Research Program on Forecasting

Jointly Evaluating the Federal Reserve's Forecasts of GDP Growth and Inflation

Edward N. Gamber
Department of Economics and Business
Lafayette College
Easton, PA 18042
gambere@lafayette.edu

Tara M. Sinclair
Department of Economics
George Washington University
Washington, DC 20052
tsinc@gwu.edu

H.O. Stekler
Department of Economics
George Washington University
Washington, DC 20052
hstekler@gwu.edu

Elizabeth Reid
Department of Economics
George Washington University
Washington, DC 20052
ereid@gwu.edu

RPF Working Paper No. 2008-002
http://www.gwu.edu/~forcpgm/2008-002.pdf

Original Version:*
April 15, 2008

This Version:
March 8, 2011

RESEARCH PROGRAM ON FORECASTING
Center of Economic Research
Department of Economics
The George Washington University
Washington, DC 20052
http://www.gwu.edu/~forcpgm

*This paper was originally titled: Multivariate Forecast Evaluation Based on Linear Policy Rules

Jointly Evaluating the Federal Reserve's Forecasts of GDP Growth and Inflation

Tara M. Sinclair[1]
Department of Economics
George Washington University
Washington, DC 20052
tsinc@gwu.edu

Edward N. Gamber
Department of Economics and Business
Lafayette College
Easton, PA 18042
gambere@lafayette.edu

H.O. Stekler
Department of Economics
George Washington University
Washington, DC 20052
hstekler@gwu.edu

Elizabeth Reid
Department of Economics
George Washington University
Washington, DC 20052
ereid@gwu.edu

JEL Codes: C53, E37, E52, E58

Keywords: Evaluating Forecasts, Macroeconomic Forecasts, Loss Function, Inflation Forecasting, GDP Growth Forecasting, Monetary Policy

This Draft:
March 8, 2011

Abstract

In this paper we jointly evaluate the Federal Reserve staff forecasts of U.S. real output growth and the inflation rate assuming the forecasts are to be used as inputs for the Taylor rule. Our simple methodology generates "policy forecast errors" which have a direct interpretation for the impact of forecast errors on the target interest rate given by the Taylor rule. Without interest rate smoothing, we find that, on average, the Taylor rule target interest rate would have been approximately a full percentage point away from the intended target because of errors in forecasting output growth and inflation. Our results are robust to changes in the forecast horizon and to changes in the weights on the variables in the policy rule.

[1] Corresponding author. Phone: 202-994-7988. Fax: 202-994-6147.

The authors wish to thank Mike Bradley, Dean Croushore, Fred Joutz, Kajal Lahiri, George Monokroussos, Keith Ord, Michael Owyang, Roberto Sameniego, Peter Tinsley, and Harry Watson for helpful discussions. We also thank participants in the GWU Brown Bag Seminar Series on Forecasting, the 16th Federal Forecasters Conference, the 28th Annual International Symposium on Forecasting, the 1st Macroeconomic Forecasting Conference in Rome, West Virginia University Economics Department Seminar, Richmond University Economics Department Seminar, University of Memphis Economics Department Seminar, University at Albany Economics Seminar, and Wesleyan University Economics Department Seminar for their comments and suggestions. Warren Carnow, Alexander Kapinos, and Blake Saville provided excellent research assistance. Generous support for this project was provided by the Institute for International Economic Policy at the Elliott School of International Affairs and the GW Institute of Public Policy.

1. Introduction

Forecasts are an essential input into all decisions involving the future. The accuracy of these forecasts can determine the adequacy, validity, and credibility of both macro and micro economic decisions. It is for this reason that the accuracy of these forecasts is evaluated. However, there are different perspectives or approaches for conducting these evaluations. In one approach, the forecaster's accuracy is evaluated while another method considers the value of the forecast to its users.

