

Econometric Policy Evaluation and Inverse Control

Michael K. Salemi*
Department of Economics
CB# 3305, Gardner Hall
University of North Carolina
Chapel Hill, NC 27599-3305
Michael_Salemi@unc.edu

Extended Version
September, 2006

This version of the paper contains information referred to but not included in the version published in JMCB. To preserve readability and comparison of the two versions, I have included the supplemental information in several appendices.

JEL Categories
C32, C61, E31, E32, E52, E61

Abstract

The traditional approach to monetary policy evaluation entails a first step in which structural parameters are estimated and a second in which the performance of alternative policy rules is studied. This paper combines the two steps of the traditional approach into one by estimating a structural model subject to the restriction that the central bank chooses the policy rule that minimizes expected loss. The structure is a forward-looking New Keynesian model in which equilibrium values of output and inflation depend on expectations of future values of those variables. Analysis of U.S. data between 1965 and 2001 support the hypotheses that the sample contains two policy regimes, that the Fed placed far greater weight on stabilizing inflation after 1980, and that improvements in policy were available in both regimes. The unified approach also sharpens estimates of structural parameters.

*I have benefitted from discussion with Evan Anderson, Rouben Atoian, Larry Christiano, Richard Froyen, Hans Genberg, Gregory Givens, Søren Johansen, George Korenko, Peter Kugler, Catalin Stefanescu, George Tauchen, Christopher Waller, and Chao Wei. I thank participants at UNC-CH, the University of Basel, the European Institute, the European Monetary Forum Conference in Honor of Sir Alan Walters and the Graduate Institute for International Studies for opportunities to present earlier versions of this paper. Some of this project was completed while I was a visiting professor at the Graduate Institute. I also thank the referees and editor of this journal for their many helpful suggestions on earlier drafts.

Monetary policy rules are naturally amenable to modern econometric policy evaluation methods..... When using these methods, researchers first build a structural model of the economy, consisting of mathematical equations with estimated numerical parameter values. They then test out different rules by simulating the model stochastically.... One monetary policy rule is better than another...if the simulation results show better economic performance. (Taylor, 1998).

The two-step approach to policy evaluation, described by Taylor in the quoted passage, has been around for a long time. Indeed, Marschak describes a two-step approach to tax policy in Chapter 1 of Hood and Koopmans (1953). Fuhrer (1997a) provides a typical implementation of the two-step approach as it is currently employed. He estimates a New Keynesian model treating the coefficients of the policy rule as free parameters. Under the maintained hypothesis that the model is structural, he fixes parameters at their estimated values, computes an inflation-variance-output-variance policy frontier by varying policy-rule coefficients, and finds that the observed policy rule achieved output and inflation variances close to the frontier. If the model is not structural, conclusions about alternative policies are suspect (Fuhrer, 1997b).

My paper unifies the two-step approach by estimating a structural model subject to the testable restriction that the policy rule minimizes a specified central bank loss function. Estimation under the unified approach searches over the parameters of a model and loss function for values that reconcile a policy rule that fits data with one that minimizes expected loss.

Combining the two steps of the policy evaluation program yields several benefits. First, the unified approach provides a test of the hypothesis that the policy rule minimizes expected loss. Second, it provides direct estimates of the parameters of the central bank loss function and, in this sense, is an example of inverse control. Third, imposing restrictions in the course of estimation can sharpen estimates of the structural parameters as shown in Section 3. The cost of the unified approach is the greater computational burden that results from nesting the optimal-policy algorithm of the central bank within the estimation algorithm of the researcher.

My paper is not the first application of inverse control to monetary policy. Salemi (1995) estimates a VAR subject to the restriction that the policy-variable equation minimizes expected loss. He finds that the Fed placed greater weight on stabilizing output in two regimes prior to 1979 and greater weight on stabilizing inflation between 1982 and 1992 and cannot reject the hypothesis that the Fed policy rule minimized expected loss. Korenko (1998) uses inverse control to study Federal Reserve Policy with a version of Dornbusch's overshooting model similar to that of Papell (1989). He treats 1973:04 through 1995:12 as a single policy regime and finds that the Fed placed much higher weight on output stabilization than on other objectives. Dennis (2002) estimates the parameters of a backward-

looking model composed of an IS equation, a Phillips curve, and a Federal Reserve loss function subject to the restriction that the coefficients of the reaction function minimize loss. He finds for both Martin-Burns and Volcker-Greenspan regimes that the Fed placed more weight on stabilizing output than on stabilizing inflation and even more weight on stabilizing interest rate changes. Favero and Ravelli (2003) estimate loss-function parameters and a backward-looking model and find, in contrast to Dennis, that the Fed placed more weight on stabilizing inflation than on stabilizing output in both regimes.

Unlike Dennis and Favero-Ravelli, my paper uses a forward-looking structure that allows for the kind of private responses to changes in policy that the Lucas critique warns about. The use of a forward-looking structure complicates policy analysis by requiring re-computation of the state transition equation for every contemplated change in the policy rule.

My chief findings are: first, that since 1965 the Fed consistently emphasized inflation stabilization over output stabilization; second, that it especially did so during the Volcker-Greenspan regime; and third, that there was significant room for improvement in the performance of policy. I also show that the unified approach sharpens the estimates of important structural parameters. Section 1 describes the structural model and explains how to compute its reduced form and the expected loss associated with a policy rule. Section 2 describes the data used in the analysis—quarterly U.S. data for 1965:I through 2004:IV. Section 3 presents estimates of the model and discusses what they imply about Federal Reserve policy. Section 4 concludes. Appendix A provides a description of the algorithms used to generate the paper's findings.

1. A FORWARD-LOOKING ECONOMIC MODEL

I consider a central bank that wishes to stabilize the time paths of the output and inflation by controlling the interest rate. Let y be the output gap, the difference between output and its long run growth path. Stabilizing output means keeping y close to zero. Stabilizing inflation and the interest rate means keeping them close to target paths. Let p and r be differences of inflation and the interest rate from target values.

The model for y , p and r is composed of an IS curve, a Phillips curve, and a rule for the short-term interest rate. While the model is consistent with a representative-agent general-equilibrium model, none of the issues taken up in this paper concern the deeper structural parameters. Because the form of the model I consider is consistent with more than one representative agent model, I focus on the expectations-augmented linear difference equations that describe local movements of y , p , and r .

$$\begin{aligned}
(1) \quad & y_t = \lambda E_t y_{t+1} + a_1 y_{t-1} + a_2 y_{t-2} - b(r_t - E_t p_{t+1}) + u_t \\
(2) \quad & p_t = \beta y_t + \alpha_1 E_t p_{t+1} + \alpha_2 p_{t-1} + v_t \\
(3) \quad & r_t = \theta_1 y_{t-1} + \theta_2 p_{t-1} + \theta_3 r_{t-1} + \theta_4 y_{t-2} + w_t .
\end{aligned}$$

Equation (1) can be obtained by combining the linearized Euler equation that characterizes a representative household's optimal choice between consumption and saving and the market clearing condition for output. As explained by Clarida, Gali, and Gertler (1999), the presence of expected future output in the IS equation results from the desire of households to smooth consumption. When households expect higher consumption in the future, they consume more in the present raising aggregate demand and introducing a positive association between current and expected future output. The presence of lagged output in the IS equation can be explained by habit persistence and delays between decision making and consumption. Woodford (1996), Rotemberg and Woodford (1997) and Bernanke, Gertler, and Gilchrist (1998) provide details and Svensson (2000) adapts the model to an open economy.

Equation (2) is a Phillips curve. If the coefficient on lagged inflation is zero, (2) is a new Phillips curve as defined by Gali and Gertler (1999), Clarida, Gali, and Gertler (2000), Svensson (2000) and many others. The foundation for the new Phillips curve is a model where monopolistically competitive firms adjust their prices on a staggered basis (Calvo, 1983). When it can, a firm adjusts its price to maximize expected profits taking into account the restriction it faces on future price adjustments and expected future prices of its competitors. With Calvo's approach, the inflation rate depends on the representative firm's marginal cost of production and the expected future inflation rate. The new Phillips curve results when the output gap is used as a proxy for marginal cost.

If the coefficient on lagged inflation is not zero, (2) is a Phillips curve modified to account for inflation inertia. A variety of foundations have been used to account for the presence of lagged inflation in the Phillips curve. Gali and Gertler extend the Calvo model to allow for a subset of firms that set their prices at time t equal to the average price set by their competitors at time $t-1$ with an adjustment for intervening inflation. In Christiano, Eichenbaum and Evans (2003), all firms are Calvo price setters but those that cannot re-optimize in the current period index their price to lagged inflation. In Rotemberg and Woodford (1997), a fraction of the firms that re-optimize in the current period must wait until the next period to charge their new price. Lagged inflation may appear in the Phillips curve to account for serial correlation in supply shocks as assumed by Clarida, Gali, and Gertler. Lagged inflation appears in the Phillips curve of Fuhrer and Moore, and others, as the result of wage contracts.

Equation (3) is the policy rule or reaction function that explains how the central bank sets the short term interest rate¹. It is essentially the same as (9) of Fuhrer and Moore (1995a) or (4) of Fuhrer (1997a). The policy rule permits the central bank to react to all the variables in the state vector, but only to lagged values of those variables. McCallum (1997, p.356) argues that central bank does not know current-quarter output when it sets the interest rate. The coefficients of the reaction function are assumed to be fixed throughout a policy regime so that policy is time consistent (Clarida, Gali, and Gertler, 1998 - 1999). For now, I defer discussion of how the coefficients of the policy rule are chosen and turn attention to solution of the model.

I compute the model's reduced form with the approach of Blanchard and Kahn (1980). Defining the state vector to be $\mathbf{X}_t = (y_t, p_t, r_t, y_{t-1})'$ equations (1) - (3) may be written as:

$$(4) \quad \begin{bmatrix} \mathbf{X}_t \\ \mathbf{E}_t \mathbf{y}_{t+1} \\ \mathbf{E}_t \mathbf{p}_{t+1} \end{bmatrix} = \mathbf{B} \begin{bmatrix} \mathbf{X}_{t-1} \\ \mathbf{y}_t \\ \mathbf{p}_t \end{bmatrix} + \mathbf{D} \begin{bmatrix} \mathbf{u}_t \\ \mathbf{v}_t \\ \mathbf{w}_t \end{bmatrix}$$

where \mathbf{B} and \mathbf{D} are (6x6) and (6x3) matrices the elements of which are non-linear functions of the parameters of equations (1) - (3). If the number of unstable eigenvalues of \mathbf{B} equals the number of forward-looking variables in (4), then the unique saddle path for the model given by

$$(5) \quad \mathbf{X}_t = \mathbf{G} \mathbf{X}_{t-1} + \boldsymbol{\phi}_t$$

where \mathbf{G} is a (4x4) matrix with elements that depend on the eigenvectors of \mathbf{B} and where \mathbf{N}_t is the (4x1) vector of reduced-form errors with typical element $\mathbf{N}_{j,t}$ and with $\mathbf{N}_{4,t} = 0$. Let \mathbf{S} be the (4x4) covariance matrix of \mathbf{N}_t . The upper-left (3x3) block of \mathbf{S} is the covariance matrix of $\mathbf{N}_{j,t}$, $j=1,2,3$. The remaining elements of \mathbf{S} equal zero. Let $\mathbf{C} = [C_{ij}]$ be the (3x3) upper-triangular Choleski decomposition of the non-zero block of \mathbf{S} . Because $\mathbf{N}_{j,t}$ are linear combinations of serially uncorrelated structural errors, \mathbf{N}_t is serially uncorrelated².

The Optimal Simple Monetary Policy Rule

¹Hansen and Sargent (1980) explain why an error may be present in a policy rule. An adaptation of their argument to the current setting is given by Salemi (p. 421).

²A detailed description of the solution is presented in Appendix B. In the course of estimation, I imposed the saddle-path restriction via a penalty function.

I assume that the central bank chooses \mathbf{z}_1 through \mathbf{z}_4 , the coefficients of the policy rule, to minimize expected loss (7):

$$(6) \quad \Lambda = E_t \sum_{j=1}^{\infty} \delta^j X'_{t+j} W X_{t+j}$$

where W is a (4x4) matrix of non-negative weights that determine the relative importance to the central bank of its various stabilization objectives and δ is the central bank's time rate of discount. For the baseline model, W is diagonal with non-zero elements W_y , W_p , and W_r , the weights associated with stabilizing output, inflation, and the interest rate. To check robustness of the baseline results, I consider cases where the bank wishes to stabilize changes in (rather than the level of) the interest rate and the growth rate (rather than the level) of output and modify W appropriately. I also consider a case in which all the elements of W are treated as free parameters. Because only the relative sizes of weights are identified, I normalize W_p to 1.0. For simplicity, I fix δ at 0.99³. In what follows, a policy regime is a set of loss-function weights and the reaction-function coefficients that minimize expected loss given those weights. Soderlind (1999) calls a policy rule with coefficients chosen to minimize expected loss an "optimal simple rule."

