

UC San Diego

Recent Work

Title

Bias in Federal Reserve Inflation Forecasts: Is the Federal Reserve Irrational or Just Cautious?

Permalink

<https://escholarship.org/uc/item/6v28v0b6>

Author

Carmona, Carlos Capistran

Publication Date

2005-07-01

Bias in Federal Reserve Inflation Forecasts: Is the Federal Reserve Irrational or Just Cautious?*

Carlos Capistrán Carmona[†]

July, 2005

Abstract

Inflation forecasts of the Federal Reserve systematically under-predicted inflation before Volcker and systematically over-predicted it afterward. Furthermore, under quadratic loss, commercial forecasts have information not contained in those forecasts. To investigate the cause, this paper recovers the loss function implied by Federal Reserve's forecasts. It finds that the cost of having inflation above an implicit time-varying target was larger than the cost of having inflation below it for the period since Volcker, and that the opposite was true for the pre-Volcker era. Once these asymmetries are taken into account, the Federal Reserve is found to be rational. (*JEL* C53, E52)

*I am grateful to Graham Elliott and to Allan Timmermann for their advice. Special thanks to Clive W.J. Granger for many useful conversations. I thank Daniel Chiquiar, Charles Elkan, Marjorie Flavin, James D. Hamilton, Nir Jaimovich, Craig R.M. McKenzie, George Monokroussos, Valerie A. Ramey, Christina D. Romer, Barbara Rossi, Kevin K. Sheppard and seminar participants at Colegio de México, Banco de México, Bank of Canada, the 2005 Society for Computational Economics conference in Washington D.C., and UCSD for helpful comments. Support provided by CONACYT and UC-MEXUS is gratefully acknowledged

[†]Department of Economics, University of California, San Diego. 9500 Gilman Drive, La Jolla, CA 92093-0508. Email: ccapistr@econ.ucsd.edu.

1 Introduction

One of the most important objectives of the Federal Reserve is to achieve stable prices. However, because inflation responds to monetary policy only after a lag, the Federal Reserve needs to make decisions based on forecasts of future inflation behavior. As noted by Chairman Greenspan, a successful monetary policy depends on the Federal Reserve's ability to produce forecasts that accurately reflect the information available at the time a decision has to be made.

The general perception in economics, supported by Romer and Romer (2000) and Sims (2002), is that Federal Reserve inflation forecasts are quite good. The Romers find that Federal Reserve forecasts of inflation are unbiased, and conclude that the forecasts are rational. They also find that if one had access to inflation forecasts from the Federal Reserve and from commercial forecasters the optimal combination would be to dispose of the commercial forecasts and use only Federal Reserve forecasts, a result maintained by Sims.¹ These results imply that the Federal Reserve uses information efficiently and that it has more information than commercial forecasters.

However, closer inspection of a data set that extends the one used by the Romers and by Sims shows that rationality can be rejected. This is not because of the new data, but because there is a change in behavior in Federal Reserve's forecast errors that coincides with Paul Volcker's appointment as Chairman and that was previously overlooked. The forecasts systematically under-predicted inflation before Volcker and systematically over-predicted it afterward (i.e., they are biased). Moreover, once this change in behavior is taken into account, Federal Reserve inflation forecasts do not seem to efficiently incorporate information contained in inflation forecasts from the Survey of Professional Forecasters, an important group of commercial forecasters. In particular, the forecasts from the Federal Reserve seem to miss information contained in the consensus forecast and in the spread across the surveyed forecasters. These results are consistent regardless of whether real-time or revised data are used for the actual values of inflation.

The bias found in Federal Reserve inflation forecasts is statistically significant and, at about half a percent for the sample since Volcker, is also economically significant. Biased forecasts are typically considered irrational because they denote a failure to incorporate information contained in past forecast errors. However, this is not always true. The property that rational forecasts have to be unbiased follows from the well-known result that under a quadratic loss function the optimal forecast is the conditional mean.² But the optimal forecast is not the conditional mean if the loss function is asymmetric in the sense that errors of the same magnitude but of different signs imply different costs. Under asymmetric

¹The Romers also find that commercial forecasts are unbiased, and conclude that they are rational.