The traditional approach for evaluating accuracy has been to measure the magnitude of the error that is associated with the forecasts of *each* particular variable.[1] If an individual or organization forecasts both the inflation rate and the GDP growth rate for next year, the accuracy of the forecast of each variable is calculated separately. This is the approach that has been previously used to evaluate the forecasts of the Federal Reserve's staff (Clements et al., 2007; Joutz and Stekler, 2000; Romer and Romer, 2000; Sims, 2002; Stekler, 1994). [2]

There is, however, a problem that is associated with this traditional forecast evaluation approach: the loss function may contain more than one argument. This paper addresses this problem, focusing on the forecasts of the Federal Reserve (Fed). The Fed is concerned with both inflation and output growth. This suggests that the accuracy of both variables should be evaluated simultaneously. Sinclair et al. (2010) provided the procedures for jointly evaluating the directional accuracy of the Fed's forecasts but did not consider how to jointly evaluate the quantitative predictions.

[1] Usually mean square error is the error metric because it is associated with a quadratic loss function.

[2] While most evaluations have focused on the accuracy of a particular variable, there have also been discussions about the appropriate procedures for jointly evaluating multiple forecasts (See Hymans, 1968; Clements and Hendry and subsequent comments, 1993; Eisenbeis, et al., 2002; and Komunjer and Owyang, 2007) or for comparing forecasts of different variables (Lahiri and Sheng, 2010). However, these evaluations involving multiple variables were not done in the context of policy loss functions or decision rules.

If the intended use of the Fed's forecasts is known, it may be possible to jointly evaluate the accuracy of the quantitative growth rate and inflation forecasts. In the case of monetary policy, there has been considerable discussion that some form of the Taylor Rule is a good approximation for the decision rule used by the Fed in setting their target interest (i.e. federal funds) rate (see, for example, the collection of articles in Taylor, 1999). By assuming that this is the decision rule, we can jointly evaluate the GDP and inflation forecasts of the Fed and estimate the size of the errors resulting from the use of these forecasts in the Taylor rule. We ask the following question: Given this specific policy rule, what would be the impact of forecast errors on the implied federal funds rate target?[3]

Our work is in the spirit of evaluating forecasts by considering the economic costs of prediction errors (Clements, 2004; Granger and Pesaran, 2000a and 2000b; and Pesaran and Skouras, 2002). It is also related to the work by Orphanides (2001, 2003) on the impact of real-time data on monetary policy. Our contribution is to draw these two lines of research together to specifically evaluate forecasts in the context for which they are prepared.

We next present data used in this analysis. We evaluate the Greenbook forecasts that are made by the staff of the Board of Governors of the Fed. These forecasts are produced before each meeting of the Federal Open Market Committee and contain projections on the economy up to eight quarters into the future. We then present a methodology for evaluating policy forecast errors assuming that the Fed follows the Taylor rule. Subsequent sections present the results, robustness checks with respect to the specification of the policy rule, and our conclusions.

[3] There is much debate in the literature on the Taylor Rule as to the specific form of the rule. We show in section 6 that our results are robust to the weights chosen for the variables in the policy rule. We further show that interest rate smoothing will reduce the impact of the current forecast errors.

2. Data

We examine the Greenbook forecasts for each quarter from 1965.4 through 2002.4, but, for some analyses, also divide the data into two subperiods, 1965.4-1979.2 and 1979.3-2002.4.[4] These periods correspond to the pre-Volcker and Volcker-Greenspan eras at the Fed. The projections used in this analysis are the annualized growth rate of real output (GNP from 1965 to 1991 and GDP from 1992 on) and the annualized inflation rate (based on the implicit price deflator through the first quarter of 1996, then the chain-weighted price index from 1996.2 on).[5] We only analyze the current quarter and one quarter ahead projections.

