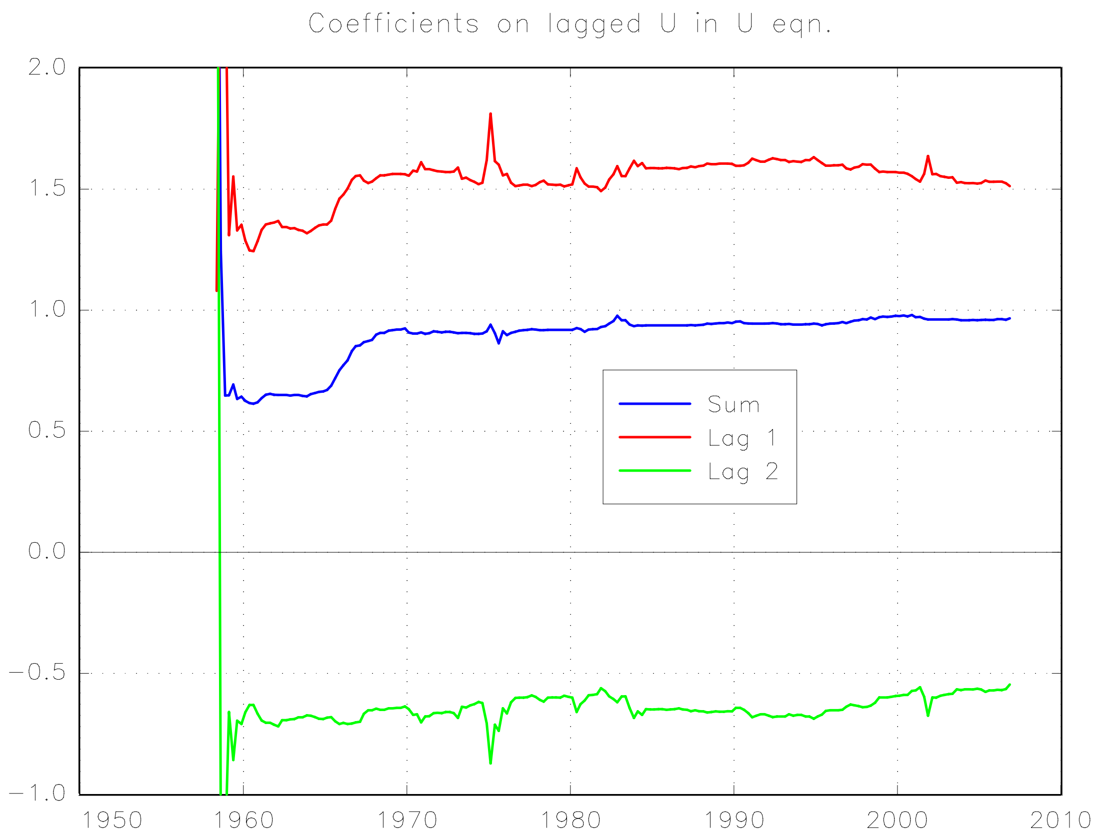

Figure 8

The unexpected inflation variable does not add much to lagged unemployment itself, since the second lag is only occasionally significant, while the first lag is never significant, and is only included because the second lag is. Nevertheless, the sum of the two coefficients, plotted in Figure 9 below, generally has the expected negative sign.