Rotemberg and Woodford show that, under certain conditions, the expected lifetime utility of a representative household is proportional to a loss function like (6) and the elements of W are non-linear functions of deep parameters of the model. If the central bank chooses the policy that maximizes representative-household utility, the parameters of the policy rule are likewise non-linear functions of the model parameters.

In contrast, I treat the elements of W as free parameters that carry information about the objectives of central bank policy. I take this approach for two reasons. First, I treat (1)–(3) as the structure of the economy rather than taking a stand on the specific model that gives rise to (1)–(3). Second, as will be clear from the empirical analysis, stabilization of the interest rate appears to be an important objective of U.S. monetary policy. But maximizing representative household utility does not require stabilizing the interest rate.

Why does it make sense to assume that the central bank commits to an optimal simple rule? The benefits of commitment are well known. Policy effectiveness depends on public expectations of future policy (Kydland and Prescott, 1977) and policy is more effective when its future course is predictable. Thus, commitment to a rule permits the central bank to distribute "policy medicine" over time. For

³Preliminary tests showed that estimation results were robust to varying δ between .975 and .995.

example, if a central bank wishes to offset inflation that will result from a supply shock, it can raise interest rates moderately provided that it maintains higher rates for a period of time. Lacking commitment, a higher initial rate increase is necessary because the public doubts that the central bank will sustain the increase.

The optimal simple rule is not the optimal commitment policy. The optimal commitment policy is the state-contingent plan for the interest rate that minimizes expected loss subject to the state transition equation implied by the economy's structure and the initial state vector. Soderlind (1999) shows that a state-space representation for the optimal commitment plan exists provided the state vector is augmented to include the co-state variables from the optimization. But writing the optimal commitment plan as a function only of lagged values of output, inflation, and the interest rate gives an equation for the interest rate that depends on the entire history of those variables.

I assume that the central bank follows an optimal simple rule rather than the optimal commitment policy for two reasons. First, it is not feasible for a central bank to provide the public in advance with a complete listing of relevant contingencies (Woodford, 2002) and it is thus difficult for the public to verify that the central bank is following the optimal commitment policy rather than engaging in discretion. It is easy for the public to ascertain whether or not the central bank is following a simple rule. Second, Atoian, Givens, and Salemi (2003) study the relative performance of optimal simple rules and optimal commitment policies in the setting of a structural model identical to (1)–(3). They find that when inflation stability is the dominant objective of policy, loss under the optimal simple rule is nearly as small as loss under optimal commitment.

Svensson (2003) argues that commitment to a simple policy rule is an unattractive hypothesis because the simple rule allows no room for policy-maker judgements and extra-model information, because the simple rule may be a good approximation to optimal policy in normal circumstances but sub-optimal in unusual ones, and because the simple rule does not allow for the arrival of new information about economic mechanisms or shocks. Svensson believes that optimal policy is better described as commitment to a “targeting rule” that “specifies operational objectives, that is, the target variables, the targets, and the loss function to be minimized” (p. 448).

The approach of this paper meets some, but not all, of Svensson's objections. In this paper, the objectives, the target variables, and the targets are set out explicitly as Svensson recommends. The variables and coefficients that appear on the right-hand-side of the policy rule are unrestricted (save that the variables must be elements of the economy's state vector). Coefficient values are endogenous and chosen to minimize expected loss. The rule is not assumed to hold exactly and one interpretation of the error term, w_t , is that it includes idiosyncratic wisdom of the policy authority. The form of the policy

rule is also consistent with Svensson's idea that policy makers respond to forecasts of target variables since those forecasts lie in the space spanned by the state vector. On the other hand, I make no attempt to allow for learning by the policy maker or to determine whether "unusual events" help explain the evolution of the interest rate.

Computing the Optimal Simple Rule

In the standard linear regulator problem, the state transition equation is invariant with respect to the policy rule and one may compute loss-minimizing policy-rule coefficients by iterating the Matrix Riccati equations to convergence (Sargent, 1987 and McGrattan, 1990).⁴ Finding the coefficients for the optimal rule when structural equations are forward looking is a more complex problem.

From (5), the moving average representation for the state vector X_t is $(I - GL)^{-1} N_t$ where L is the lag operator. It follows that expected loss may be decomposed into two parts:

$$(7) \quad \begin{aligned} \Lambda &= \text{trace} \left[W \sum_{k=1}^{\infty} \delta^k E_t(X_{t+k} X_{t+k}') \right] \\ &= \text{trace} \left[W (M_t + N_t) \right] \end{aligned}$$

where $M_t = \sum_{k=1}^{\infty} \delta^k E_t(X_{t+k} - E_t X_{t+k})(X_{t+k} - E_t X_{t+k})'$ and $N_t = \sum_{k=1}^{\infty} \delta^k (E_t X_{t+k})(E_t X_{t+k})'$.

\mathcal{Q}_t is the discounted sum of forecast error variances of X computed when policy is set. N_t is the discounted sum of quadratic terms in expected departures of X from its target. I assume that the policy authority ignores N_t when it chooses its policy rule either because the economy begins on its target path or because the authority follows the timeless perspective introduced by Woodford (1999). Because N_t is serially uncorrelated, $E_t(X_{t+k} - E_t X_{t+k})(X_{t+k} - E_t X_{t+k})' = \Omega + G\Omega G' + \dots + G^{k-1}\Omega(G^{k-1})'$ and

⁴Backus and Driffill (1985) show how to adapt the Riccati equations to a model with endogenous, forward looking expectations and provide an alternative to the procedure used herein.

$$\begin{aligned}
\mathbf{M}_t &= \delta \mathbf{\Omega} + \delta^2 (\mathbf{\Omega} + \mathbf{G} \mathbf{\Omega} \mathbf{G}') + \dots + \delta^{k+1} (\mathbf{\Omega} + \mathbf{G} \mathbf{\Omega} \mathbf{G}' + \dots + \mathbf{G}^{k-1} \mathbf{\Omega} (\mathbf{G}^{k-1})') + \dots \\
(8) \quad &= \frac{\delta}{(1-\delta)} [\mathbf{\Omega} + \delta \mathbf{G} \mathbf{\Omega} \mathbf{G}' + \delta^2 \mathbf{G}^2 \mathbf{\Omega} (\mathbf{G}^2)' + \dots]
\end{aligned}$$

To compute expected loss associated with a set of policy rule coefficients, I compute the sequence of partial sums in (8) until it converges.⁵ To find the optimal policy, I search for the policy-rule coefficients that minimize trace $(\mathbf{W} \mathbf{M}_t)$. Equation (8) implies that the loss-minimizing policy-rule coefficients depend, through \mathbf{G} , on the parameters of (1)-(3). Because agents are forward looking, \mathbf{G} depends on the policy-rule coefficients. Thus, the parameters of (1)-(3) and the policy-rule coefficients must be computed simultaneously.

Maximum Likelihood Estimation: SURF and SORF

Collect the structural parameters into vector $\mathbf{S} = (\mathbf{a}_1, \mathbf{a}_2, \mathbf{b}, \mathbf{c}_1, \mathbf{c}_2, \mathbf{c}_{1,1}, \dots, \mathbf{c}_{3,3})$, collect the coefficients of the policy rule into vector $\mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3, \mathbf{z}_4)$ and collect the free parameters of the expected loss function into vector $\mathbf{W} = (\mathbf{W}_y, \mathbf{W}_r)$.

The first step of the two-step approach to econometric policy evaluation estimates \mathbf{S} and \mathbf{z} while treating \mathbf{z} as unrestricted. I refer to this estimation strategy as SURF: Structural estimation with an Unrestricted Reaction Function. In Section 3, I report SURF estimates of \mathbf{S} and \mathbf{z} and analyze policy with these estimates.

The hypothesis that the policy-rule coefficients minimize expected loss induces a function, $\mathbf{z} = \mathbf{g}(\mathbf{S}, \mathbf{W})$. The unified approach estimates \mathbf{S} and \mathbf{W} subject to the restriction, $\mathbf{z} = \mathbf{g}(\mathbf{S}, \mathbf{W})$. I refer to this estimation strategy as SORF: Structural estimation with an Optimal Reaction Function. In Section 3, I report SORF estimates of \mathbf{S} and \mathbf{W} and the implied value of \mathbf{z} . Appendix A provides an explanation of the algorithm I use to compute SORF estimates.

2. DATA DEFINITIONS AND PRELIMINARY EMPIRICAL ANALYSIS

This and the next section describe what happens when I take my model and policy hypothesis to U.S. quarterly data for 1965:I through 2001:IV. To use a new Keynesian model for policy analysis in 1965 is to assume that the Federal Reserve possessed an understanding of the forward-looking nature of

⁵Anderson *et al* (1996) discuss alternative ways to compute equations like (8).

economic behavior a decade before Lucas published his critique. While it might be preferable to assume that the Fed learned as time passed, I follow standard practice and assume that the Fed knew the model and its parameters.

A first step is choosing measures for y , p , and r . Except for population, the data come from *Fred*, the St. Louis Federal Research Bank data source. Quarterly population numbers are interpolations of annual values found in the 2002 *Economic Report of the President*. My output measure is real GDP (in 1996 chained dollars) divided by civilian population. I estimate the long run growth path of per capita output by fitting a trend to the natural log of real GDP per capita (Figure 1a). Variable y , the output gap, is the difference between the natural logarithm of real GDP per capita and its trend.

Measuring target output with a trend may avoid pitfalls that attach to more complicated models of potential output. Ghysels, Swanson and Callan (2002), Orphanides (2000) and Orphanides and Williams (2002) argue that monetary policy should be analyzed in the context of contemporaneous central bank estimates of potential output. They caution that if the Fed stimulated the economy in response to slow output growth only to conclude later that potential output growth had also slowed, a researcher working with revised data would misunderstand Fed motives.

I assume that the potential output evolves like a constant-coefficient time trend. Assuming that potential output grows smoothly does not automatically imply that the Fed estimates it accurately since the Fed estimate of the trend coefficient may vary through time. I assume that the Fed knows the true value of the trend coefficient which is constant during the sample period. If policy makers use more complicated models of potential output, my estimate of y may differ from theirs.

The inflation rate is the annualized percentage change in the chained GDP deflator. Figure 1b shows that the inflation rate trended up between 1965 and 1980, trended sharply downward between 1980 and 1985, and trended gradually downward between 1986 and 2001. The interest rate is the annualized secondary market yield of three-month treasury bills. Figure 1c shows that the interest rate followed a trend path similar to that of inflation. To estimate target values for the inflation rate and interest rate, I fit continuous, piecewise-linear trends that allow for trend coefficient changes in 1980:I and 1986:I but constrain trend coefficients to be the same in the inflation rate and interest rate equations. The resulting target paths are displayed in the figures. Variables p and r are the residuals from the two trend regressions.

My target-path assumptions imply that the Fed acclimated itself to rising inflation during the 1960's and 1970's. For example, the Fed would have judged an inflation rate of four percent to be above target in 1967, at target in 1970, but below target during the remainder of the 1970's. The paths imply that the Fed rapidly decreased target inflation between 1980 and 1986 and then gradually lowered the

inflation target by about 1.5 percentage points over the next 15 years. The paths imply a constant target value for the real rate of interest.

An often-used alternative hypothesis is that Fed targets were constant and equal to sample average values (.042 for the inflation rate and .063 for the Treasury bill rate). This alternative is unattractive because it implies that the inflation rate and interest rate were below target in every quarter after 1991:IV with the implication that a superior policy would have raised the average inflation rate by 1.5 percentage points per year. The approach used in this paper to estimate target values for output, inflation, and the interest rate is reasonable but far from ideal. While it would be preferable to estimate target values, estimating them along with other model parameters is beyond the scope of the current project.

There is no evidence of a unit root in y , p , or r . For y , one would reject the hypothesis of a unit root at the five-percent level using the augmented Dickey-Fuller statistic with lag lengths 2 to 4. For p and r one would reject a unit root at the five percent level using the augmented Dickey-Fuller or the Phillips-Perron tests with lag lengths 0 through 4.

How many policy regimes occur between 1965 and 2001? A clear candidate for regime change is 1980:I, the first full quarter of the anti-inflation policy introduced by Paul Volcker. Table 1 reports linear projections of y , p , and r on the state variables of my model. The estimates show that reaction-function coefficients changed in 1980:I. The p-value associated with the relevant F statistic is .039. Given the changes in the target paths that occurred in 1986, it would be reasonable to test for a third policy regime but there are too few data points to estimate a separate regime for 1980-85.