We focus on short horizons because the Greenbook forecasts have sometimes been based on an assumed (possibly varying) path for monetary policy. At other times, however, the forecasts assumed that monetary policy would remain unchanged over its forecast horizon (Reifschneider and Tulip, 2007). Since the assumed path for monetary policy associated with each Greenbook forecast is not known, a possible complication arises when analyzing longer-term forecasts. The current quarter and one-quarter ahead forecasts are too short of a time horizon, however, to be affected by the Fed's future path for monetary policy. Therefore, regardless of whether the Fed assumes a constant path or a varying path for monetary policy, the current and one-quarter-ahead forecasts will be unaffected by those assumptions.

There are at least two forecasts per quarter, each made in various months of the quarter. Consequently the forecasts made for the current and next quarter have leads of zero to five months to the end of the relevant quarter. The horizons of the forecasts vary between zero and five months: horizons of (a) zero to two months and (b) three to five months for the current and next quarters, respectively. Because forecasts were not made at all horizons in every quarter, the

[4] The Greenbook data are only available with at least a 5-year lag.
[5] We use the compounded annual rate of change to construct the growth rates.

number of observations differs between horizons. The actual values were the data available approximately 90 days after the end of the quarter to which they refer.[6]

3. The Fed's Policy Rule

> *"Even if the actual policy process is far more sophisticated than any simple rule could completely describe, the [Taylor] rule often provides a reasonably good approximation of what policymakers decide and a framework for thinking about policy actions."*
>
> --Federal Reserve Board Vice Chairman Donald L. Kohn (2007)

In order to analyze the effects of the Fed's forecasts errors on its selection of a target interest rate, it is necessary to first establish the relevant policy rule. We assume the Fed either follows the original Taylor rule (Taylor, 1993) or uses it as a guide in setting monetary policy. Although Taylor originally proposed his rule as an empirical description of past Fed policy actions, Woodford (2001a, 2001b) has shown that the Taylor rule can also be theoretically justified as a prescriptive rule. Clarida et al (2000) suggested that it should be forward-looking, employing expected future values of both output and inflation. As pointed out by Orphanides (2001), even if the Fed is using current values of both output and inflation, the data the Fed has at the time they are setting policy is not the actual values, but predictions of those values. Rather than focusing on the various issues associated with this rule, we start with the fixed weights proposed by Taylor (1993). We show in Section 6 that our results are robust to the range of weights considered in the literature.

[6] The use of the real time data avoids definitional and classification changes. Our results are robust to when the data published 45-60 days after the end of the quarter are used. Our real-time data were obtained from the ArchivaL Federal Reserve Economic Data (ALFRED®), maintained by the Federal Reserve Bank of St. Louis at http://alfred.stlouisfed.org/.

According to the forward-looking Taylor rule, the Fed sets a target federal funds rate, i_t^{Tf}, based on equation (1), where, the superscript "f" denotes that the target is based on forecasted variables.[7] The Fed's policy decision $(P_{t,t+h}^f)$ is written as:

$$P_{t,t+h}^f = i_t^{Tf} = r^* + \pi_{t+h}^f + 0.5(\pi_{t+h}^f - \pi^*) + 0.5(\ln Y_{t+h}^f - \ln Y^*), \qquad (1)$$

where r* is the equilibrium real interest rate, π^* is the Fed's implicit inflation rate target, and Y* is potential output. While the output gap (i.e. lnY – lnY*) is typically used in the Taylor rule, the growth rate of real GDP is typically used in forecast evaluation. For our application we replace the output gap by the GDP growth rates, yielding the following policy rule:

$$P_{t,t+h}^f = i_t^{Tf} = r^* + \pi_{t+h}^f + 0.5(\pi_{t+h}^f - \pi^*) + 0.5(y_{t+h}^f - y^*), \qquad (1')$$

where y* is the potential output growth rate.[8] The Fed forecasts both inflation, π_{t+h}^f, and output growth, y_{t+h}^f, h periods ahead.