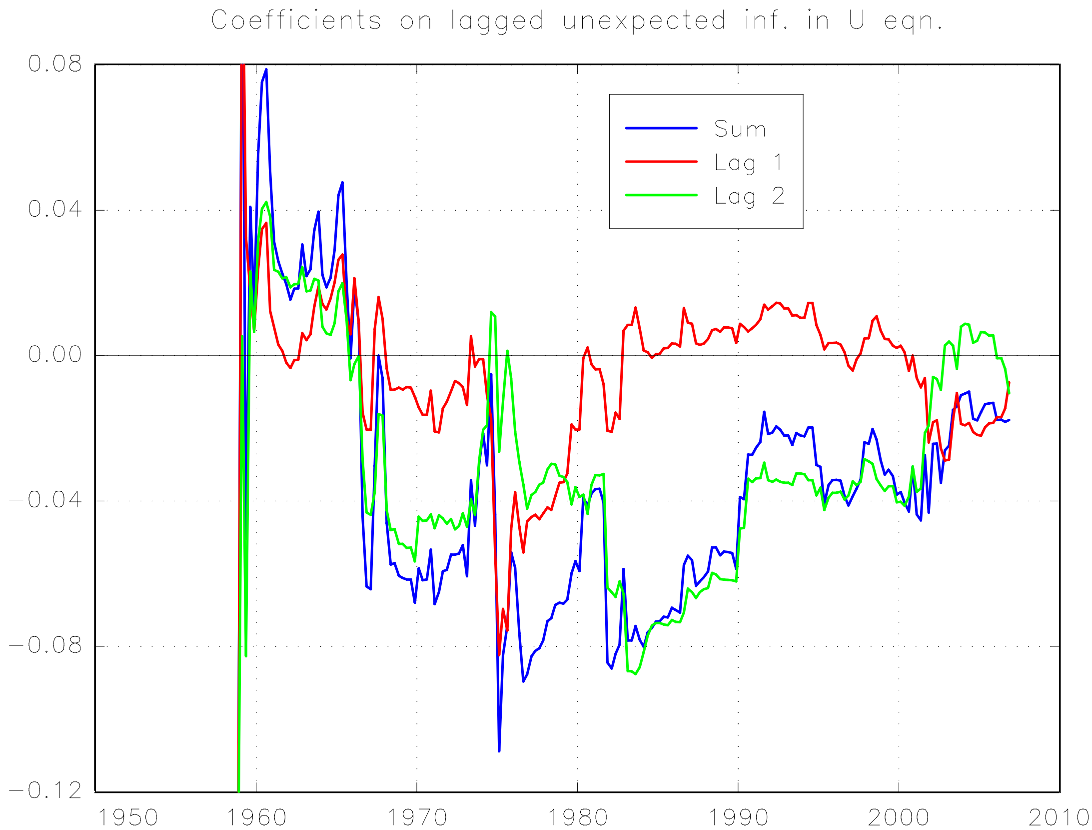

Figure 9

Since the unexpected inflation variable by assumption has mean 0, the natural unemployment rate can be computed as the constant term, divided by 1 minus the sum of the unemployment lag coefficients. Figure 10 below shows the filter estimate, based only on past unemployment data, alongside the smoother estimate, based on both past and future data, for comparison. The filter reflects what agents at the time could have perceived the natural rate to be, while the smoother is the best estimate, given the benefit of hindsight. Naturally the two estimates coincide at the last date observed. These estimates confirm the conclusion of Orphanides and Williams (2006) that the Fed generally underestimated the natural rate in the years leading up to 1980, and generally overestimated it afterwards.