²The theory of Rational Expectations says that rational agents have expectations that are optimal forecasts. See Muth (1961) on the definition of Rational Expectations and Mishkin (1981) on using it for testing the rationality of forecasts.

loss the optimal forecast is the mean plus an optimal bias term.^{3,4}

Most papers that test rational expectations using forecasts as proxies for expectations, such as those by the Romers and Sims, implicitly assume quadratic loss.⁵ But, does it make sense for the Federal Reserve to have symmetric preferences? Some authors have argued that it does not. Blinder (1998) recalls his experience as Vice-Chairman of the Federal Reserve and explains that a central bank is more likely to “... take far more political heat when it tightens preemptively to avoid higher inflation than when it eases preemptively to avoid higher unemployment” (Blinder 1998, p. 19). Nobay and Peel (2003) provide anecdotal support for the argument that both the European Central Bank and the Bank of England may have asymmetric preferences. Both papers indicate that for an inflation targeting central bank inflation below the target is less costly than inflation above it. Finally, Ruge-Murcia (2000) finds evidence that, in practice, Canada’s central bank “... may attach different weights to positive and negative inflation deviations from the target.” (Ruge-Murcia 2000, p. 1). In a later paper, Ruge-Murcia (2003) finds empirical evidence of asymmetric costs for Canada, Sweden and the United Kingdom.

This paper follows this literature in postulating a model of an inflation targeting central bank with asymmetric preferences which is used to reconcile the evidence of inefficient use of information on the part of the Federal Reserve. To make the model a better description of the behavior of the Federal Reserve, a time-varying inflation target is used, as in Gürkaynak, Sack, and Swanson (2005). The model shows that a negative bias in the forecasts (systematic over-prediction) is rational if the central bank is cautious in the sense that inflation above the target is considered more costly than inflation below the target. The mechanism at work is the following: Take an inflation targeting central bank that sets its monetary policy instrument so that the forecast of inflation equals the target, as in Svensson’s (1997) “inflation forecast targeting” framework. If for the central bank inflation above the target is as costly as inflation below it (i.e., the central bank has symmetric loss), then it would set its instrument so that the expected value of inflation equals the target. In this case the forecast coincides with the expected value of inflation. However, if inflation above the target is more costly than inflation below it (i.e., the central bank has asymmetric loss), then the central bank would, as a precautionary move, set the instrument so that the expected value of inflation is below the target. In this case, the forecast does not coincide with the expected value of inflation and hence a rational bias exists.

To investigate if the empirical evidence is consistent with an asymmetric-cost Federal Reserve, this paper backs out the Federal Reserve’s loss function as implied by its forecasts. The method used is to derive moment conditions from the model under an asymmetric quadratic loss function that nests the traditional quadratic loss as a special case. Elliott, Komunjer and Timmermann (2005) suggest this method to test for the presence of asymmetric costs and, jointly, to test for rationality. The moment conditions for each forecast horizon are

³See Christoffersen and Diebold (1997), Granger (1969, 1999), and Zellner (1986).

⁴Other papers present evidence that the evaluation of forecasts depends on the loss function. Leitch and Tanner (1991) find that forecasts that appear to be bad forecasts under traditional measures, like mean squared error, are not so under other measures, like the profits they generate to firms that use them. Keane and Runkle (1990), analyzing commercial price forecasts, indicate that a biased forecasts is consistent with rationality under asymmetric loss.

⁵See the survey by Stekler (2002).

used to form a system of equations and Hansen's (1982) Generalized Method of Moments (GMM) is applied to estimate the parameter that rationalizes the bias and to test if it implies a symmetric loss. A constant, past forecast errors, and forecasts from the Survey of Professional Forecasters are used as instruments. Tests of over-identifying restrictions based on the minimal value of the sample criterion function (Hansen's J test) are used to perform rationality tests.