There is not strong evidence that the coefficients of the y and p equations changed in 1980. The p-values associated with the relevant F statistics are .42 for the y equation and .49 for the p equation. The p-value associated with the likelihood ratio statistic for the hypothesis of no shift in the parameters of either equation is .38.⁶ Because the model is forward-looking, changes in the policy rule will cause across-regime changes in the parameters of the reduced form equations. But there is not evidence of large changes in reduced form parameters that could indicate a structural break.

⁶The finding of no shift in the coefficients of the p and r equations is robust to allowing for different error covariance matrices in each regime. The p-value for that likelihood ratio statistic is .35.

As equations (7) and (8) make clear, loss-minimizing values of the policy rule coefficients depend on the covariance matrix of shocks (\mathbf{S}) as well as on the reduced form parameters (\mathbf{G})⁷. It is therefore important to ask if the Fed faced a different distribution of shocks in the two regimes. The p-value associated with the likelihood ratio test of the hypothesis that the residual covariance matrix of the Table 1 projections is the same in the two regimes is less than .001. Based on the sample error covariance matrices, the Fed faced larger output and inflation shocks in the first regime than in the second. Inflation and output shocks were negatively correlated in the first regime but uncorrelated in the second. Interest rate shocks were uncorrelated with output and inflation shocks in the first regime, but positively correlated with both, and strongly correlated with output, in the second. Overall, the evidence suggests that the covariance matrix of shocks differed in the two regimes.

In summary, preliminary analysis of the data for y , p , and r suggests two policy regimes, the second beginning in 1980:I. The parameters of (1)–(3) will be held constant across the regimes but the coefficients of the policy rule and the covariance matrix of shocks will be permitted to change.

3. ECONOMETRIC POLICY EVALUATION OF U.S. MONETARY POLICY BETWEEN 1965 AND 2001

Tables 2 and 3 reports SURF and SORF estimates of the parameters of (1)–(3). As a measure of parameter-estimate significance, I report the p-value of the likelihood ratio for the hypothesis that the true value of the parameter is zero. A similar approach is used by Gallant, Hsu, and Tauchen (1999). Because the likelihood surface is not smooth, derivative-based procedures produced unreasonable estimates of standard errors. The estimation algorithm is described in Appendix A.

Estimates of Structural Parameters

SURF results support several findings. First, the estimate of β is sizable and significant implying that expected future output plays an important role in determining aggregate demand. Second, estimates of a_1 and a_2 imply that shocks have sustained effects on output. Third, the estimate of μ_1 is very small and insignificant which results in a backward-looking Phillips curve and casts doubt on the sticky price mechanism of Calvo. Fourth, the estimate of b is small but significant implying that changes in the interest rate do cause changes in aggregate demand. Fifth, the estimate of δ is large and significant which implies that a positive output gap raises the inflation rate. SURF estimates of the reaction function coefficients will be discussed shortly.

⁷While certainty equivalence does not hold for the optimal simple rule, Svensson and Woodford (2003) make clear that it does hold for the optimal time-consistent policy.

The SORF estimates differ in several important ways. The SORF estimate of β is smaller although still highly significant. The SORF estimate of b is larger implying that a change in the interest rate has a larger effect on aggregate demand. SORF estimates of both μ_1 and μ_2 are positive and highly significant indicating support for the forward-looking Phillips curve implied by the Calvo price-setting mechanism. One cannot reject the hypothesis that μ_1 and μ_2 sum to one as in Fuhrer and Moore (1995) and SORF estimates were obtained with that restriction imposed. While still highly significant, the SORF estimate of δ is much smaller implying that it is more difficult for the Fed to control inflation by changing aggregate demand.

As the p-values reported in Table 2 indicate, μ_1 is more precisely estimated under SORF than under SURF, so that the SORF estimates more strongly support a forward-looking Phillips curve. Likelihood-surface plots, reported in the extended version of the paper, show that a wide range of values for μ_1 fit the data equally well under SURF but that μ_1 is very precisely estimated under SORF. It appears that the optimal-policy restriction contains information that sharpens the estimate of μ_1 . Why? As will be discussed below, the interest rate is highly persistent in both regimes. The best way to reconcile a persistent interest rate process with the restriction that the process resulted from an optimal policy is that both occurred in an economy where agents are forward looking.

Estimates of Policy-Rule Coefficients and Tests of Policy Optimality

Table 3 reports estimates of policy rule coefficients obtained via OLS, SURF, SORF and a two-part procedure that will be explained shortly. Table 3 supports several conclusions. First, there is clear evidence of a change in policy regime in 1980 in the form of substantial and significant across-regime differences in policy-rule coefficients. The coefficient on lagged inflation (α_2) is small, negative, and insignificant in the first regime but large, positive, and highly significant in the second. Second, SORF estimates of the policy rule coefficient on lagged inflation are positive in both regimes and larger than either the SURF or OLS estimate in regime two. This suggests that the Fed should have responded more strongly to inflation than it did—an issue that we take up shortly. Third, under SORF the best fitting policy-rule coefficients are obtained for values of W_y and W_r that are very small implying that inflation stabilization was the dominant objective of policy in both regimes. The weight placed on output stabilization is .005 in regime one and .001 in regime two, neither significantly different from zero. The weight placed on interest rate stabilization is small but significant in each regime and slightly smaller in the second.

Was Federal Reserve policy optimal? A test of policy optimality is a test of the hypothesis that SORF fits the data as well as SURF, that is, that the SORF and SURF policy-rule coefficients are the

same. The answer is no. For regime one, the relevant likelihood ratio statistic is 11.80 with a p-value of 0.003. For regime two, the statistic and its p-value are 11.1 and 0.004. Thus, the hypothesis that the policy rule coefficients that fit the data best were those that minimized quadratic expected loss as defined by (6) is soundly rejected.

Interest Rate Smoothing

As Table 3 shows, the reaction-function coefficient for the lagged interest rate is always large relative to the coefficients for inflation and output so that, according to the definition provided by Clarida, Gali and Gertler (2000), the Fed engages in “interest rate smoothing.” Goodfriend (1991) suggests that the motive for interest rate smoothing is fear of disruption of financial markets. Sack and Weiland (2000) suggest three competing explanations for interest rate smoothing. First, the Fed may realize that rules with persistence are more effective for stabilizing output and inflation than rules without persistence in economies where forward-looking expectations matter. Second, the Fed may be uncertain about the accuracy of the data on which they condition policy. Third, the Fed may be uncertain about the values of structural parameters. Lansing (2002) cautions that smoothing can arise spuriously if the lagged rate of interest is correlated with measurement errors in the data.

While the reaction functions reported in Table 3 imply that the Fed smooths interest rates, SURF estimates suggest that the Fed placed little weight on interest rate stability *per se*. Thus, my estimates are consistent with Sack and Weiland’s first explanation for smoothing. The large coefficient on lagged interest rates required to make the model’s reduced form fit the data can be explained by the finding that forward-looking expectations are empirically important and the fact that optimal policies require persistent responses to shocks.

Test of the New-Keynesian Model

SURF estimation places one over-identifying restriction on the model’s reduced form. Log likelihood for the SURF estimates is 1460.5. Log likelihood for the linear projections reported in Table 1 is 1464.6. The p-value for the likelihood ratio test of the hypothesis that the SURF estimates explain the data as well as the unrestricted linear projections is .004. Sign restrictions imposed on parameters in the course of estimation are not responsible for the small p-value since very nearly the same SURF log likelihood is obtained when the sign restrictions are relaxed. The saddle path restriction is likewise not responsible because relaxing it raises SURF log likelihood by very little. Strictly speaking, the data reject the class of new Keynesian models consistent with (1)–(3).

Why does the New Keynesian model fit the data less well than OLS projections? There appear to be two reasons. First, the OLS estimates imply, for both regimes, that a positive shock to the interest rate initially raises the inflation rate. It takes about eight quarters before a negative effect of the shock on inflation appears. The New Keynesian model employed in this paper cannot account for this sort of inflation response. Second, in the first regime, the responses of output to an inflation shock and of inflation to an output shock are more sustained than the responses implied by the New Keynesian model. The following subsection provides a more detailed description of the impulse response functions implied by OLS, SURF and SORF estimates.

Comparison of Impulse Response Functions for OLS and Structural Estimation

While the data reject the model, the model may still explain much of the variation in the data. To check this possibility, I compare the impulse response functions (IRF's) implied by SURF, SORF, and unrestricted linear projections.⁸ Figures 2 and 3 present the IRF's for regimes one and two. The confidence intervals displayed in the figures were computed with a procedure explained in the extended version of the paper.

The figures make clear that the SURF and SORF estimates duplicate many features of the unrestricted OLS impulse response functions. For example, OLS, SURF and SORF estimates all imply that a positive shock to output keeps output above its trend for 12 quarters in the first regime and longer in the second. All three imply that a positive shock to inflation produces moderate inflation persistence with inflation returning to trend in about six quarters.⁹ All three imply that a positive inflation shock causes output to fall below trend for a substantial period of time. Both SURF and SORF match the OLS

⁸In keeping with my hypothesis that monetary policy is conditioned only on lagged values of output and inflation, I choose y-p-r as the within period causal ordering for the impulse response function. The ordering implies that a shock to output affects only output contemporaneously, a shock to inflation affects both inflation and output contemporaneously, and a shock to the interest rate affects all variables contemporaneously. The ordering makes very little difference in the first regime because the off-diagonal elements of the error shock covariance matrix are small. The ordering makes a difference in the second regime because the contemporaneous correlation between output and the interest rate is .42.

⁹The OLS impulse response functions imply less inflation persistence than reported by Fuhrer and Moore (1995b) who do not extract a segmented trend from inflation as I do. The high degree of inflation persistence reported by Fuhrer and Moore may result from not correcting inflation data for changes in target inflation rates. Alternatively, Dittmar, Gavin and Kydland (2002) account for inflation persistence by showing that a Taylor-type interest rate rule translates output persistence into inflation persistence. Since my reaction function is a Taylor-type rule, the Dittmar-Gavin-Kydland explanation is not consistent with my finding of less inflation persistence.

output response to an inflation shock better in regime two than in regime one. OLS, SURF, and SORF IRF's all show that the interest rate rises quickly after a positive output shock and remains above trend for at least six quarters.

There are differences in the impulse response functions. Both OLS and SURF estimates imply that the interest rate falls after a positive inflation shock in regime one and rises after a positive inflation shock in regime two. SORF estimates imply that the interest rate rises after a positive inflation shock in both regimes and that the interest rate response is much larger than predicted by the OLS estimates in the second regime. A second difference is related to the first. SORF estimates imply that the inflation response to a positive output shock is far smaller than implied by either the OLS or SURF estimates.

Comparing the IRF's suggest that actual and optimal responses of the interest rate to output shocks were nearly the same and, given the SORF estimates of the structural parameters, would have completely damped inflation resulting from those shocks. The comparison also implies that actual interest rate responses to inflation shocks were not optimal. A loss-minimizing policy would have raised rather than lowered the interest rate after inflation shocks in regime one and raised it more in regime two.

Figure 4 compares the responses of output and inflation to an interest rate shock. The r-equation residuals are positively correlated with the y-and-p equation residuals in the second regime. On the maintained hypothesis that the policy authority reacts to economic conditions only with a lag, this correlation between must be due either to responses of output and inflation to an interest rate surprise or to responses of y, p, and r to an un-modeled shock. To obtain the model's prediction of responses of output and inflation to a pure interest rate shock, I ignore the contemporaneous correlations when computing the IRF's reported in Figure 4.

Figure 4 shows that in both regimes SURF and SORF estimates can account for the "hump-shaped" response of output to an interest rate surprise. For both SURF and SORF and in both regimes, peaks occur about 8 quarters after the interest rate shock. OLS estimates imply that the inflation rate first rises and then falls after an interest rate shock while SURF and SORF estimates show that the inflation rate declines immediately after an interest rate shock and remains below trend for many quarters. SORF estimates imply a milder inflation response than do SURF estimates.

SURF and SORF estimates are well behaved in another sense. For both regimes, Q statistics indicate no sign of serial correlation y and p equation residuals. For the r equation, Q statistics show no sign of serial correlation in the first regime but suggest some serial correlation in the second.

Because SURF and SORF estimates of the new Keynesian model explain much of the variation in the data, I conclude that it is useful to conduct policy evaluation using those estimates.