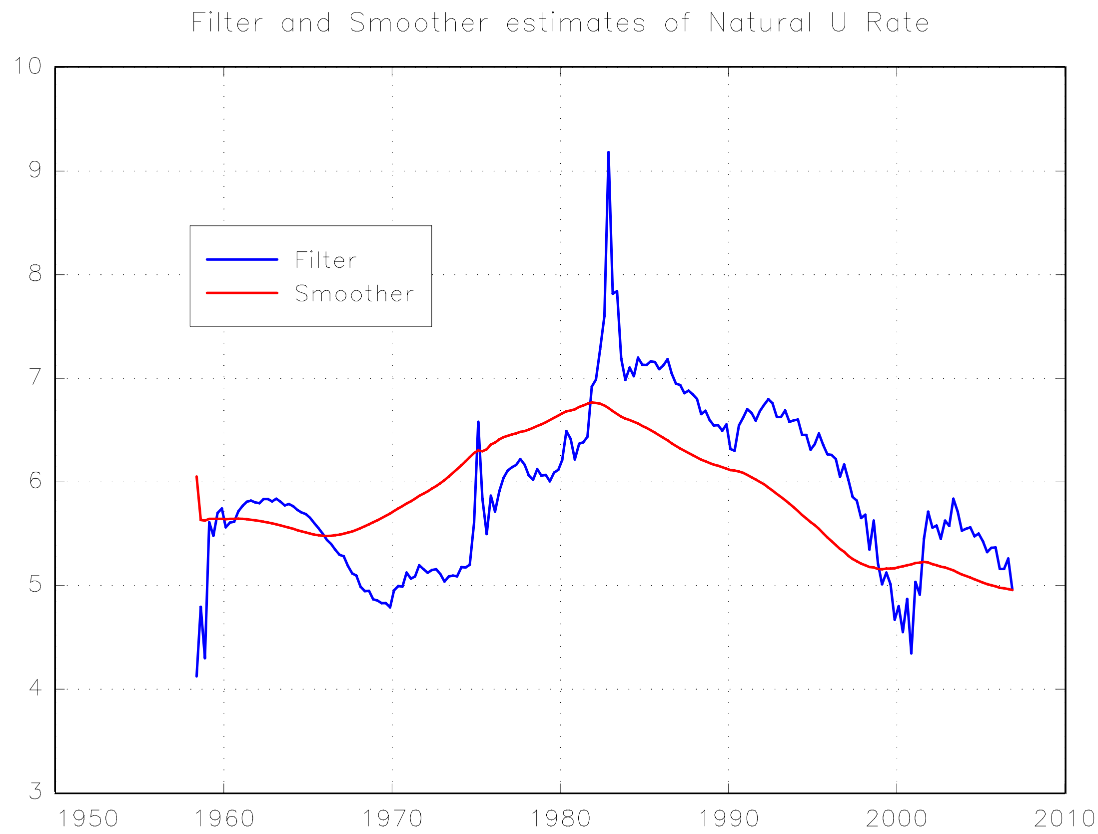

Figure 10

Figure 11 below shows the filter predictions of unemployment for the coming quarter, alongside actual unemployment. When unemployment is very high or low, the coming quarter's rate is generally predicted to be below or above the current rate, respectively, as the rate reverts gradually to the current natural rate. A few exceptions occur, however, because when the rate changes quickly, the inertia in the rate of change can induce it to continue briefly in the same direction before eventually being attracted to the natural rate.

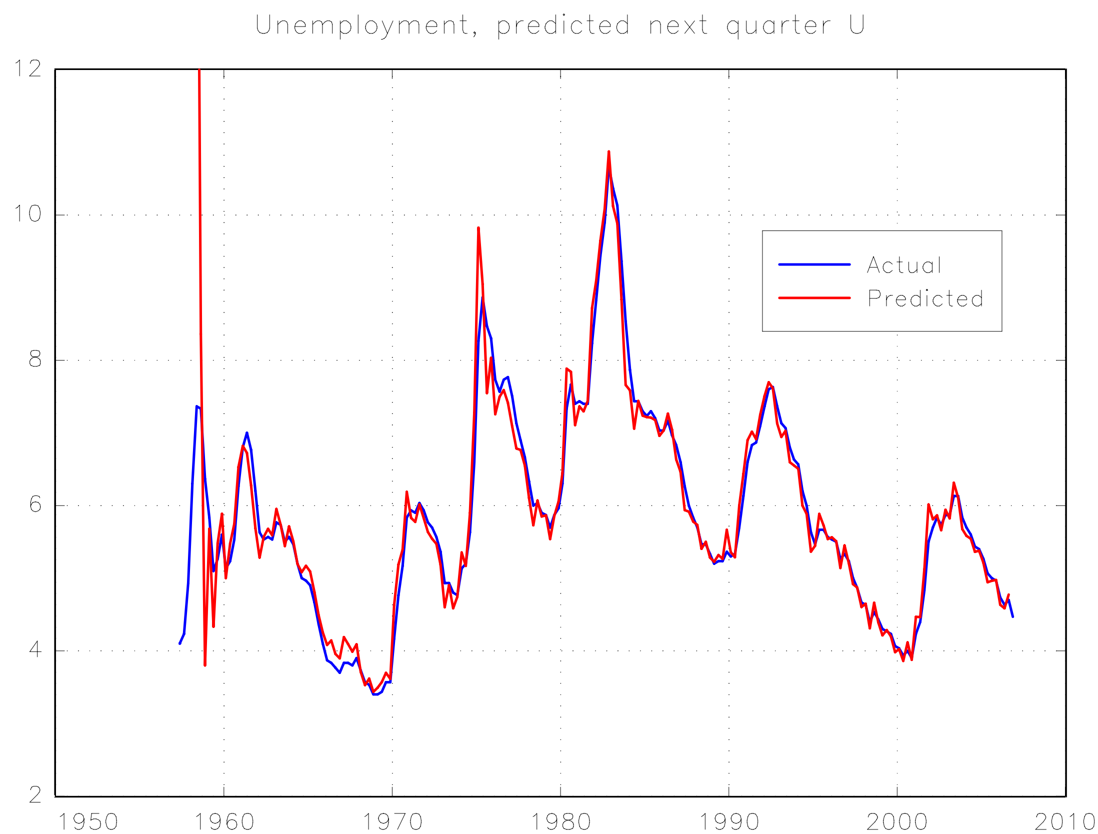

Figure 11

Figure 12 below shows the predicted next quarter unemployment gap, defined as the filter-predicted unemployment rate minus the filter-estimated natural rate, alongside the estimated current unemployment gap. It may be seen that there is not a great difference between the two series.