The actual outcomes, however, will in general differ from the Fed's forecasts. If the actual values for π_{t+h} and y_{t+h} were instead included in the Taylor rule (i.e. if the Fed had perfect forecasts or perfect foresight), a potentially different target federal funds rate would have been suggested by the policy rule. Consequently, under perfect foresight the implied interest rate target would have been:

$$P_{t,t+h} = i_t^T = r^* + \pi_{t+h}^A + 0.5(\pi_{t+h}^A - \pi^*) + 0.5(y_{t+h}^A - y^*), \qquad (2)$$

[7] Following Orphanides (2001), we assume that the Fed uses the Greenbook forecasts in their decision rule.
[8] If we assume that potential output, Y*, is either known or a constant, then we can difference out potential output which permits us to use the growth rate in order to construct the policy rule. A similar approach has been used by Hamilton et al. (2009). For a discussion of the role of real time output gap estimates and the Taylor rule, see Orphanides (2001).

where π^A_{t+h} and y^A_{t+h} represent the actual realizations of π_{t+h} and y_{t+h}. The difference between i^{Tf} and i^T measures the difference in the Fed funds rate suggested by the Taylor rule that occurs because of inaccurate forecasts and thus represents the implied policy forecast error, PFE_t:

$$PFE_t = i^T_t - i^{Tf}_t = 1.5\left(\pi^A_{t+h} - \pi^f_{t+h}\right) + 0.5\left(y^A_{t+h} - y^f_{t+h}\right). \qquad (3)$$

The differences, $\left(\pi^A_{t+h} - \pi^f_{t+h}\right)$ and $\left(y^A_{t+h} - y^f_{t+h}\right)$, are the Fed's forecast errors for the inflation rate and real output growth respectively. Given the PFEs, the evaluation procedures are then similar to those used in judging individual forecast errors. [9]

4. Evaluation Procedure: Standards of Comparison

There are two sets of forecasts that we use as standards of comparison: the PFEs that would have occurred if (1) naïve forecasts or (2) the median forecasts from the Survey of Professional Forecasters[10] had been used to set policy. In each case, the comparison uses the same weights and has the same number of observations as were used in constructing the PFEs of the Fed forecasts that we are evaluating. We calculate the mean absolute policy forecast error (MAPFE) and root mean squared policy forecast error (RMSPFE) for the full sample as well as for a sub-sample where we have the median forecasts from the Survey of Professional Forecasters (SPF) available for comparison.

The naïve forecast assumes that the same growth rate of output and the same inflation rate as were observed in the previous quarter will occur in the future period(s).[11] We also use the

[9] For a discussion of the impact of using projections on the estimates of the Taylor rule, see Orphanides and Wieland (2008). We do not use the interest rate, the equilibrium real interest rate, the Fed's implicit inflation rate target, or potential output in our analysis. Considerable research has gone into estimating these variables (e.g. Clark and Kozicki, 2005, for the equilibrium real interest rate and Leigh, 2008, for the Fed's implicit inflation rate target). While these variables may be time-varying, we treat them here as known to the policymaker and they thus drop out of our expression. This assumption allows us to isolate the impact of forecast errors.

[10] These data come from the Federal Reserve Bank of Philadelphia (for a description of this survey, see Croushore, 1993).

[11] We construct these naïve "forecasts" from the real-time data, published in the Greenbook, of GDP growth and the inflation rate for the quarter before the Greenbook forecasts were made. This assures that the same information as

median forecasts from the SPF to construct forecasts comparable to our Greenbook data. Since the SPF forecasts are only reported once per quarter (in the middle of the second month of each quarter), they are comparable to the Fed forecasts made with one and four month horizons. The SPF forecasts are available beginning with the fourth quarter of 1968.

5. Results

For the entire sample 1965:4-2002:4, the mean absolute policy forecast error (MAPFE) based on Fed forecasts at the zero month horizon is 101 basis points (Table 1). This can be interpreted as the amount by which the implied target interest rate based on the Taylor rule would have differed by using forecasts rather than the actual data in the policy rule.[12] According to this analysis, the implied target interest rate based on the Taylor rule was at least a full percentage point off because of errors in forecasting growth and inflation at the zero horizon.[13] We also tested for bias, but the null of no bias was usually not rejected.