Policy Evaluation with SURF Estimates

I use two strategies to draw out policy conclusions from the SURF estimates. In the first, I fix the parameters of S at their SURF estimates and compute loss-minimizing policy-rule coefficients for various loss function weights. Figure 5 reports the ratio of policy loss (**7**) computed for the SURF-estimated rule to policy loss computed for the counterfactual optimal rule¹⁰. For regime one, the loss ratio is smallest (1.1) when very high weight is placed on stabilizing inflation and very low weight is placed on stabilizing output and the interest rate. When substantial weight is placed on stabilizing output, the observed rule does twice as badly as the counterfactual optimal rule. When substantial weight is placed on stabilizing the interest rate, the observed rule performs even more poorly. For regime two, the story is different in one important way. The loss ratio is extremely high when substantial weight is assigned to stabilizing output. It is still the case that the loss ratio is smallest when the weight placed on inflation stabilization is very large although the minimum loss ratio (1.7) is larger than for regime one.

Based on the relative performance of observed and optimal rules, I conclude that the Federal Reserve assigned higher weight to stabilizing inflation than to other objectives and that it assigned little if any weight to stabilizing output in the second regime. Otherwise, one must conclude that the rule that fits the data best is grossly inferior to others that were available.

The second strategy is a version of the two-part policy program—SORF estimation with structural parameters constrained to SURF estimates. Because changes in the coefficients of the policy rule imply changes in the residuals of all three equations, the two-part analysis treats policy weights and the residual covariance matrices as “free parameters” and finds values that produce the highest log likelihood.

The results of the two-part analysis are reported in Table 3 and imply that the Fed placed slightly more weight on stabilizing output than on stabilizing inflation in the first regime, dramatically lowered that weight in the second regime, and substantially lowered the weight placed on interest rate stabilization in the second regime. The weights are not precisely estimated. A plot of the log likelihood surface shows that the data contain little information about W_y . The reduced form that results from the two-part analysis fits the data far less well than SURF or SORF which is further evidence in favor of a unified approach to policy evaluation.

Policy Evaluation with SORF Estimates

¹⁰In Figure 5, the loss-function weights are normalized by requiring that they sum to one. Thus, along the diagonal of the floor of the figure, the weight on interest-rate stabilization is zero, while in the back-left corner it is one.

What other policy conclusions can be drawn from the unified approach and its resulting SORF estimates? As already mentioned, the SORF estimates of W_y and W_r indicate that the Fed placed much greater weight on stabilizing inflation than on stabilizing output or interest rates in both regimes and placed more weight on inflation stabilization in the second regime. A plot (not included) of log likelihood shows that the loss-function weights are sharply estimated in the sense that log likelihood falls a great deal when W_y and W_r are raised above their point estimates. In the context of equations (1)–(3), it is not possible to reconcile optimal policy with the hypothesis that the Fed regarded output stabilization as important as inflation stabilization.

As reported above, the unrestricted policy rule fits the interest rate better than the rule restricted to be loss minimizing. Using SORF estimates of the loss function weights, I find that OLS policy rule coefficients imply that expected loss was 3.1 percent higher than the loss of an optimal policy in regime one and .5 percent higher in regime two. Contrasting this finding with the comparable finding using SURF parameter estimates, provides additional evidence for the unified approach to policy evaluation. Using SURF estimates of S leads to the conclusion that a change in the policy rule could have lowered loss by 70 percent in the second regime. Using SORF estimates implies that while policy was not optimal in either regime, improving it would not have greatly reduced loss.

Robustness Checks: Alternative Hypotheses about Policy Objectives

Here I check the robustness of my findings by considering alternative hypotheses about Federal Reserve policy objectives and the form of its loss function. One possible mis-specification involves the way that interest rate stability is defined. For Svensson (1999) and Favero and Ravelli (2003), interest rate stability means low variation in interest rate changes rather than small departures of interest rate levels from target. Table 4 reports parameter estimates when the interest rate first difference (Δr) is the dependent variable of the policy rule and an argument of the Federal Reserve loss function. Both SURF and SORF estimates are very close to their Table 2 counterparts. Log likelihood values are slightly higher but SURF log likelihood remains significantly larger than SORF log likelihood. For SORF, Q statistics indicate that the residuals are not serially correlated in either regime. For SURF, it appears that the interest rate equation residuals are serially correlated in the second regime but not the first.

Table 5 reports policy rule coefficients for the Δr case. There are five policy rule coefficients because the state vector comprises y_{t-1} , p_{t-1} , r_{t-1} , y_{t-2} and r_{t-1} . The corresponding loss function has four arguments and three weights to be estimated. As Table 5 shows, both SURF and SORF continue to detect the substantial increase in the coefficient on lagged inflation that occurred in regime two. SURF, SORF and OLS all produce similar estimates of the other parameters of the policy rule. For regime one,

SORF estimates imply that stabilizing π was the most important policy objective and that stabilizing inflation was more important than stabilizing output. For regime two, the estimates imply that inflation stabilization was by far the most important policy objective. SURF and SORF log likelihood are again significantly smaller than OLS log likelihood. In sum, the important findings associated with a policy rule for r are robust to re-defining policy in terms of π .

I next consider the possibility that the Fed defines output stability in terms of the growth rate of output rather than the level. If lags in the IS schedule are due to habit persistence in the representative agent utility function and if the Federal Reserve loss function is an approximation to the representative agent utility function, the Fed will want to stabilize changes rather than levels of output. To check this alternative specification requires only re-definition W , the (4x4) matrix of loss function weights from equation (6).¹¹ SORF estimation of this alternative specification produces parameter estimates that are almost identical to those reported in Table 2. To conserve space, I do not report the details. Log likelihood is slightly lower and estimates of the weights placed by the Fed on interest rate and output growth stabilization in regime two are slightly higher.

As a final robustness check, I investigate whether differences between SURF and SORF policy-rule coefficients are due to the assumption that W is diagonal. I fix the parameters in S at their SURF estimates and search over the nine¹² free elements of W for values that minimize the sum of squared differences between SURF and SORF policy-rule coefficients.

The investigation leads to three conclusions. First, it is not possible to reconcile SURF and SORF policy-rule coefficients for either regime. Significant difference remain between the SURF estimates and the SORF estimates associated with the best W . Second, it is difficult to tell a coherent story about first regime policy objectives. The W that minimizes squared differences between SURF and SORF rule coefficients implies that $\mathbf{X}'_t \mathbf{W} \mathbf{X}_t \approx (\mathbf{y}_t - \mathbf{y}_{t-1})^2 - .15 (\mathbf{y}_t - \mathbf{y}_{t-1}) \mathbf{p}_t$. This loss function implies that the Fed stabilizes output growth but that inflation **lowers** loss whenever growth is above target. Third, the results for the second policy regime largely confirm earlier findings. The W that minimizes squared differences between SURF and SORF rule coefficients for the second regime implies that $\mathbf{X}'_t \mathbf{W} \mathbf{X}_t \approx (\mathbf{p}_t)^2 + .21 (\mathbf{y}_t - \mathbf{y}_{t-1}) \mathbf{p}_t$. This loss function implies that the Fed cares primarily

¹¹The hypothesis that the Fed cares about the growth rate rather than the level of output implies that W_y , the weight on output growth stabilization, appears in the (1,1) and (4,4) position of W and that $-W_y$ appears in the (4,1) and (1,4) positions.

¹²There are ten free elements in a (4x4) symmetric matrix but I continue to impose $W_p = 1$ as a normalization.

about stabilizing inflation and places greater weight on inflation stability when output growth and inflation are both above target values.

4. CONCLUSIONS

In the unified approach to econometric policy evaluation, parameters of a forward-looking model and a central-bank loss function are estimated simultaneously subject to the restriction that the central-bank policy rule minimizes expected loss. This paper demonstrates that the unified approach is tractable and provides evidence that estimation subject to the optimal policy restriction both alters and sharpens estimates of the model's parameters. This finding casts doubt on policy conclusions reached via the widely used two-part approach and should not be surprising since optimal policy is very different when agents are forward looking than when they are not.

The paper provides evidence that stabilization of the inflation rate was a far more important policy objective of the Federal Reserve than stabilization of output both before and during the Volcker-Greenspan era. This finding contradicts Dennis (2002) and supports Favero and Ravelli (2003) both of whom use a unified approach with a backward looking structural model. Interestingly, a version of the two-part approach to policy evaluation implies that output and inflation stabilization were equally important prior to 1980. However, the reduced form implied by the two-part approach fits the data very poorly. The finding that output stabilization was a relatively unimportant policy objective is robust to specifying the interest rate rule in either level or first-difference form and the output stabilization objective in either level or growth form.

There is evidence that the Fed did not follow an optimal policy. SORF estimates that impose the optimal-policy restriction do not fit the data as well as the SURF estimates that do not. Comparison of SURF and SORF impulse response functions for regime one suggests that the Fed accommodated inflation shocks when it should have leaned against them. A similar comparison for regime two suggests that the Fed responded to second-regime inflation shocks by moving interest rates in the right direction but less strongly than it should have. The accommodation of inflation shocks in the first regime is not evidence that the Fed attached great weight to output because SORF estimates of W_y are very small. A comparison of the policy rule that fits the data best with the rule that minimizes loss suggests that improving policy could have lowered loss by 3.1 percent in regime one and .5 percent in regime two.

There are several promising directions for future research with the unified approach. It would be interesting to apply the approach with a broader class of central bank objective functions. Central bank preferences may be asymmetric and algorithms introduced in this paper can be easily altered to accommodate asymmetric loss functions. Central banks may respond to far-from-target realizations of

inflation or output by raising the weight placed on stabilizing the variable that has strayed. The unified approach is well suited to testing hypotheses about such “endogenous” regime changes.

While the unified approach is tractable when the minimization algorithm of the central bank is nested within the estimation algorithm, computation could be speeded with a one-pass algorithm. I am currently experimenting with an algorithm that combines the first-order conditions from the policy problem with moment conditions that characterize best-fit parameter estimates.

Finally, it will be interesting to take the unified approach to data for different countries and to different and larger economic models. Givens (2004) considers a representative agent model in which wages and prices are both sticky and finds that delegating policy authority to a central banker who cares about wage stabilization produces policy outcomes almost as good as provided by the timeless perspective policy. Atoian (2004) studies the performance of optimal simple rules using data for Canada and an open economy model and finds that the Bank of Canada has assigned greater weight to stabilizing inflation in recent years.

APPENDIX A: Description of the Inverse Control Algorithm (This appendix is an expanded version of its counterpart in the published paper.)

The parameters of the forward looking model are divided into three groups: S , W , and $\mathbf{2}$. Vector S contains the parameters of the IS schedule, Phillips curve, and Choleski decomposition of the reduced form error-covariance matrix, vector W contains the parameters of the central bank loss function, and vector $\mathbf{2}$ contains the coefficients of the central bank reaction function. W_r , W_y , and W_p , measure the importance to the central bank of interest rate, output, and inflation stability. As a normalization, W_p is fixed at 1.0. Because preliminary work suggested that it would be difficult to identify the central bank time rate of discount, β is fixed at 0.99.

The SURF algorithm estimates S and $\mathbf{2}$ without restricting $\mathbf{2}$. For given S and $\mathbf{2}$ the algorithm computes the Blanchard and Kahn solution for the model, uses it to compute predicted values for y_t , p_t and r_t , and computes residuals and log likelihood using standard formulas. Because some values of S and $\mathbf{2}$ yield B matrices that do not satisfy the Blanchard and Kahn roots condition, the maximand used in SURF estimation is log likelihood minus a penalty.

After experimenting with alternatives, I chose a penalty-function approach to handle values of S and $\mathbf{2}$ that violate the Blanchard and Kahn root condition. Let J_1 through J_6 be the eigenvalues of matrix B arrayed in ascending order. Suppose a unique saddle path solution exists. Then $|J_k| < 1.0$ for $k = 1, \dots, 4$ and $|J_k| > 1.0$ for $k = 5$ and 6. Let I_k be the indicator function that equals zero if the root condition is satisfied for the k^{th} root and 1 if it is not. The penalty function is

$$\text{Penalty} = \bar{p} \sum_{k=1}^6 I_k (|J_k| - 1.0)^2$$

where \bar{p} is a constant chosen to guarantee that the Penalty is large relative to log likelihood. If the root condition is satisfied, Penalty = 0. If it is not satisfied, Penalty is a smooth, positive function of the difference between out-of-bounds roots and 1.0. When the root condition is violated, equation (6) can evaluate to complex numbers. To handle this problem, I set each elements of sub-matrix $G_{1,1}$ equal to the real part of the corresponding element of $(K_{2,2})^{-1} K_{2,1}$. When the root condition is met, the conversion has no effect. Test runs showed that the hill climber quickly learned to keep parameters in bounds.