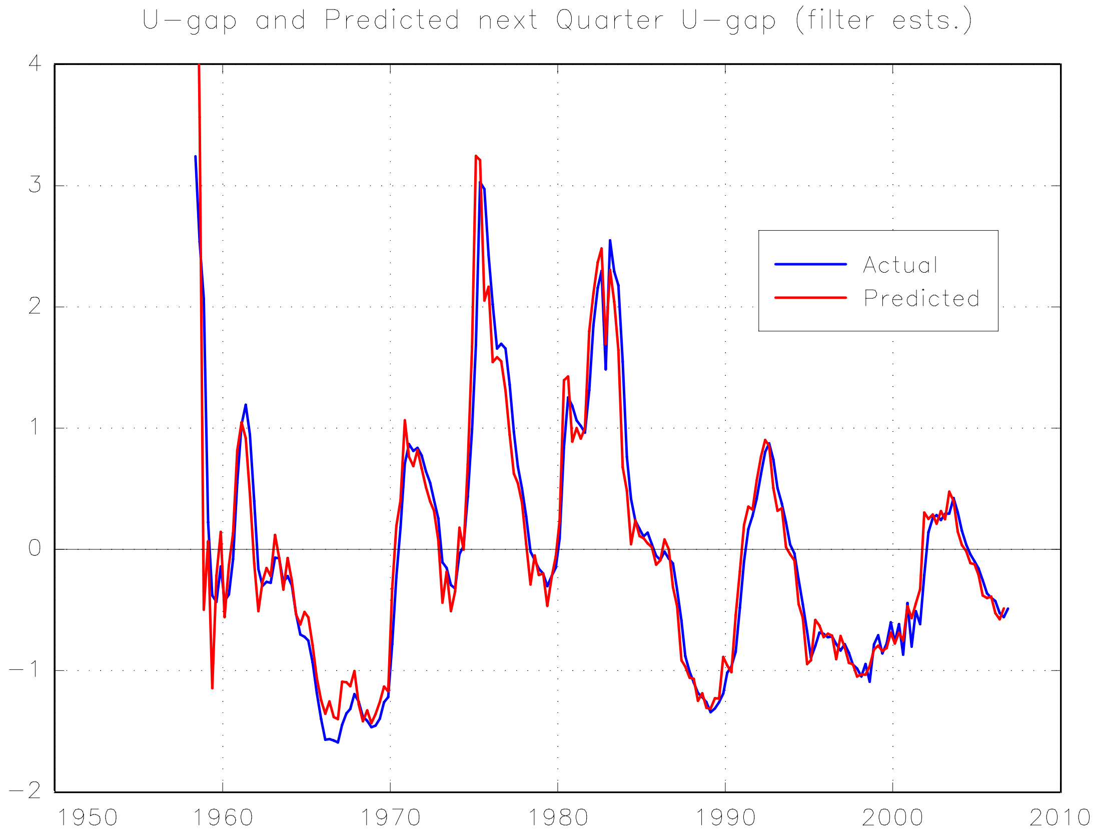

Figure 12

6. The Taylor Equation

We are now in a position to estimate a “forward-looking” version of the Fed’s Policy Response function, by using ALS to regress the Federal Funds rate on a constant, expected inflation for the coming quarter, the expected unemployment gap for the coming quarter, and up to 4 lags of the Funds rate itself, as follows:

$$R_t = (a_t + b_t \hat{E}_t \pi_{t+1} + c_t \hat{E}_t Ugap_{t+1})(1 - \sum_{j=1}^{p} d_{tj}) + \sum_{j=1}^{p} d_{tj} R_{t-j} \qquad (12)$$

Since we now want to make an historical assessment of what the Fed’s policy was on each date, under the assumption that it was not too dissimilar immediately before or immediately after, we will now focus on *smoother* estimates of the coefficients. In order to give the unemployment gap estimates a few quarters to stabilize, we do not attempt to begin this regression until 1960Q1. Figure 13 below shows the data used for this Taylor Rule regression.