Table 1 also presents the root mean squared policy forecast errors (RMSPFE) for the samples and horizons we analyze in our subsequent tests. As expected, both the MAPFE and RMSPFE in general increase with the length of the forecast horizon and the MAPFE and RMSPFE are both smaller than those of the naïve and SPF forecasts in all cases. This result is consistent with previous findings from evaluations of a single variable.

A comparison of the results for the Pre-Volcker sample (Panel B) and the Volcker-Greenspan sample (Panel C) suggests that there has been a small reduction of both MAPFE and

---

was available to the Fed forecasters at the time the predictions were made was used in constructing this alternative forecast.

[12] Bernanke (2010) pointed out that using forecasts in the Taylor rule instead of actual values may explain the difference in the interpretation of the stance of monetary policy from 2002 to 2006. Based on forecasts it appears that the Fed was close to the recommended policy in the Taylor rule. Therefore the intended policy of the Fed was to approximately follow the Taylor rule. If, as Taylor (2007) points out, however, we use the actual values for output and inflation in the Taylor rule, it appears that monetary policy was much looser than prescribed by the Taylor rule.

[13] This result holds even if we look at the 4-quarter average of the forecasts instead of just the current quarter. The volatility of the forecast errors is reduced with the 4-quarter average, but the size of the errors increases and offsets the volatility reduction.

RMSPFE in the more recent period. The similarity of these results with the full sample results, however, suggests that our findings for the full-sample are not dependent upon a specific time period.[14] Furthermore, the improvement is also present for the naïve forecasts, which suggests no relative improvement on the part of the Fed.

6. Alternative Specifications

In our analysis we have chosen to focus on the policy rule as initially formulated in Taylor (1993). Subsequently, a number of authors have estimated the policy rule that was actually used by the Fed (e.g. Clarida et al, 2000; Orphanides, 2001). Two key issues are the size of the policy rule weights and the role of interest rate smoothing.[15]

6A. Alternative Policy Rule Weights

One focus of this literature, and the part relevant to our analysis, is the weights associated with inflation and output in the Taylor rule. We performed a simple calculation to examine the effect on the PFE of using weights other than those of the original Taylor Rule. The PFE is a weighted sum of the individual forecast errors, consequently we ask: Given the individual forecast errors, what weights would minimize the impact of the forecast errors on the implied target interest rate? We thus determined the weights that minimize the RMSPFE of the Fed's forecasts, assuming that the weights summed to two.[16] We calculated different weights for each horizon and allowed for both the case where the weights are constant for the full sample and allowing the weights to differ between the pre-Volcker and Volcker-Greenspan periods. In all

[14] These subsamples are representative of several alternative subsamples that we explored.

[15] Another concern, beyond the scope of this paper, is that the optimal Taylor rule may not be linear (Nobay and Peel, 2003). There has been a substantial amount of research in evaluating forecasts in the case of nonlinear loss functions (e.g. Elliott et al., 2005 and 2008; Patton and Timmermann, 20071), including evidence that the Fed's loss function may not be linear, with respect to inflation (Capistran, 2008) and with respect to output growth (Patton and Timermann, 2007b) when the forecasts are evaluated separately.

[16] Almost identical results were found when we instead minimized the MAPFE. The restriction that the coefficients sum to two is consistent with the literature where every estimate or prescriptive rule we found suggests the coefficients sum to something greater than or equal to two. We chose the lower bound so that we are truly minimizing the RMSPFE.

cases the weights are very similar to the Taylor weights. These weights do not necessarily imply anything about actual Fed behavior, but they allow us to set a lower bound on the impact of forecast errors on the implied target interest rate.[17] The results are very similar to what we found using the Taylor weights. Based on this analysis we still conclude that, without smoothing, the impact of forecast errors alone on the implied target interest rate based on the Taylor rule is approximately 100 basis points.