The hill climber used was PATTERN from Version 6 of the GQOPT Library of Fortran optimization programs (Goldfeld and Quandt, 1972). PATTERN is a direct search algorithm that combines exploratory searches parallel to the parameter-space axes and “pattern” searches in directions found

successful in recent iterations. For smooth likelihood functions, derivative-based programs such as DFP often outperform derivative-free programs. However, preliminary tests showed that for my problem PATTERN worked much better than DFP. For given test samples, PATTERN always achieved higher log likelihood than DFP. Across those samples, PATTERN estimates displayed less bias than DFP estimates.

Because GQOPT folklore also holds that the performance of PATTERN can be very sensitive to the initial step size of the search, I re-started PATTERN many times with different initial step sizes. Let ISS_j be the initial step size for run j of PATTERN. Based on preliminary tests, I chose 0.10 for ISS_1 and set $ISS_j = .5 ISS_{j-1}$ for $j = 2, \dots, 8$. I then repeated the loop of eight re-starts twice for a total of 24 calls to PATTERN per estimation. For each call, I set initial values of S and W parameters equal to the final values from the previous call. SURF estimates typically converged for j between 2 and 4. SORF estimates frequently required more than 8 re-starts to converge.

SORF computes maximum likelihood estimates of S and W subject to the restriction that the coefficients of the central bank reaction function minimize expected loss, that is that $\mathbf{2} = g(S, W)$. SORF nests the central bank minimization problem inside the econometrician's estimation problem. For SORF estimation, I used PATTERN and the same formula to compute Penalty. To impose $\mathbf{2} = g(S, W)$, I used the IMSL subroutine DUMINF to find the values for $\mathbf{2}$ that minimized expected loss as defined by (7) - (9). DUMINF was called twice with different starting values. The first call initially set each element of $\mathbf{2}$ to a small positive number. The second initially set each element of $\mathbf{2}$ to the previous best estimate. Test runs showed that most often the two calls to DUMINF produced the same final value for $\mathbf{2}$. When the two calls produced different values, the one associated with lower expected loss was chosen.

I experimented with two strategies for setting SORF starting values for S and W . The first strategy used the same starting values for S that were used for SURF and naive values for W . The second strategy started SORF at SURF estimates of S . In test runs, I discovered that the first strategy outperformed the second.

A Monte Carlo procedure was also used to construct the confidence intervals reported in Figures 2-4.

Table A
Results of Monte Carlo Study Based on 101 Samples
Each with 58 Observations in Regime One and 88 Observations in Regime Two

Parameter	True Value	Estimation with Unrestricted Reaction Function Coefficients		Estimation with Optimal Reaction Function Coefficients	
		Mean	Std. Dev.	Mean	Std. Dev.
$\ln(\boldsymbol{\theta})$	-1.82	-3.23	3.80	-1.35	0.16
a_1	1.10	1.03	0.23	0.97	0.08
a_2	-0.30	-0.27	0.10	-0.26	0.06
$\ln(b)$	-3.18	-3.76	1.97	-3.45	0.38
$\ln(\boldsymbol{\pi}_1)$	-0.57	-2.79	4.30	-0.61	0.05
$\ln(\boldsymbol{\$})$	-8.18	-11.05	5.90	-6.80	0.77
Regime 1 Reaction Fctn. Coefs.					
y_{t-1}	0.37	0.38	0.11	0.39	0.10
p_{t-1}	0.14	0.19	0.07	0.16	0.05
r_{t-1}	0.91	0.88	0.05	0.90	0.03
y_{t-2}	-0.38	-0.38	0.12	-0.39	0.10
$\ln(W_y)$	-5.37	—	—	-3.39	1.60
$\ln(W_r)$	-4.34	—	—	-2.54	1.40
Regime 2 Reaction Fctn. Coefs.					
y_{t-1}	0.21	0.19	0.15	0.22	0.15
p_{t-1}	0.41	0.44	0.10	0.43	0.08
r_{t-1}	0.85	0.83	0.05	0.82	0.04
y_{t-2}	-0.26	-0.24	0.15	-0.27	0.15
$\ln(W_y)$	-6.77	—	—	-4.59	1.50
$\ln(W_r)$	-5.09	—	—	-3.60	1.20
Log Likelihood*	1449.0*0	1456.20	16.50	1455.20	14.70

* True log likelihood is computed with the standard formula and the true covariance matrices.

APPENDIX B: Detailed Description of the Solution (The equation numbers are those in the body of the text.)

To find the model's reduced form, I follow the approach of Blanchard and Kahn (1980). As in the text, define X_t as $X_t = (y_t \ p_t \ r_t \ y_{t-1})'$ and write the model as:

$$(4) \quad \begin{bmatrix} X_t \\ E_t y_{t+1} \\ E_t p_{t+1} \end{bmatrix} = B \begin{bmatrix} X_{t-1} \\ y_t \\ p_t \end{bmatrix} + D \begin{bmatrix} u_t \\ v_t \\ w_t \end{bmatrix}$$

where

$$B = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ \theta_1 & \theta_2 & \theta_3 & \theta_4 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ \frac{b\theta_1 - a_1}{\lambda} & \frac{b(\theta_2\alpha_1 + \alpha_2)}{\alpha_1\lambda} & \frac{b\theta_3}{\lambda} & \frac{b\theta_4 - a_2}{\lambda} & \frac{\alpha_1 + b\beta}{\lambda\alpha_1} & \frac{-b}{\alpha_1\lambda} \\ 0 & -\frac{\alpha_2}{\alpha_1} & 0 & 0 & -\frac{\beta}{\alpha_1} & \frac{1}{\alpha_1} \end{bmatrix}$$

and

$$D = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ -\frac{1}{\lambda} & \frac{b}{\alpha_1\lambda} & \frac{b}{\lambda} \\ 0 & -\frac{1}{\alpha_1} & 0 \end{bmatrix}$$

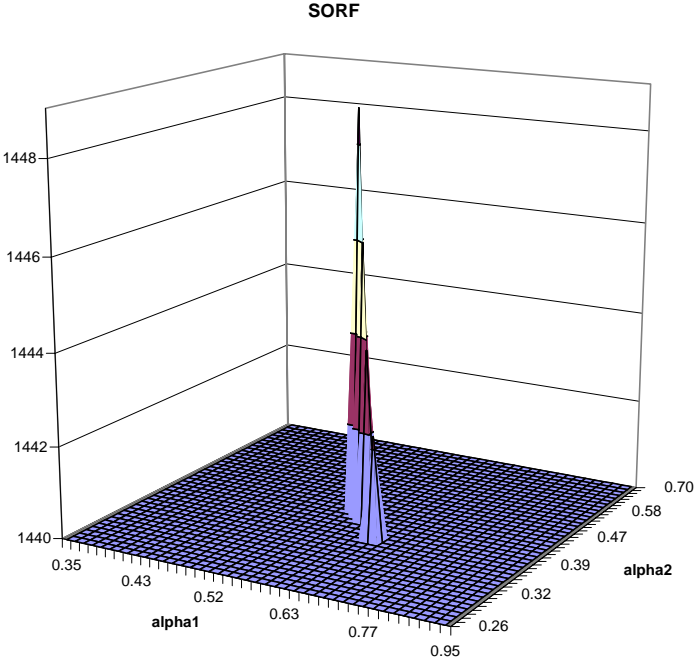
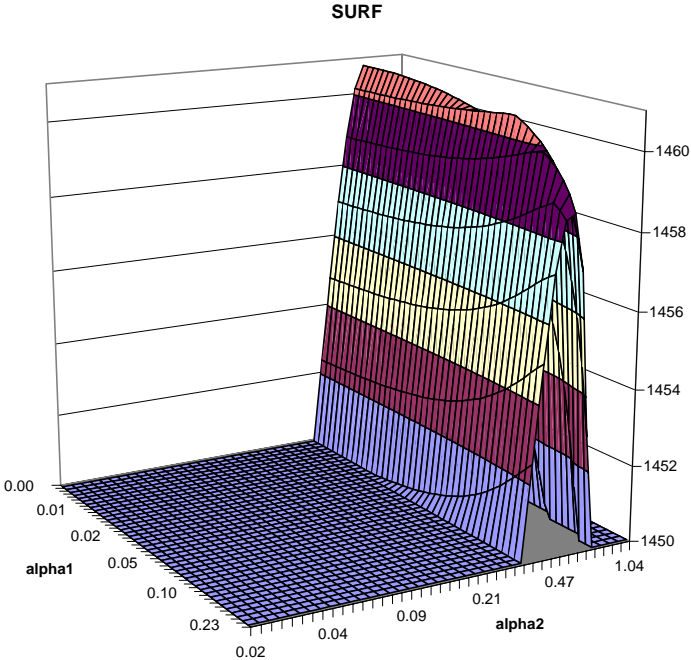
In the language of Blanchard and Kahn, X_{t-1} is the vector of predetermined variables, and y_t and p_t are "not predetermined". Blanchard and Kahn uncouple B so that $B = K^{-1} J K$ where J is a diagonal matrix containing the eigenvalues of B in ascending order and where the columns of K^{-1} are eigenvectors of B .

If the number of unstable eigenvalues of B equals the number of not-predetermined variables and if $K_{2,2}$ is non-singular, then there exists a unique saddle path for the model given by

$$(5) \quad X_t = GX_{t-1} + \phi_t = \begin{bmatrix} -K_{2,2}^{-1} K_{2,1} \\ \theta_1 & \theta_2 & \theta_3 & \theta_4 \\ 1 & 0 & 0 & 0 \end{bmatrix} X_{t-1} + \begin{bmatrix} \phi_{1,t} \\ \phi_{2,t} \\ \phi_{3,t} \\ 0 \end{bmatrix}$$

where G is (4x4), $K_{2,2}$ and $K_{2,1}$ are the (2x2) lower right block and the (2x4) lower left block of K (Blanchard and Kahn, Proposition 1). For this paper's model, a unique saddle path exists if B has two eigenvalues with absolute values greater than one. No solution exists if there are too many unstable eigenvalues. A multiplicity of solutions can exist if there are too few unstable eigenvalues.


APPENDIX C: Likelihood Surface Plots for Parameters α_1 and α_2 , the Coefficients of Expected Future Inflation and Past Inflation in the Phillips Curve




The figures clearly indicate that α_1 is far more precisely estimated with the SORF procedure.

APPENDIX D: Plots of Log Likelihood That Show that Loss Function Weights Are Precisely Estimated

Regime 1



Regime 2



LITERATURE CITED

- Anderson, Evan, Lars P. Hansen, Ellen R. McGrattan, and Thomas J. Sargent, "Mechanics of forming and estimating dynamic linear economies," in, Amman, Hans M., David A. Kendrick, and John Rust, eds., *Handbook of Computational Economics* 1, Handbooks in Economics 13, Elsevier Science, North-Holland, Amsterdam, 1996, 171-252.
- Atoian, Rouben, "Econometric Investigation of Policy Preference Evolution: The Bank of Canada's Case," University of North Carolina mimeo, 2004
- Atoian, Rouben V., Gregory E. Givens, and Michael K. Salemi, "Policy evaluation with a forward-looking model," in Minford, Patrick, ed., *Money Matters: Essays in Honour of Alan Walters*, 2003, Edward Elgar, Cheltenham, U.K., 294-316.
- Backus, David and John Driffill, "The consistency of optimal policy in stochastic rational expectations models," CEPR Discussion Paper No. 124, August, 1986.
- Bernanke, Ben, Mark Gertler, and Simon Gilchrist, "The financial accelerator in a quantitative business cycle framework," NBER Working Paper 6455, March, 1998.
- Blanchard, Olivier J. and Charles M. Kahn "The solution of linear difference models under rational expectations," *Econometrica* 48, 1980, 1305-11.
- Calvo, Guillermo A., "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics* 12, 1983, 383-98.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," Revision of NBER Working Paper 8403, 2003.
- Clarida, Richard, Jordi Gali, and Mark Gertler, "The science of monetary policy: a new Keynesian perspective," *Journal of Economic Literature* XXXVII, 1999, 1661-1707.
- _____, "Monetary policy rules and macroeconomic stability: Evidence and some theory," *The Quarterly Journal of Economics*, 115, 2000, 147-80.
- Clark, Peter B., Charles A. E. Goodhart and Haizhou Huang, "Optimal monetary policy rules in a rational expectations model of the Phillips Curve," mimeo, September, 1996.
- Dennis, Richard, "The policy preferences of the US Federal Reserve," Federal Reserve Bank of San Francisco mimeo, April, 2002.
- Dittmar, Robert, William T. Gavin, and Finn E. Kydland, "Inflation persistence and flexible prices," Federal Reserve Bank of St. Louis Working Paper 2001-010c, 2002.
- Favero, Carlo A. and Riccardo Rovelli, "Macroeconomic stability and the preferences of the Fed: A formal analysis, 1961-98," *Journal of Money, Credit, and Banking*, 35, 2003, 545-56.
- Fuhrer, Jeffrey C., "Inflation/output variance trade-offs and optimal monetary policy," *Journal of Money Credit and Banking* 29, 1997a, 214-234.