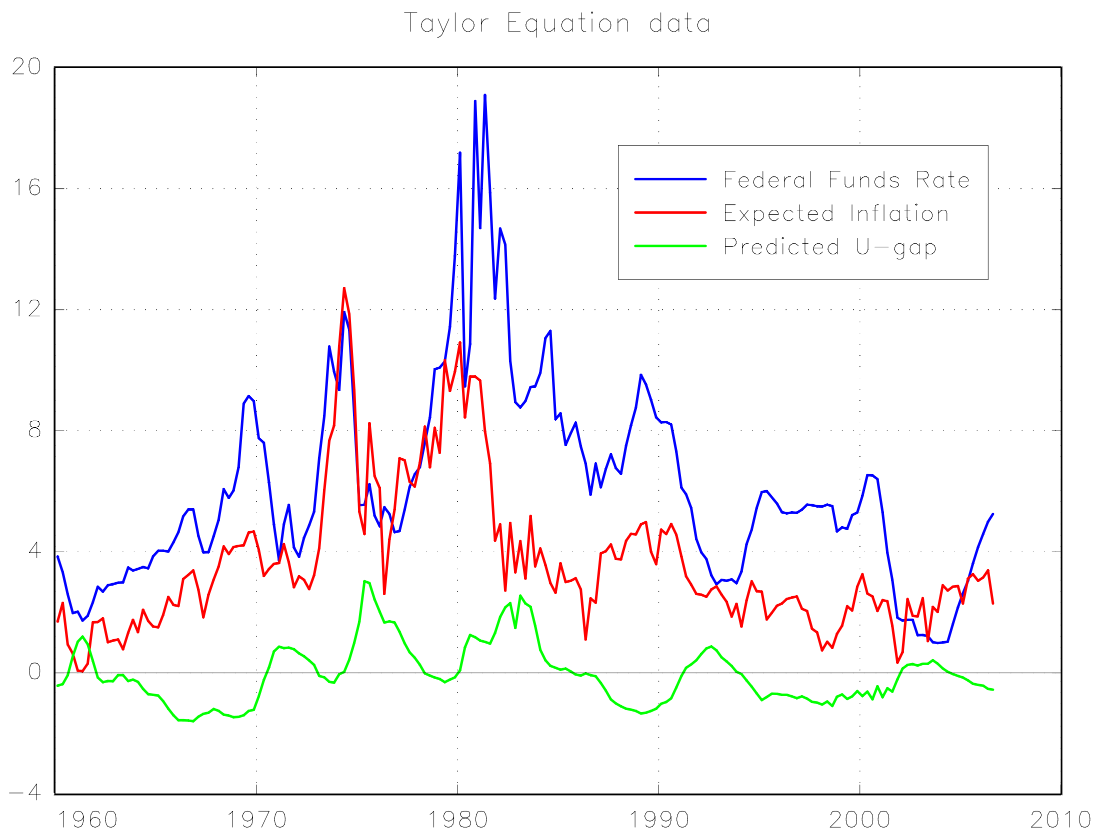

Figure 13

Our model selection rule permits eliminating the fourth lag of the Funds rate, but not the third lag, i.e. we use p = 3 in (12). The estimates are

n = 187

T_∞ = 20.7 (quarters)

LR = 11.58.

The asymptotic effective sample size, 20.7 quarters = 5.1 years, is substantially smaller than for the first two equations, though probably not significantly so, given the large range of uncertainty of these estimates. The LR statistic again easily permits rejection of parameter constancy.

Figure 14 below shows the smoother estimates of the three coefficients on lags of the Funds rate, along with their sum. It may be seen that there was a substantial increase in the persistence of the Funds rate between 1975 and 1980, from a sum of around 0.6 to a little over 0.8. Since 1995 there has been a gradual further increase, to nearly 0.9.