6B. Interest Rate Smoothing

According to Giannoni and Woodford (2002), the application of certainty equivalence requires interest rate smoothing. Therefore, suppose the policy rule is modified to allow for interest rate smoothing. The policy forecast error in this case would be a simple proportion of the policy forecast error without smoothing:

$$PFE_t = i_t^T - i_t^{Tf} = (1-\rho)\left(1.5\left(\pi_{t+h}^A - \pi_{t+h}^f\right) + 0.5\left(y_{t+h}^A - y_{t+h}^f\right)\right), \qquad (3')$$

Estimates of the smoothing parameter, ρ, have varied widely in the literature (e.g. Clarida et al, 1998, find ρ = 0.8, whereas Rudebusch, 2001, argues that the smoothing parameter may be negligible in practice). For example, if the smoothing parameter is 0.5, then the impact of the forecast errors on the policy decision is reduced by half. Thus, the larger the forecast errors are expected to be, the greater the benefit of smoothing.[18]

7. Conclusions

In this paper we developed a simple methodology in order to evaluate the impact of forecast errors on the Fed's monetary policy as characterized by a policy rule that is assumed to take the form of a Taylor rule. We find that the Fed's policy forecast error is smaller than the

[17] In particular, the specific weights found by Clarida et al (2000) were in the set of potential weights we considered. Their weights resulted in a higher MAPFE than those that minimized the MAPFE. The weights, RMSPFEs, and MAPFEs are all available from the authors upon request.
[18] This assumes, however, that the past interest rate was set in an appropriate way. As noted by Orphanides and Williams (2007), interest rate smoothing can also cause past errors to persist.

errors that would have resulted from naïve forecasts or SPF predictions. Nevertheless, the mean absolute policy forecast error of the Fed forecasts is approximately 100 basis points. The results for selected subsamples were similar to those for the entire period and a simple calculation showed that the weights placed on the two variables that minimize the PFE were similar to those contained in the original Taylor Rule. Thus we find the robust result that, without interest-rate smoothing, forecast errors contribute a mean absolute error of approximately 100 basis points to target interest rates set by using a Taylor Rule.

References

Bernanke, Ben S. (2010). "Monetary Policy and the Housing Bubble" speech at the Annual Meeting of the American Economic Association. Atlanta, Georgia, January 3, 2010. External Link http://www.federalreserve.gov/newsevents/speech/bernanke20100103a.htm.

Capistran, C. (2008). Bias in Federal Reserve inflation forecasts: Is the Federal Reserve irrational or just cautious? Journal of Monetary Economics, 55, 1415-1427.

Clarida, R., Gali, J. and Gertler, M. (1998). Monetary Policy Rules in Practice: Some International Evidence. European Economic Review 42, 1033–67.

Clarida, R., Gali, J. and Gertler, M. (2000). Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. The Quarterly Journal of Economics 115(1), 147-180.

Clark, T. E. and Kozicki, S. (2005). Estimating Equilibrium Real Interest Rates in Real Time. North American Journal of Economics and Finance 16(3), 395-413.

Clements, M. P. (2004). Evaluating The Bank Of England Density Forecasts Of Inflation. The Economic Journal, 114 (October), 855–877.

Clements, M. P. and Hendry, D. F. (1993). On the Limitations of Comparing Mean Square Forecast Errors, and subsequent comments. Journal of Forecasting, 12(8), 617-676.

Clements, M. P., Joutz, F. and Stekler, H. O. (2007). An Evaluation of the Forecasts of the Federal Reserve: A Pooled Approach. Journal of Applied Econometrics 22(1), 121-136.

Croushore, D. (1993). Introducing: The Survey of Professional Forecasters. Federal Reserve Bank of Philadelphia Business Review, November/December, 3-13.

Eisenbeis, R., Waggoner, D. and Zha, T. (2002). Evaluating Wall Street Journal Survey Forecasters: A Multivariate Approach. Business Economics, 37(3), 11-21.