- _____, "Comment on Rotemberg and Woodford," *Op. Cit.*, 1997b, 346-355.
- Fuhrer, Jeffrey C. and George R. Moore, "Monetary policy trade-offs and the correlation between nominal interest rates and real output," *American Economic Review* 85, 1995a, 219-239.
- _____, "Inflation persistence," *Quarterly Journal of Economics*, February, 1995b, 129-59.
- Gali, Jordi and Mark Gertler, "Inflation dynamics: a structural econometric analysis," *Journal of Monetary Economics* 44, 1999, 195-222.
- Gallant, A. Ronald, Chien-Te Hsu, and George Tauchen, "Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance," *The Review of Economics and Statistics*, 81, 1999, 617-631.
- Ghysels, Eric, Norman R. Swanson and Myles Callan, "Monetary Policy Rules with Model and Data Uncertainty," *Southern Economic Journal*, 69, 2002, 239-65.
- Givens, Gregory E., "Revisiting the Delegation Problem in a Sticky Price and Wage Model," University of North Carolina mimeo, 2004.
- Goodfriend, Marvin, "Interest rates and the conduct of monetary policy," *Carnegie-Rochester Conference Series on Public Policy*, XXXIV, 1991, 7-30.
- Goodfriend, Marvin and Robert G. King, "The new neoclassical synthesis and the role of monetary policy," in Bernanke, Ben S. and Julio J. Rotemberg, eds., *NBER Macroeconomics Annual 1997*, MIT Press, Cambridge, MA, 1997, 231-83.
- Hansen, Lars P. and Thomas J. Sargent, "Formulating and estimating dynamic linear rational expectations models," *Journal of Economic Dynamics and Control* 2, 1980, 7-46.
- Klein, Paul, "Using the generalized Schur form to solve a multivariate linear rational expectations model," *Journal of Economic Dynamics and Control* 24, 2000, 1405-23.
- Korenko, George, "Inferring Federal Reserve preferences using inverse control," Ph.D. Dissertation, Department of Economics, University of North Carolina at Chapel Hill, 1998.
- Lansing, "Real-time estimation of trend output and the illusion of interest rate smoothing," *Economic Review*, Federal Reserve Bank of San Francisco, 2002, 17-35.
- Marschak, Jacob, "Economic Measurements for Policy and Prediction," Chapter 1 in Hood, William C. and Tjalling C. Koopmans, *Studies in Econometric Method*, Cowles Commission for Research in Economics, Yale University, New Haven, Connecticut, 1953, 1-26.
- McCallum, Bennett T., "Comment on Rotemberg and Woodford," *Op. Cit.*, 355-361.
- McGrattan, Ellen R., "Solving the stochastic growth model by linear quadratic approximation," *Journal of Business and Economic Statistics* 8, 1990, 41-44.

- Orphanides, Athanasios, "Activist stabilization policy and inflation: the Taylor rule in the 1970s," mimeo, February, 2000.
- Orphanides, Athanasios, and J. C. Williams, "Robust monetary policy rules: The case of unknown natural rates of interest and unemployment," Federal Reserve Board mimeo, March, 2002.
- Papell, David H., "Monetary policy in the United States under flexible exchange rates," *American Economic Review* 79, 1989, 1106-16.
- Rotemberg, Julio J. and Michael Woodford, "An optimization-based econometric framework for the evaluation of monetary policy," in Bernanke, Ben S. and Julio J. Rotemberg, eds., NBER Macroeconomics Annual 1997, MIT Press, Cambridge, MA, 1997, 297-346.
- Sack, Brian and Volker Weiland, "Interest-rate smoothing and optimal monetary policy: A review of recent empirical evidence," *Journal of Economics and Business*, 52, 2000, 205-28.
- Salemi, Michael K., "Revealed preference of the Federal Reserve: using inverse-control theory to interpret the policy equation of a vector autoregression," *Journal of Business and Economic Statistics* 13, 1995, 419-433.
- Sargent, Thomas J., *Dynamic Macroeconomic Theory*, Harvard University Press, Cambridge, MA, 1987.
- Söderlind, Paul, "Solution and estimation of RE macromodels with optimal policy," *European Economic Review* 43, 1999, 813-23.
- Svensson, Lars E. O., "Inflation targeting: Some extensions," *Scandinavian Journal of Economics*, 101, 1999, 337-361.
- _____, "Open-economy inflation targeting," *Journal of International Economics*, 50, 2000, 155-83.
- _____, "What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules," *Journal of Economic Literature*, XLI, 2003, 426-77.
- Svensson, Lars E. and Michael Woodford, "Optimal policy with partial information in a forward-looking model: Certainty-equivalence redux," NBER Working Paper 9430, January 2003.
- Taylor, John B., "Estimation and Control of a Macroeconomic Model with Rational Expectations," *Econometrica* 47, 1979, 1267-86.
- _____, "Applying academic research on monetary policy rules: an exercise in translational economics," The Harry G. Johnston Lecture, Macro, Money, and Finance Research Group Conference, Durham University, Durham England, revised February, 1998.
- Woodford, Michael, "Control of the public debt: A requirement for price stability?," NBER Working Paper 5684, 1996.
- _____, "Interest and Prices." Princeton University Working Paper, 2002

_____, "Optimal Monetary Policy Inertia," NBER Working Paper No 7261, 1999.

Table 1
Coefficient Estimates for Linear Projections of Model Variables on the State Vector
1965. II – 2001. IV

Independent Variables	Dependent Variable					
	Output Gap		Inflation Rate		Interest Rate	
y_{t-1}	1.065 (0.12)	1.200 (0.08)	-.183 (0.15)	-.098 (0.10)	.174 (0.14)	.198 (0.10)
p_{t-1}	-.136 (0.07)	-.147 (0.06)	.523 (0.09)	.557 (0.07)	-.086 (0.09)	.043 (0.07)
r_{t-1}	-.091 (0.05)	-.083 (0.06)	.032 (0.06)	.031 (0.04)	.952 (0.06)	.892 (0.04)
y_{t-2}	-.109 (0.12)	-.229 (0.08)	.324 (0.15)	.181 (0.11)	-.107 (0.15)	-.196 (0.10)
$D(2) y_{t-1}$.264 (0.16)	--	.182 (0.21)	--	.071 (0.19)	--
$D(2) p_{t-1}$	-.055 (0.11)	--	.064 (0.15)	--	.331 (0.14)	--
$D(2) r_{t-1}$.011 (0.07)	--	-.028 (0.08)	--	-.137 (0.08)	--
$D(2) y_{t-2}$	-.239 (0.17)	--	-.269 (0.21)	--	-.189 (0.20)	--
R^2	.948	.946	.477	.464	.802	.787
P-value of F Statistic	0.42		0.49		0.039	

The table reports the coefficients of unrestricted linear projections of the output gap, inflation rate, and interest rate on the four state variables of the New Keynesian model (two lags of the gap, one lag each of inflation and the interest rate). $D(2)$ is a dummy variable that is 0.0 for regime one (1965:I–1979:IV) and 1.0 for regime two (1980:I–2001:IV). The data are described in the text and displayed in Figure 1. The p-value is for the F test of the hypothesis that the coefficients of the four state variables have the same values in regimes one and two.

Table 2 Maximum Likelihood Estimates of Model Parameters When the Interest Rate Level Is the Dependent Variable of the Policy Rule				
	Estimation with Unrestricted Reaction Function SURF		Estimation with Optimal Reaction Function SORF	
Model Parameters	Estimate	Significance (p-value)	Estimate	Significance (p-value)
8	.433	.029	.162	.003
a₁	.722	.050 ⁴	1.10	.578 ⁴
a₂	-.156	<.001	-.300	<.001
b	.0076	<.001	.042	<.001
"₁	.0001	1.00	.567 ⁵	<.001
"₂	.624	<.001	.433 ⁵	<.001
\$.055	.043	.0003	<.001
Log Likelihood	1460.53		1449.14	

1. The table reports maximum likelihood estimates of the parameters of the following model.

$$y_t = \lambda E_t y_{t+1} + a_1 y_{t-1} + a_2 y_{t-2} - b(r_t - E_t p_{t+1}) + u_t$$

$$p_t = \beta y_t + \alpha_1 E_t p_{t+1} + \alpha_2 p_{t-1} + v_t$$

$$r_t = \theta_1 y_{t-1} + \theta_2 p_{t-1} + \theta_3 r_{t-1} + \theta_4 y_{t-2} + w_t .$$

SURF estimation places no restrictions on the parameters of the reaction function. **SORF** estimation requires that the parameters of the reaction function minimize the loss function of the Federal Reserve.

2. The model parameter values reported in this table are assumed to be the same for both policy regimes. The parameters of the interest rate equation (**2₁ - 2₄**) are permitted to change across regimes. Parameter estimates for the reaction function are reported in Table 3.

3. Parameter-estimate significance is measured with the p-value of the likelihood ratio test statistic for the hypothesis that the true value of the parameter is very small (0.0001).

4. The p-value reported in this case is for a test of the hypothesis that $a_1 = 1.0$.

5. Estimation was undertaken subject to the restriction $"_2 = 1.0 - "_1$.

		Table 3 Policy Rule Coefficient Estimates When the Interest Rate Level Is the Dependent Variable of the Policy Rule						
Reaction Function Coefficients		OLS		SURF		Two-Part	SORF	
		Coefficient	P-value	Coefficient	P-value	Coefficient	Coefficient	P-value
Regime One	2₁	.174 (0.14)	.22	.277	.018	.036	.321	---
	2₂	-.086 (0.09)	.32	-.023	.716	-.0012	.135	---
	2₃	.952 (0.06)	<.001	.929	<.001	.798	.919	---
	2₄	-.107 (0.15)	.47	-.227	.069	.131	-.330	---
Loss Function Wt. on y		---	---	---	---	1.050	0.0047	.29
Loss Function Wt. on r		---	---	---	---	0.658	0.013	<.001
Regime 2	2₁	.245 (0.14)	.07	.186	.152	-.264	.197	---
	2₂	.245 (0.11)	.02	.356	<.001	.263	.406	---
	2₃	.815 (0.05)	<.001	.832	<.001	.761	.849	---
	2₄	-.295 (0.14)	.04	-.240	.081	.390	-.256	---
Loss Function Wt. on y		---	---	---	---	0.0003	0.0012	.244
Loss Function Wt. on r		---	---	---	---	0.0110	0.0062	<.001
Log Likelihood		1464.59		1460.53		1380.34	1449.14	

1. The table reports maximum likelihood estimates of the parameters of the following model.

$$\begin{aligned}
 y_t &= \lambda E_t y_{t+1} + a_1 y_{t-1} + a_2 y_{t-2} - b(r_t - E_t p_{t+1}) + u_t \\
 p_t &= \beta y_t + \alpha_1 E_t p_{t+1} + \alpha_2 p_{t-1} + v_t \\
 r_t &= \theta_1 y_{t-1} + \theta_2 p_{t-1} + \theta_3 r_{t-1} + \theta_4 y_{t-2} + w_t.
 \end{aligned}$$

SURF estimation places no restrictions on the parameters of the reaction function. **SORF** estimation requires that the parameters of the reaction function minimize the assumed loss function of the Federal Reserve. The values of the reaction function coefficients reported for **SORF** are those induced jointly by the model parameters, the loss function weights, and the requirement that the policy rule minimizes loss.

2. Loss function weights measure the relative importance to the Fed of stabilizing output and the interest rate. The weight on stabilizing inflation is normalized to 1.0

3. Parameter-estimate significance is measured with the p-value of the likelihood ratio test statistic for the hypothesis that the true value of the parameter is very small (0.0001).

4. The **Two Part** procedure estimates structural parameters in a first step and then estimates stabilization-objective weights in a second step while holding structural model parameters constant.