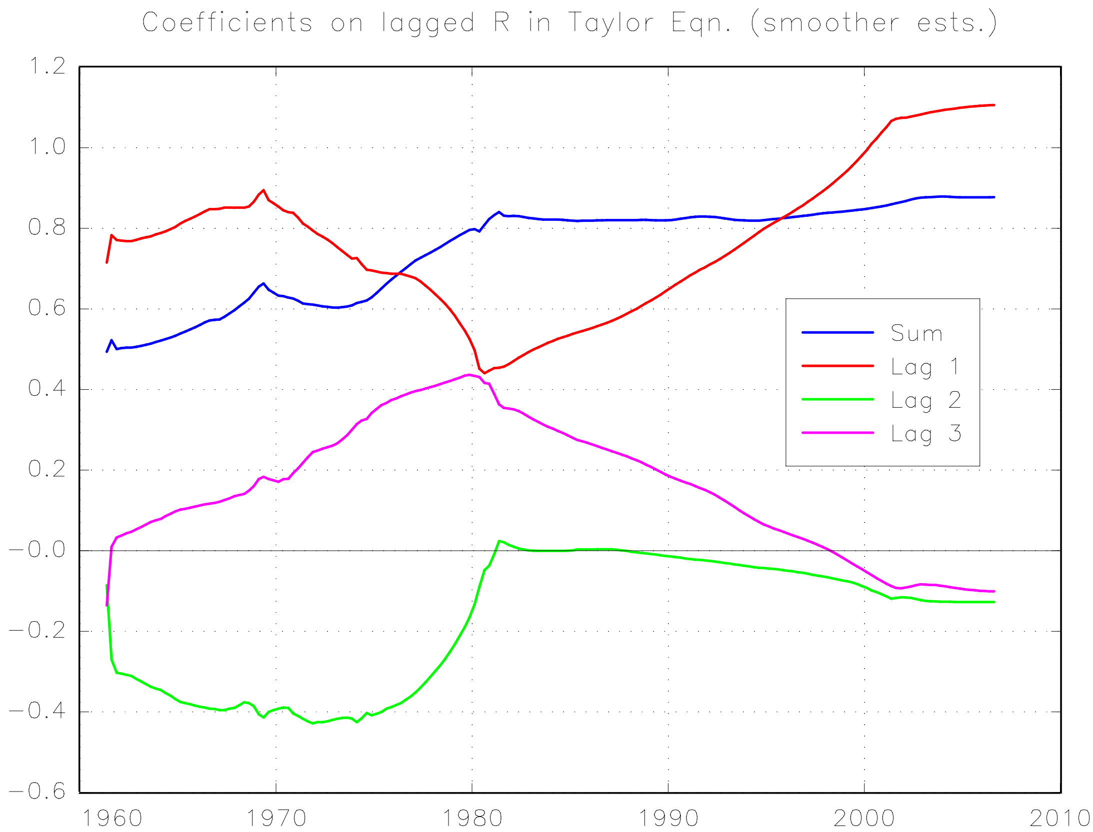

Figure 14

Figure 15 below show the smoother-generated estimates of the long-run Taylor Rule coefficients, computed as the smoother estimates of the corresponding raw coefficient, divided by 1 minus the sum of the lagged Funds rate coefficients. It may be seen that the coefficient on expected inflation was near or even below 1.0 prior to 1975, but rose sharply to 2.0 or even higher by 1980. It fell off to about 1.6 in the late 90s, but has been back above 1.8 since 2003.

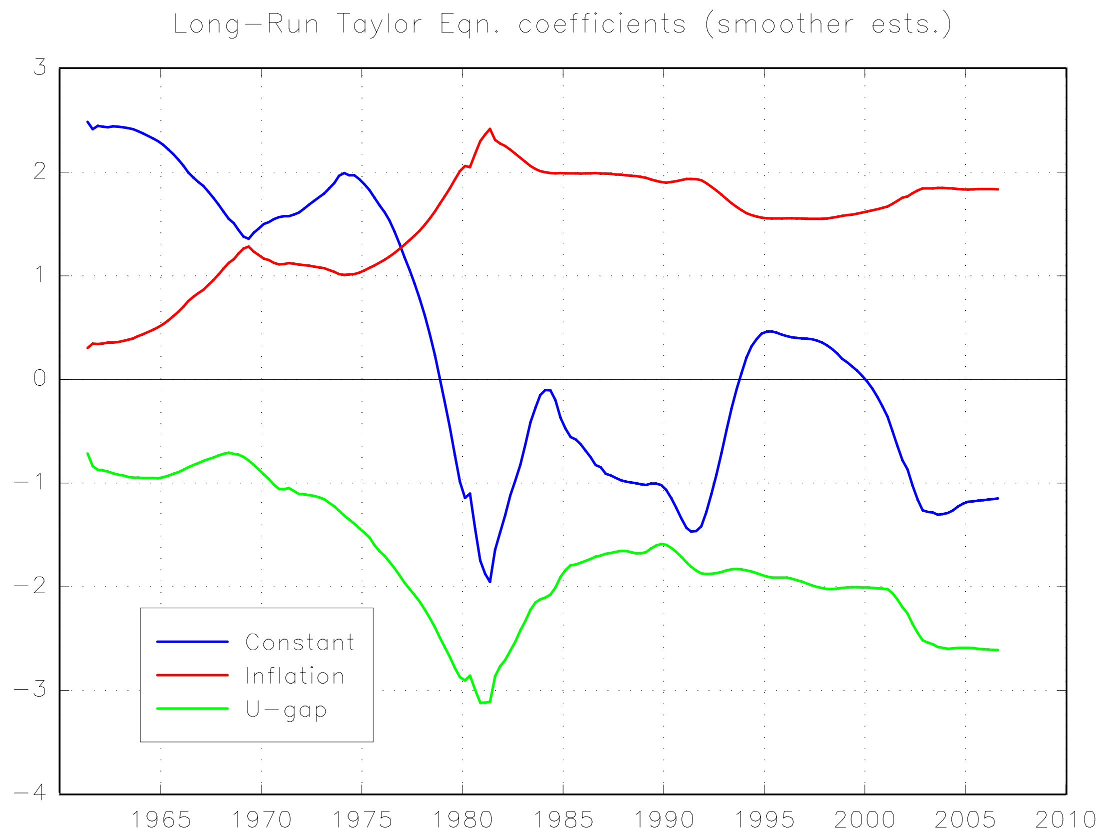

Figure 15

The long-run coefficient on the expected U gap is negative throughout. This would imply a positive coefficient on the output gap, if the U gap were used as a negative proxy for the y-gap. Since lagged U is already being used to form the inflation forecasts, the presence of the U-gap in the Taylor rule actually represents an activist policy, rather than just a supplementary term to help forecast inflation, as might be appropriate if past inflation rather than our bivariate forecast were used directly in the equation.

It may be seen that policymakers were actually softest on unemployment at the height of the Volcker credit crunch in 1980-82. In the late 80s and 90s, they were almost twice as activist as during the 60s, and since 2003 they have almost reached their 1982 peak.