The empirical results are that starting with Volcker's appointment as Chairman, the Federal Reserve's cost of under-prediction is four times the cost of over-prediction. For the pre-Volcker era the result is that the cost of under-prediction was a third of that of over-prediction, thus supporting the presence of asymmetric costs in both periods. These results imply that for the Federal Reserve since Volcker the cost of having inflation above the target was larger than the cost of having inflation below it, and that the opposite was true for the pre-Volcker era. Hence, this paper provides an empirical reason to move away from quadratic loss. The over-identification tests are not able to reject the hypothesis that, once the asymmetries are taken into account, the Federal Reserve is using information efficiently both before and after Volcker.

The finding of an asymmetric-cost Federal Reserve has several implications. First, it implies a deflationary bias during the Volcker-Greenspan periods, as equilibrium inflation is pushed below the target as a result of the precautions taken to avoid high inflation. The opposite will be true for the pre-Volcker period. Second, higher order moments of inflation affect the mean of inflation in equilibrium. This is a result of the Federal Reserve having to consider summary statistics not only of the location of inflation, but also of its dispersion, as the optimal bias is a function of the latter. Finally, Federal Reserve inflation forecasts are not the Federal Reserve's expected values of inflation, and a de-biasing mechanism is needed to recuperate the expected values from the forecasts.

The paper proceeds as follows. In section 1, the empirical properties of the Federal Reserve forecast errors are analyzed under quadratic loss and their biases and lack of encompassing of commercial forecasts are documented. To rationalize this evidence, in section 2 the loss function implied by Federal Reserve inflation forecasts is backed out and evidence is found of asymmetric costs of under- and over-prediction. Once these costs are taken into account, the Federal Reserve is found to be using information efficiently. Section 3 discusses the implications of asymmetric costs and considers alternative explanations for the empirical findings (e.g., learning by the Federal Reserve), and argues that they have difficulties to explain the duration and the change of sign of the bias. Section 4 is the conclusion.

2 Empirical Evidence on the Properties of Federal Reserve Inflation Forecasts

Federal Reserve forecasts are contained in the "Green Book" prepared by the staff of the Board of Governors before each meeting of the Federal Open Market Committee (FOMC). The forecasts are made with an assumption about monetary policy, and are judgmental in the sense that they are not the direct output of an econometric model, but the product

of judgmental adjustments made to forecasts obtained from econometric models.⁶ It is the policy of the Federal Reserve (the Fed) to release the forecasts to the public with a five year lag. The Federal Reserve of Philadelphia has put together a series of Green Book forecasts of inflation and output from November 1965 to May 1998, but instead of giving all the forecasts available they present the forecasts closest to the middle of each quarter so as to make the series comparable to the Survey of Professional Forecasters (SPF) and other surveys.⁷ This is convenient because FOMC's meetings have not always been as regular as they are today.⁸ The analysis presented throughout this paper uses the data at the quarterly frequency.⁹

The Green Book contains forecasts for more than 50 variables. This paper uses inflation forecasts for the output deflator.¹⁰ The forecast horizon varies from just the current quarter to as many as nine quarters ahead. In this paper forecasts for the current quarter are labeled forecasts at horizon zero, and only forecasts up to four-quarters-ahead are used because longer horizons do not contain enough data to confidently perform econometric analysis.

In any forecasting exercise the value used as the actual value for the variable of interest, inflation in this paper, can be taken either as the first value released (if available), which is typically referred to as real-time data, or the latest revision of the data.¹¹ In general, it is not clear which value the producer of a forecast is actually targeting, and arguments can be made for either real-time or revised data. For example, one can argue that the Federal Reserve is interested in forecasting the "true" value of inflation, so that evaluation of Fed's forecasts should be done with fully revised data. On the contrary, for commercial forecasters one can argue that they are interested in the accuracy of the forecasts as seen when the data are first released, so that evaluation of commercial forecasts should be done with real-time data. In this paper all the results are reported for both data sets, using the second revision as real-time data and the latest available revision as fully revised data.¹² Sims (2002) points out that it is worth to compare the results with both sets to see if the analysis is sensitive to which variable is used to construct actual values.¹³ There is more to say about real-time versus revised data, but the discussion is postponed until section II where it can be framed in the context of the theoretical model.