Table 4 Maximum Likelihood Estimates of Model Parameters First Difference of the Interest Rate Is the Dependent Variable of the Policy Rule				
	Estimation with Unrestricted Reaction Function SURF		Estimation with Optimal Reaction Function SORF	
Model Parameters	Estimate	Significance (p-value)	Estimate	Significance (p-value)
8	.441	.024	.185	.047
a₁	.704	.037 ⁴	1.08	
a₂	-.145	<.001	-.290	<.001
b	.007	<.001	.062	<.001
"₁	.003	1.00	.594	<.001
"₂	.621	<.001	.399	<.001
\$.056	<.001	.0002	.141
Log Likelihood	1462.94		1451.36	
<p>1. The table reports maximum likelihood estimates of the parameters of the following model.</p> $y_t = \lambda E_t y_{t+1} + a_1 y_{t-1} + a_2 y_{t-2} - b (r_t - E_t p_{t+1}) + u_t$ $p_t = \beta y_t + \alpha_1 E_t p_{t+1} + \alpha_2 p_{t-1} + v_t$ $\Delta r_t = \theta_1 y_{t-1} + \theta_2 p_{t-1} + \theta_3 \Delta r_{t-1} + \theta_4 y_{t-2} + w_t$ <p>SURF estimation places no restrictions on the parameters of the reaction function. SORF estimation requires that the parameters of the reaction function minimize the loss function of the Federal Reserve.</p> <p>2. The model parameter values reported in this table are assumed to be the same for both policy regimes. The parameters of the interest rate equation (2₁ - 2₄) are permitted to change across regimes. Parameter estimates for the reaction function are reported in Table 5.</p> <p>3. Parameter-estimate significance is measured with the p-value of the likelihood ratio test statistic for the hypothesis that the true value of the parameter is very small (0.0001).</p> <p>4. The p-value reported is for a test of the hypothesis that a₁ = 1.0.</p>				

Table 5
Policy Rule Coefficient Estimates
The Interest Rate First Difference Is the Dependent Variable of the Policy Rule

Reaction Function Coefficients	OLS		SURF		SORF	
	Coefficient	P-value	Coefficient	P-value	Coefficient	P-value
Regime One 2 ₁	.146 (.14)		.249	.028	.074	---
2 ₂	-.074 (.09)		-.015	.78	.010	---
2 ₃	.218 (.16)		.151	.254	.002	---
2 ₄	-.064 (.06)		-.077	.065	-.121	---
2 ₅	-.094 (.15)		-.210	.085	-.027	
Loss Function Weight on y	---	---	---	---	0.0001	1.00
Loss Function Weight on) r	---	---	---	---	87.4	<.001
Loss Function Weight on r					.8e-06	.355
Regime 2 2 ₁	.231 (.15)		.285	.039	.121	---
2 ₂	.241 (0.11)		.399	<.001	.396	---
2 ₃	.022 (0.11)		-.186	.052	-.142	---
2 ₄	-.188 (0.05)		-.159	.002	-.168	---
2 ₅	-.281 (.15)		-.336	.018	-.145	
Loss Function Weight on y	---	---	---	---	0.0002	<.001
Loss Function Weight on) r					-0.0176	.055
Loss Function Weight on r	---	---	---	---	0.0017	<.001
Log Likelihood	1475.88		1462.94		1451.36	

1. The table reports maximum likelihood estimates of the parameters of the following model.


$$\begin{aligned}
 y_t &= \lambda E_t y_{t+1} + a_1 y_{t-1} + a_2 y_{t-2} - b(r_t - E_t p_{t+1}) + u_t \\
 p_t &= \beta y_t + \alpha_1 E_t p_{t+1} + \alpha_2 p_{t-1} + v_t \\
 \Delta r_t &= \theta_1 y_{t-1} + \theta_2 p_{t-1} + \theta_3 \Delta r_{t-1} + \theta_4 r_{t-1} + \theta_5 y_{t-2} + w_t .
 \end{aligned}$$

2. Loss function weights measure the relative importance to the Fed of stabilizing output and the interest rate. The weight on stabilizing inflation is normalized to 1.0

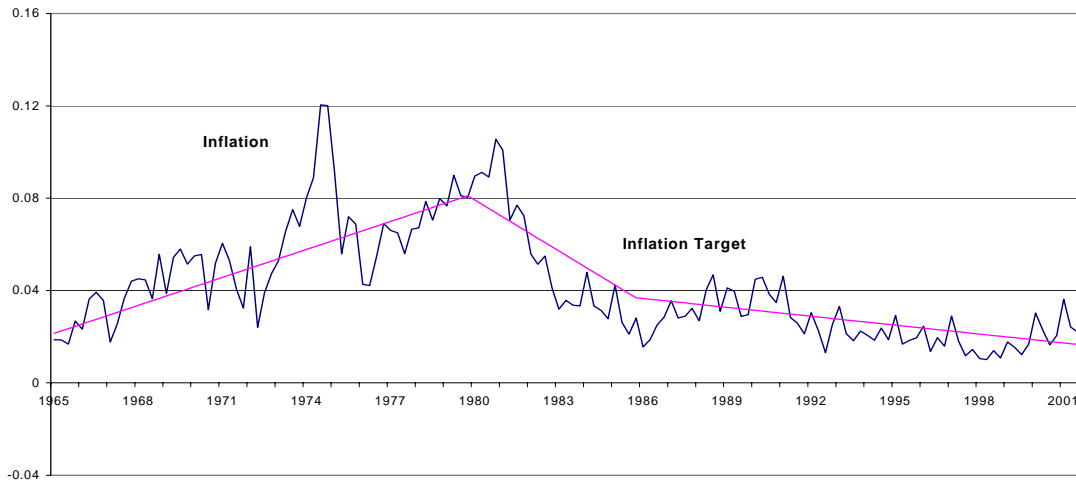
3. Parameter-estimate significance is measured with the p-value of the likelihood ratio test statistic for the hypothesis that the true value of the parameter is very small (0.0001).

Figure 1
Actual and Target Values of Variables

a: Output Gap



b: Inflation Rate



c: Interest Rate


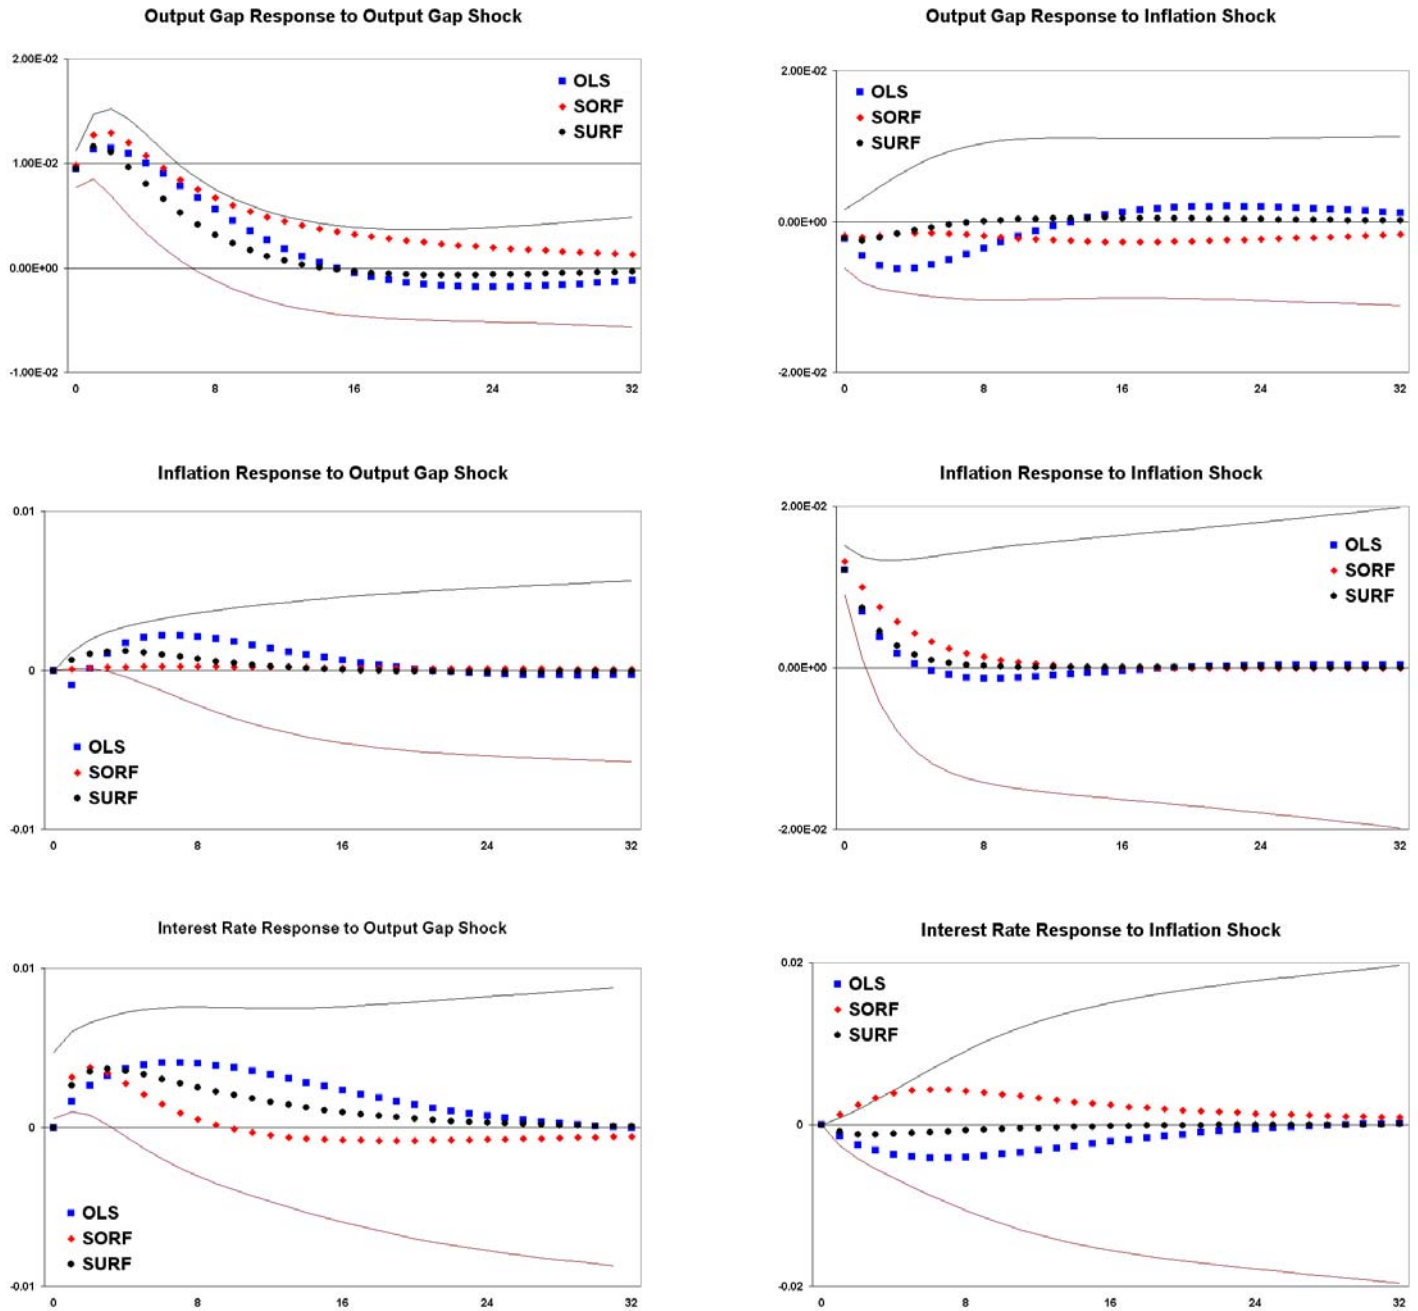
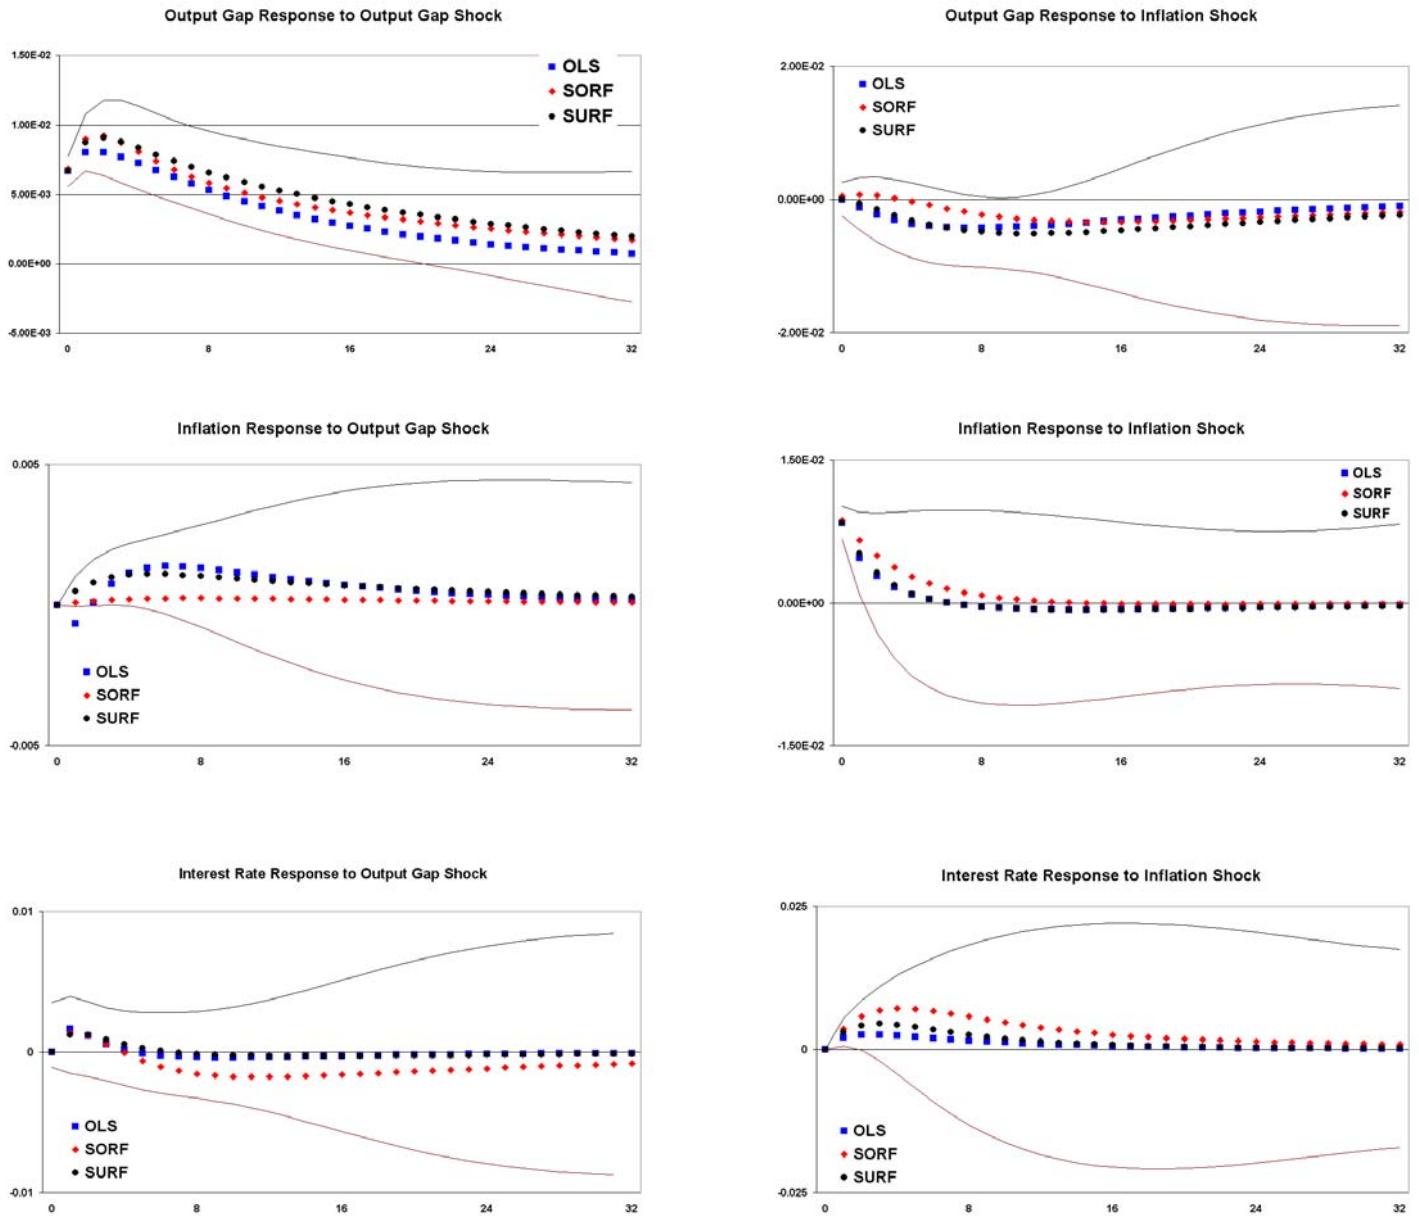