Elliott, G., Komunjer, I. and Timmermann, A. (2005). Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss? Journal of the European Economic Association, 6(1), 122-157.

Elliott, G., Komunjer, I. and Timmermann, A. (2008). Estimation and Testing of Forecast Rationality under Flexible Loss. Review of Economic Studies, 72, 1107-1125.

Giannoni, M. P. and Woodford M. (2002). Optimal Interest Rate Rules: II. Applications. NBER Working Paper No. 9420.

Granger, C. W. J. and Pesaran, M. H. (2000a). A decision-based approach to forecast evaluation. In (W.S. Chan, W. K. Li, and H. Tong, eds). Statistics and Finance: An Interface. London: Imperial College Press.

Granger, C. W. J. and Pesaran, M. H. (2000b). Economic and statistical measures of forecast accuracy. Journal of Forecasting, 19, 537–60.

Hamilton, J. D., Pruitt, S. and Borger, S. C. (2009). The Market-Perceived Monetary Policy Rule. Board of Governors of the Federal Reserve System International Finance Discussion Papers Number 982.

Hymans, Saul H. (1968). Simultaneous Confidence Intervals in Econometric Forecasting. Econometrica, 36(1), 18-30.

Joutz, F. and Stekler, H.O. (2000). An Evaluation of the Predictions of the Federal Reserve. International Journal of Forecasting, 16, 17-38.

Kohn, D. L. (2007). “John Taylor Rules” a speech at the Conference on John Taylor's Contributions to Monetary Theory and Policy. Federal Reserve Bank of Dallas, Dallas, Texas, October 12, 2007.

Komunjer, I. and Owyang M. T. (2007). Multivariate Forecast Evaluation and Rationality Testing. Federal Reserve Bank of St. Louis Working Paper No. 2007-047.

Lahiri, K. and Sheng, X. (2010). Learning and Heterogeneity in GDP and Inflation Forecasts. International Journal of Forecasting (special issue on Bayesian Forecasting in Economics), 26, 265-292.

Leigh, D. (2008). Estimating the Federal Reserve's implicit inflation target: A state space approach. Journal of Economic Dynamics and Control, 32(6), 2013-2030.

Nobay, A. Robert and Peel, David A. (2003). Optimal Discretionary Monetary Policy in a Model of Asymmetric Central Bank Preferences. The Economic Journal, 113 (489), 657–665.

Orphanides, A. (2001). Monetary Policy Rules Based on Real-Time Data. The American Economic Review, 91(4), 964-985.

Orphanides, A. (2003). Monetary Policy Evaluation with Noisy Information. Journal of Monetary Economics, 50, 605-631.

Orphanides, A., and Weiland, V. (2008). Economic Projections and Rules of Thumb for Monetary Policy. Federal Reserve Bank of St. Louis Review, 90(4), 307-24.

Orphanides, A. and Williams, J. C. (2007). Robust Monetary Policy with Imperfect Knowledge. Journal of Monetary Economics, 54, 1406-1435.

Patton, Andrew J. and Timmermann, Allan (2007a). Properties of optimal forecasts under asymmetric loss and nonlinearity. Journal of Econometrics, 140, 884–918.

Patton, Andrew J. and Timmermann, Allan (2007b). Testing Forecast Optimality Under Unknown Loss. Journal of the American Statistical Association, 102 (480), Applications and Case Studies, 1172-1184.

Pesaran, M. H. and Skouras, S. (2002). Decision-based methods for forecast evaluation. In (M.P. Clements, and D.F. Hendry, eds) A Companion to Economic Forecasting, Oxford: Blackwell, pp. 241–67.

Reifschneider, D. and Tulip, P. (2007). Gauging the Uncertainty of the Economic Outlook from Historical Forecasting Errors. Federal Reserve Board Finance and Economics Discussion Series, 2007-60.