⁶Sims (2002) analyses both, Green Book forecasts and forecasts that are directly obtained from the econometric models.

⁷More information about Green Book forecasts and the Survey of Professional Forecasters at the Philadelphia Fed web page: <http://www.phil.frb.org/econ/forecast/index.html>

⁸The committee currently meets every six weeks.

⁹Romer and Romer (2000) use Green Book forecasts at a monthly frequency, whereas Sims (2002) uses data at the quarterly frequency. The advantage of using quarterly data, Sims points out, is that if one uses forecasts from other sources, like the SPF, then the data sets have uniform timing, something that simplifies the econometric analysis.

¹⁰The series been forecasted is quarter-to-quarter (annualized) inflation from the level of the GDP's price index. From 1965 to 1991, the index used was the price deflator implicit in the GNP, from 1992 to the third quarter of 1996 it was the price deflator implicit in the GDP, and since then it has been the GDP's chain-weighted price index. All the series are seasonally adjusted.

¹¹See Croushore and Stark (2002) (with discussion).

¹²Real-time data is taken from the Federal Reserve Bank of Philadelphia's web page: <http://www.phil.frb.org/econ/>. For more on real time data see Croushore and Stark (2000). Revised data is also taken from the real-time data base. It corresponds to the last vintage available in May 2004.

¹³Romer and Romer (2000) use the second revision, whereas Sims (2002) uses fully revised data.

2.1 Comments on the Tests Used in the Literature

Romer and Romer (2000) conclude that inflation forecasts from the Federal Reserve are rational and that they dominate commercial forecasts. They use Green Book forecasts of inflation in a sample that goes from November of 1965 to November of 1991.¹⁴ For the commercial forecasts they use forecasts taken from Data Resources Inc., Blue Chip Economic Indicators, and the SPF. For the last two they use the consensus forecast formed by taking the median across forecasters.¹⁵

The Romers reach their conclusion about rationality by estimating, for each forecaster and forecast horizon, what in the forecasting literature is known as a Mincer-Zarnowitz regression (Mincer and Zarnowitz (1969)). Let π_{t+h} denote inflation h periods after period t . For example, if t equals the first quarter of 1990 and h equals two, then π_{t+h} is actual inflation in the third quarter of 1990. In the same way, let $f_{t+h,t}$ denote the forecast of inflation made at period t for period $t+h$. Then the Mincer-Zarnowitz regression is:

$$\pi_{t+h} = \alpha + \beta f_{t+h,t} + \varepsilon_{t+h}, \quad (1)$$

and a test of rationality is that $\alpha = 0$ and $\beta = 1$.¹⁶ The Romers apply ordinary least squares (OLS) to their sample and find that inflation forecasts from commercial forecasters and from the Green Book are rational. They conclude that all the forecasters are using their information efficiently.

To fully understand what the Mincer-Zarnowitz regression tests, one can think of imposing $\beta = 1$ and then on subtracting the forecast from both sides of the regression. If the forecast error is defined as $e_{t+h,t} \equiv \pi_{t+h} - f_{t+h,t}$ the transformed regression is:

$$e_{t+h,t} = \alpha + \varepsilon_{t+h}. \quad (2)$$

Testing that $\alpha = 0$ in the last regression is equivalent to jointly testing that $\alpha = 0$ and $\beta = 1$ in the Mincer-Zarnowitz regression. If β is different from one (and, for the sake of the argument, $\alpha = 0$), a traditional t-test on α would still reject the hypothesis of rationality in equation (2) as it is testing the whole maintained hypothesis, from which the restriction $\beta = 1$ is part of. In the second regression is clear that what is being tested is if the forecast errors have a zero mean, that is, if there is no systematic bias in the forecasts. The idea is that rational forecasts should not systematically over- or under-predict because simply adding the constant α to the forecasts improves them.