If the equilibrium real interest rate that neither pushes inflation up nor down relative to expectations is constant at some level r*, the Taylor Rule (12) implies that in a steady state, in which inflation is fully expected, and $Ugap_t$ is 0, inflation will be constant at an equilibrium level of

$$\pi_t^* = (r^* - a_t)/(b_t - 1),$$

provided b_t is greater than 1. This level depends on the difference between r^* and the constant term a, and is greatly magnified if the inflation term b is close to unity.

Figure 16 below plots this equilibrium inflation rate for five values of *r** ranging from 0 to 4%. Prior to 1967, when b_t was actually less than 1, the equilibrium rate is undefined. Between 1967 and 1980, it is defined but generally offscale. Since 1980, it has been between -1% and 7%. Currently, it could be anywhere between 1% and 6%, depending on the unobserved value of *r**.

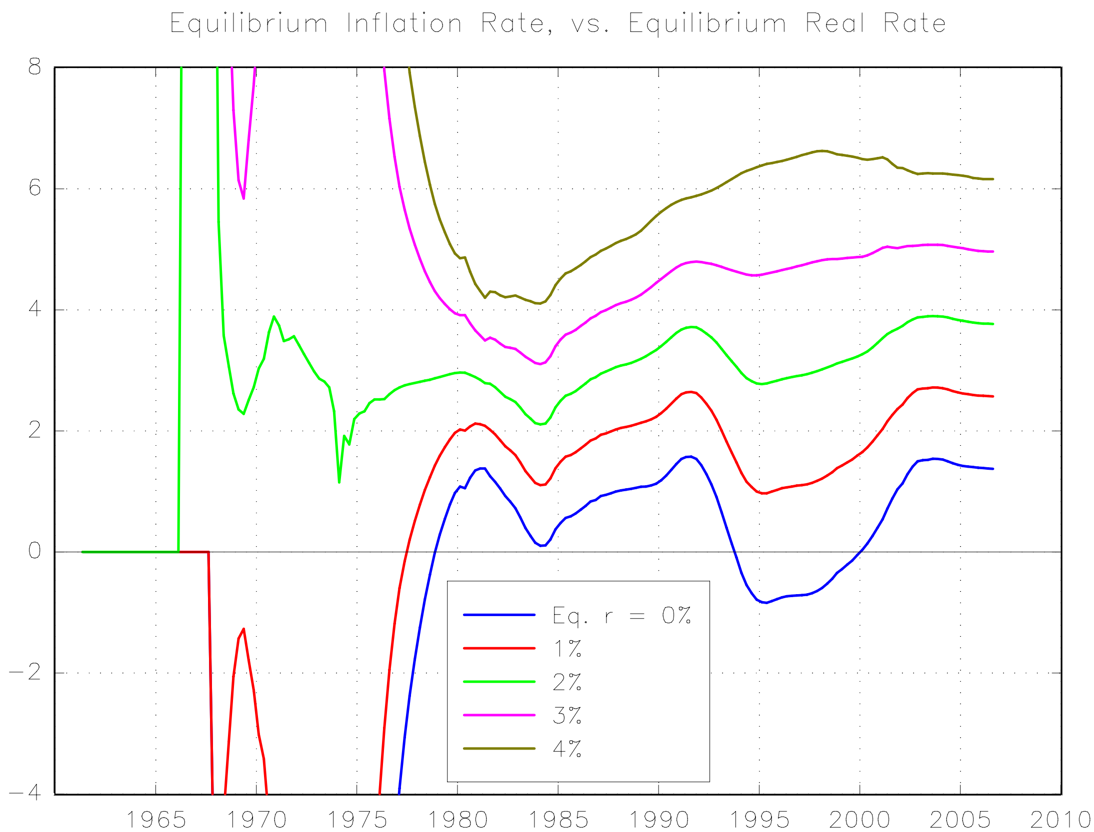

Figure 16

An intrinsic drawback of using an interest-rate rule to target inflation is that the required rule depends on the unknown level of r*. Unfortunately, the only way to determine *r** is to pursue a neutral monetary policy and see what the market comes up with. The Catch-22 is that in order to know whether a policy is neutral, we would first have to know r*!

A monetary aggregate growth rule has the analogous drawback that its outcome depends on the demand for money, which is not precisely observable either. In the end, the choice between an interest rate rule or a money growth rule depends not on which is "correct," but on the consideration of which uncertainty is less. In the spirit of Poole (1970), a fully optimized policy might in fact be a blend of an interest-rate rule and a money-growth rule.