Figure 2
Comparison of Impulse Response Functions
Regime One (1965:I–1979:IV)



The figure presents impulse response functions implied by three estimates of the model reduced form. The IRFs labeled OLS are based on the OLS projections of the output gap, inflation and the interest rate on the state variables (two lags of the gap, and one lag each of inflation and the interest rate). The IRFs labeled SURF are based on the reduced form of the New Keynesian model estimated with an unrestricted reaction function for the interest rate. The IRFs labeled SORF are based on the reduced form of the New Keynesian model estimated subject to the restriction that the reaction function minimized expected loss. The thin lines are two-standard-deviation upper and lower confidence intervals estimated by the Monte Carlo procedure described in the text.

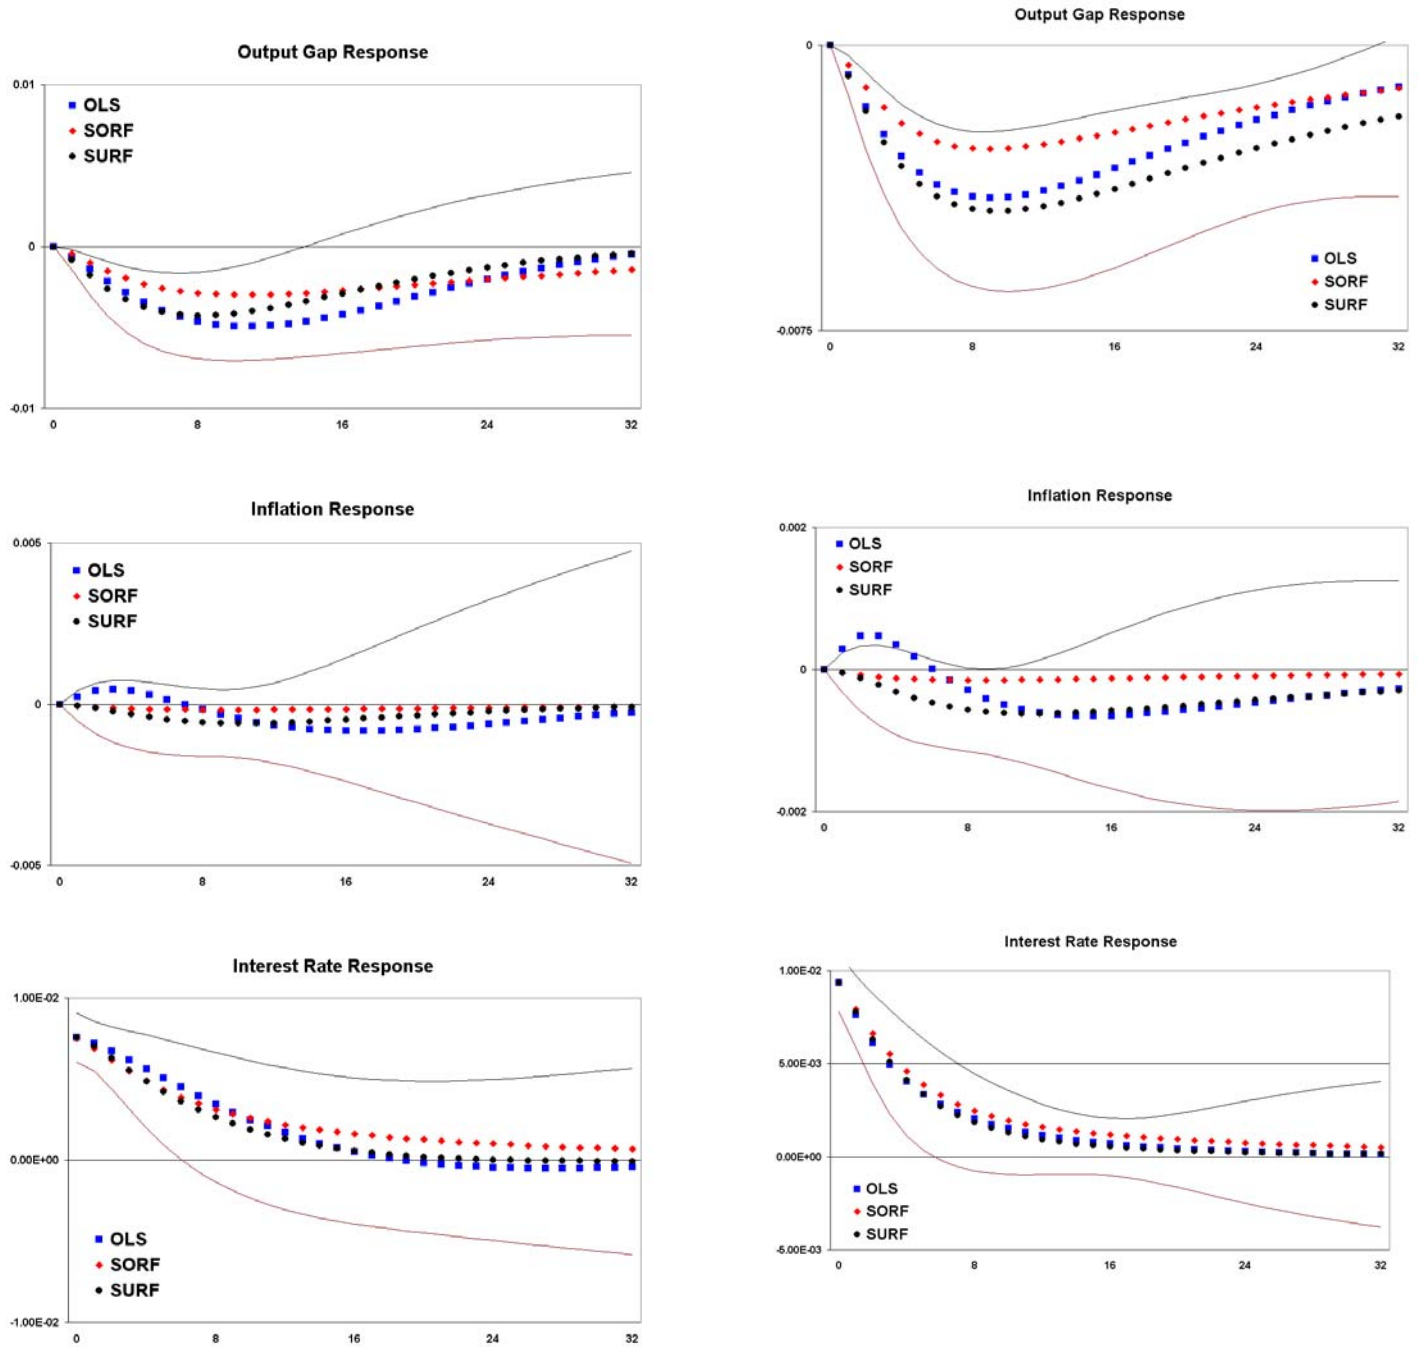
Figure 3
Comparison of Impulse Response Functions
Regime Two (1980:I–2001:IV)



The figure presents impulse response functions implied by three estimates of the model reduced form. The IRFs labeled OLS are based on the OLS projections of the output gap, inflation and the interest rate on the state variables (two lags of the gap, and one lag each of inflation and the interest rate). The IRFs labeled SURF are based on the reduced form of the New Keynesian model estimated with an unrestricted reaction function for the interest rate. The IRFs labeled SORF are based on the reduced form of the New Keynesian model estimated subject to the restriction that the reaction function minimized expected loss. The thin lines are two-standard-deviation upper and lower confidence intervals estimated by the Monte Carlo procedure described in the text.

Figure 4
Comparison of Responses to an Interest Rate Shock

Regime One (1965:I–1979:IV)



Regime Two (!980:I–2001:IV)

The figure presents impulse response functions implied by three estimates of the model reduced form. The IRFs labeled OLS are based on the OLS projections of the output gap, inflation and the interest rate on the state variables (two lags of the gap, and one lag each of inflation and the interest rate). The IRFs labeled SURF are based on the reduced form of the New Keynesian model estimated with an unrestricted reaction function for the interest rate. The IRFs labeled SORF are based on the reduced form of the New Keynesian model estimated subject to the restriction that the reaction function minimized expected loss. The thin lines are two-standard-deviation upper and lower confidence intervals estimated by the Monte Carlo procedure described in the text.

Figure 5
Ratio of Observed-Policy-Rule Loss to Optimal-Policy-Rule Loss