Upon closer inspection, Green Book forecasts before 1991 (Romers' sample) appear unbiased, but not in the random way that rationality calls for. A simple inspection of the time series of the forecast errors in Figure 1 reveals systematic positive errors (under-prediction) up until about 1979, and systematic negative errors (over-prediction) from about 1979 to about

¹⁴The Romers' sample ends in 1991 because of the lag in the release of Green Book forecasts, and to avoid the change to GNP.

¹⁵For details on the samples used for commercial forecasts and more information see the original Romer and Romer (2000) paper.

¹⁶Under the null of rationality and quadratic loss, $E_t[\varepsilon_{t+h}] = 0$. The properties of the error are discussed later in the paper.

1991.¹⁷ The specific dates change with the horizon used, but it is clear that the average of the forecast errors is close to zero because for the first part of the sample the average is positive whereas for the second part the average is negative, offsetting each other when the average is taken using the entire sample up until 1991. When Sims extended the sample to 1995, he reports finding some evidence that the Green Book inflation forecasts are (negatively) biased. Figure 1 shows that Sims's result differs from the Romers' because the tendency to over-predict inflation was maintained during the first half of the nineties.

Some objections have emerged over the years about the use of the Mincer-Zarnowitz regression to test rationality. Granger and Newbold (1986) correctly indicate that the regression is only testing a necessary condition for the optimality of the forecasts.¹⁸ It is easy to see Granger and Newbold's point that a forecast could be unbiased without been optimal. For instance, an unbiased forecast for inflation could be based just on past inflation but a better forecast, also unbiased, may be available by also using information on unemployment or output. Without further tests that make use of the forecaster's information set when testing rationality, it is certainly premature to conclude that a forecast is using all the available information in an efficient way just because it passes an unbiasedness test. In the forecasting literature optimality of a forecast is always defined with respect to the variable considered to be in the forecaster's information set (Clements and Hendry (1998)). If a constant is used in the definition, then the forecast is said to be unbiased (or weakly rational). If another variable is used, then the forecast is said to be efficient (or optimal) with respect to that variable.

Another objection about the Mincer-Zarnowitz regression is that in order for the original test to work one has to assume that the error in the regression is white noise, which is only true for optimal one-step-ahead forecasts (more on this below). The Romers deal with this issue by using autocorrelation-corrected standard errors in their estimates. But optimal forecast errors have a precise form of serial correlation, and one can directly model it. Also one may want to test for the presence of serial correlation in the forecast errors as another test of rationality.

Finally, there are some advantages of using equation (2) instead of equation (1) when testing for unbiasedness. First, only one parameter has to be estimated, which is certainly important in macroeconomics where the number of observations is small. Second, equation (1) requires the forecast to be uncorrelated with the error term for the estimators of α and β to be consistent, which is true for optimal forecasts (again, more on this below) but not for other forecasts, whereas equation (2) does not have this requirement. Third, if the variable to be forecasted is highly persistent, like inflation, then both the dependent and the explanatory variables are highly persistent in regression (1) which may cause the traditional test to over-reject the null hypothesis, as the normal distribution may be a poor approximation to the distribution of the test.¹⁹ Regression (2) does not present this problem because the dependent variable is not persistent and the explanatory variable is just a constant.

The Romers also show that Green Book forecasts dominate commercial forecasts of in-

¹⁷Orphanides (2002) reports that the Green Book forecasts are clearly biased towards under-prediction for the period 1969-1979, but he does not quantify the bias nor does he tests it.

¹⁸Which is to say that it is testing only a necessary condition for rationality because under rational expectations, expectations are optimal forecasts (Muth 1961).