It is not inconceivable that *r** varies with the macroeconomic conditions that give rise to variations in the unemployment gap. In particular, a deficiency of aggregate demand for current output (relative to future output) would tend to simultaneously cause an unanticipated fall in real interest rates (relative to earlier forward rates) and a surge in unemployment.[4] Even a purely inflation-targeting, non-activist Taylor Rule might therefore include the U-gap with a negative sign, in order to compensate for likely fluctuations in r^*.

Although we were careful to model expected inflation and unemployment as a structural VAR using potential lags of both in each equation, most of the explanatory power was in the own lags (on the variable in question) rather than the alien lags (on the other variable). It therefore would have made little difference if we had simply used univariate ALS models to forecast these two variables.

Furthermore, although we were careful, following CGG, to use a "forward-looking" Taylor Rule that employs current expectations of next quarter's values of inflation and the U-gap, the high degree of persistence of both series implies that, a "backward looking" Taylor equation that simply used current values of the inflation rate, either for the most recent quarter or averaged over the past 4 quarters, and of the unemployment gap would have given almost identical results.

7. Conclusion

Adaptive Least Squares successfully models a time-varying structural VAR reduced form for inflation and unemployment that permits us to construct numerical proxies for agents' real-time expectations of these variables. It then permits us to fit a "forward-looking" Taylor equation with time-varying coefficients.

The Likelihood Ratio statistic rejects the hypothesis of constant coefficients in all three equations. The more "fundamental" coefficients in the inflation and unemployment equations are estimated to turn over, in effect, every 9.7 or 8.6 years, respectively, while the policy-determined coefficients in the Taylor Rule change, in effect, every 5.1 years.

Smoother estimates of the Taylor Rule indicate that the coefficient on expected inflation rose from barely 1.0 to 2.0 or higher during 1975-1980. It fell to 1.6 in the 1990's, but has been nearly 2.0 since 2003. The response to the unemployment gap has been negative throughout, and was strongest near 1980 and since 2002. The equilibrium inflation rate consistent with the estimated Taylor Rule coefficients could be anywhere in the range of 1 to 6% since 2003, depending on the unobserved equilibrium real interest rate.

---

[4] See McCulloch (1981).

References

Andrews, Donald W.K., and Hong-Yuan Chen, 1994. "Approximately Median-Unbiased Estimation of Autoregressive Models." *Journal of Business & Economic Statistics* **12**: 187-204.

Cagan, Phillip, 1956. "The Monetary Dynamics of Hyperinflation," in Friedman, ed., *Studies in the Quantity Theory of Money*, University of Chicago Press.

Clarida, Richard, Jordi Galí and Mark Gertler, 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics* **115**: 147-180.

Evans, George W., and Seppo Honkapohja, 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.

Kalman, R.E., 1960. "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering, Transactions ASME Series D* **82**: 35-45.

Kim, Chang-Jin, and Charles Nelson, 2006. "Estimation of a Forward-Looking Monetary Policy Rule: A Time-Varying Parameter Model using Ex-Post Data," *Journal of Monetary Economics* **53** (Nov.): 1949-66.

Klein, Benjamin. "The Measurement of Long- and Short-Term Price Uncertainty: A Moving Regression Time Series Analysis," *Economic Inquiry* **16** (1978), 438-52.

Longbrake, Mark W., and J. Huston McCulloch, 2007. "Searching for the Output Gap: Economic Variable or Statistical Illusion?" 13th International Conference on Computing in Economics and Finance, Montreal.

McCulloch, J. Huston, 1981. "Misintermediation and Macroeconomic Fluctuations," *Journal of Monetary Economics* **8**: 103-115.

McCulloch, J. Huston, 2005a. "The Kalman Foundations of Adaptive Least Squares," Ohio State University working paper, Aug. (corr. 5/7/06), <http://www.econ.ohio-state.edu/jhm/papers/KalmanAL.pdf>

McCulloch, J. Huston, 2005b. GAUSS program ALS. Fits a linear regression by Adaptive Least Squares, estimating the signal/noise parameter by Maximum Likelihood, at <http://www.econ.ohio-state.edu/jhm/programs/ALS>.

Muth, John F., 1960. "Optimal Properties of Exponentially Weighted Forecasts," *J. of the American Statistical Assn.*, 299-306.

Orphanides, Athanasios, and John C. Williams, 2005. "The Decline of Activist Stabilization Policy: Natural Rate Misperceptions Learning, and Expectations," *Journal of Economic Dynamics and Control* (Nov.), 1927-50.

Poole, William, 1970. "Optimal Choice of Monetary Policy Instruments in a Simple Stochastic Macro Model" *Quarterly Journal of Economics* **84**: 197-216.

Sargent, Thomas J., 1971. "A Note on the 'Accelerationist' Controversy," *J. of Money, Credit and Banking* **3**: 721-5.