Romer, C. D. and Romer, D. H. (2000). Federal Reserve Information and the Behavior of Interest Rates. American Economic Review, 90, 429-457.

Rudebusch, G. D. (2001). Term Structure Evidence on Interest Rate Smoothing and Monetary Policy Inertia. Journal of Monetary Economics 49(6), 1161-1187.

Sims, C. A. (2002). The Role of Models and Probabilities in the Monetary Policy Process. Brookings Papers on Economic Activity, 2, 1-40.

Sinclair, T. M., Stekler, H.O. and Kitzinger, L. (2010). Directional Forecasts of GDP and Inflation: A Joint Evaluation with an Application to Federal Reserve Predictions. Applied Economics 42(18), 2289 - 2297.

Stekler, H.O. (1994). Are economic forecasts valuable? Journal of Forecasting 13, 495-505.

Taylor, J.B. (1993). Discretion versus Policy Rules in Practice. Carnegie-Rochester Conference Series on Public Policy, 39, 195-214.

Taylor, J. B. (1999). Monetary Policy Rules. University of Chicago Press.

John B. Taylor, (2007). Housing and monetary policy. Proceedings, Federal Reserve Bank of Kansas City, 463-476.

Woodford, M. (2001a). The Taylor Rule and Optimal Monetary Policy. American Economic Review, 91 (2), 232-237.

Woodford, M. (2001b). Inflation Stabilization and Welfare. National Bureau of Economic Research Working Paper 8071.

Table 1: Mean Absolute Policy Forecast Error (MAPFE) and Root Mean Squared Policy Forecast Error (RMSPFE)

Panel A: Full Sample (1965.4 – 2002.4)[19]

Horizon (months)	$MAPFE_{Fed}$	$MAPFE_{naïve}$	$MAPFE_{SPF}$	$RMSPFE_{Fed}$	$RMSPFE_{naïve}$	$RMSPFE_{SPF}$	N
0	1.01	1.95	---	1.31	2.77	---	159
1	1.25	2.08	1.54	1.67	2.83	1.99	110[20]
2	1.54	1.99	---	1.93	2.68	---	94
3	1.56	2.31	---	2.08	3.03	---	159
4	1.83	2.41	2.21	2.34	3.15	2.77	105[21]
5	1.80	2.38	---	2.43	3.13	---	88

Panel B: Pre-Volcker Sample (1965.4 – 1979.2)[19]

Horizon (months)	$MAPFE_{Fed}$	$MAPFE_{naïve}$	$MAPFE_{SPF}$	$RMSPFE_{Fed}$	$RMSPFE_{naïve}$	$RMSPFE_{SPF}$	N
0	1.03	2.06	---	1.35	2.98	---	68
1	1.38	2.26	1.95	1.73	3.18	2.36	54[22]
2	1.71	2.21	---	2.15	2.97	---	51
3	1.84	2.70	---	2.38	3.46	---	67
4	2.09	2.74	2.78	2.69	3.56	3.36	48[23]
5	2.38	2.81	---	3.01	3.58	---	44

Panel C: Volcker-Greenspan Sample (1979.3 – 2002.4)

Horizon (months)	$MAPFE_{Fed}$	$MAPFE_{naïve}$	$MAPFE_{SPF}$	$RMSPFE_{Fed}$	$RMSPFE_{naïve}$	$RMSPFE_{SPF}$	N
0	1.00	1.87	---	1.29	2.60	---	91
1	1.12	1.92	1.23	1.62	2.45	1.66	56
2	1.35	1.73	---	1.65	2.29	---	43
3	1.36	2.02	---	1.82	2.67	---	92
4	1.61	2.14	1.81	2.01	2.75	2.26	57
5	1.23	1.96	---	1.65	2.61	---	44

[19] SPF sample starts with forecasts made in 1968.4.
[20] N = 98 for SPF.
[21] N = 97 for SPF.
[22] N = 42 for SPF.
[23] N = 40 for SPF.