¹⁹See Cavanagh, Elliott and Stock (1995).

flation. They show this by running forecast combination regressions pairwise with the Green Book forecasts in each regression. The regression is:

$$\pi_{t+h} = \alpha + \omega^F f_{t+h,t}^F + \omega^C f_{t+h,t}^C + \varepsilon_{t+h}, \quad (3)$$

where α is a constant, ω^F is the weight assigned to the Federal Reserve forecast (denoted by $f_{t+h,t}^F$) and ω^C is the weight assigned to the commercial forecast (denoted by $f_{t+h,t}^C$). The Romers apply OLS to their sample and find that the constant and ω^C for each commercial forecaster are in general not significantly different from zero, whereas ω^F is in general not significantly different from one. According to these results, if one had access to both forecasts the best combination would be to assign a weight of one to the Federal Reserve forecasts and a weight of zero to the commercial forecasts, that is, the optimal action would be to throw away the commercial forecasts.²⁰ Sims (2002) reaches the same conclusion using a similar methodology.

Combination regressions like (3) first appeared in Granger and Ramanathan (1984), and were later used by Hendry and Chong (1986) to test what they called “forecast encompassing”. According to Hendry and Chong, a forecast “encompasses” another forecast if the weight assigned to the first forecast is not significantly different from one and the weight assigned to the second forecast is not significantly different from zero. The idea behind forecast encompassing is to test if one forecast contains information useful for another forecast of interest or not, for example, Fair and Shiller (1989) use a regression like (3) to measure information content of the forecasts.

If the objective is to combine forecasts, it is clear that equation (3) is an adequate way to proceed, and that as a by product one can obtain an encompassing test by testing if the weight assigned to the “encompassed” forecast is zero. But if the objective is to test if one forecast has information not contained in another forecast, then one can directly test for forecast encompassing. For the case of the Federal Reserve and commercial forecasters, the following regression can be used:

$$e_{t+h,t}^F = \alpha + \omega^C (e_{t+h,t}^F - e_{t+h,t}^C) + \varepsilon'_{t+h}, \quad (4)$$

where $e_{t+h,t}^i \equiv \pi_{t+h} - f_{t+h,t}^i$, $i = F, C$. The error term in regression (4) is different from that in regression (3) because the restriction $\omega^F + \omega^C = 1$ is imposed to obtain (4). An encompassing test is simply the test of $\omega^C = 0$. The reason regression (4) is a better way to test forecast encompassing relative to regression (3) is again the fact that one gains one degree of freedom, although at the expense of imposing the restriction that the weights add to one. An alternative is to use the forecast $f_{t+h,t}^C$ as the explanatory variable in regression (4), but if the variable of interest is persistent, like inflation, then the normal distribution may not be a good approximation to the distribution of the test statistic of interest. A by-product of regression (4) is that the coefficient ω^C is the weight the commercial forecaster would receive in regression (3), with $\omega^F = (1 - \omega^C)$.

²⁰Romer and Romer (2000) conclude that if both, the Fed and commercial forecasters are using all their information efficiently (because they are rational) and if Federal Reserve forecasts encompass commercial ones, then it must be that the Fed has more information. They get the same results when they analyze output forecasts.

Comparing regressions (2) and (4) one can see that the constant in regression (4) can be used to test the unbiasedness of the Fed's forecasts. One can also see that the role of the constant in regression (3) is to compensate for any bias contained in the forecasts to be combined so that the resulting combination is by construction unbiased.

2.2 Empirical Evidence

An important practical issue in evaluating Federal Reserve forecasts is that their mean seems to be different from zero and to change over time (Figure 1). To accommodate such changes, in this section allowance is made for changes in the value of the parameters of a regression that tests, under quadratic loss, unbiasedness and serial correlation of forecast errors and information content of forecast in a parsimonious way. The result is that the hypotheses of rationality and encompassing are rejected under quadratic loss.

2.2.1 A Regression to Test Rationality, Serial Correlation, and Information Content Under Quadratic Loss

Apart from testing for unbiasedness and to see if the Green Book forecasts encompass commercial forecasts there is another property of rational forecasts that is worth looking at. Under quadratic loss optimal forecast errors should have an autocorrelation structure like that of a moving average (MA) of order $(h - 1)$, where h denotes the forecasts horizon. A formal derivation can be found in Granger and Newbold (1986, 130), but the intuition is easy to convey: A forecast for $t + 2$ made at t (a two-step-ahead forecast) has to include information up to t , but any shock occurring in the two periods between t and $t + 2$ is not taken into account. At $t + 1$, another two-step-ahead forecast is going to be issued, and is going to be a forecast for $t + 3$. The second forecast contains information up to $t + 1$, but does not contain information about anything that occurs in the two periods between $t + 1$ and $t + 3$. So there is one period, from $t + 1$ to $t + 2$, for which neither forecast has information. Any event that happens in this period is going to impact both forecasts, inducing an $MA(1)$ -like behavior in the forecast errors.

How can this property be tested? If the forecast errors behave like an $MA(h - 1)$ then any autocorrelation of order h or larger has to be zero.²¹ So this property can be tested using the regression:

$$e_{t+h,t} = \gamma e_{t+h-j,t-j} + \varepsilon_{t+h}, \quad (5)$$

with $j \geq h$. The hypothesis of no serial correlation corresponds to $\gamma = 0$. The dependent variable is the same as in equations (2) and (4), which suggests that a single regression can be used to tests for unbiasedness, serial correlation, and encompassing.

In this section a single multivariate regression is used to analyze Green Book inflation forecasts. The SPF consensus forecast of inflation is used as representative of commercial

²¹See Hamilton (1994).

forecasts.²² The regression is:

$$e_{t+h,t}^F = \alpha + \gamma e_{t+h-j,t-j}^F + \omega^C (e_{t+h,t}^F - e_{t+h,t}^C) + \varepsilon_{t+h}. \quad (6)$$

OLS is applied to the available sample (1968:4 to 1998:4) for each horizon, with $j = h+1$.^{23,24} To correct for any autocorrelation in excess of $j = h+1$ and for heteroskedasticity, expected from a non-constant variance in Figures 1 and 2, autocorrelation and heteroskedasticity corrected standard errors using Newey and West's (1997) method are employed.²⁵ The results are presented in Tables 1 (real-time data) and 2 (revised data).

In terms of the bias, the sign of the estimated α is negative for all horizons and data sets, although it is only significantly different from zero for $h = 3$. A look at the forecast errors is helpful to explain the result. Figure 1 presents forecast errors for horizons one and four. When the sample employed by the Romers is extended to include most of the nineties, the systematic over-prediction of inflation (forecast errors systematically below zero) that occurred in the last part of the sample outweighs the systematic under-prediction (forecast errors systematically above zero) that occurred during the first part of the sample and instead of an average error close to zero one gets a negative average. The results from the regressions and a look at the plot of the forecast errors provide evidence to reject unbiasedness of Federal Reserve inflation forecasts.

But Figure 1 contains more information. It indicates the periods of each of the Chairmen of the Federal Reserve during the sample. One can see that the bias presents a pattern that can be associated with the Chairmen. From the beginning of the sample until about 1979, the Fed systematically under-predicts inflation. From about 1979 onwards the Fed systematically over-predicts inflation. But this coincides with Volcker's appointment as Chairman. So that Chairmen considered to have strong preferences against inflation, Volcker and Greenspan, presided over periods with negative bias, whereas Chairmen considered to be more relaxed about inflation (Chairmen before Volcker) presided over periods with positive bias.²⁶ This

²²The Survey of Professional Forecasters is conducted by the Federal Reserve Bank of Philadelphia. It was formerly known as the ASA/NBER Economic Outlook Survey. The consensus used in this paper is formed by taking the median across forecasters. The variable been forecasted is the GNP deflator prior to 1992, the GDP implicit price deflator prior to 1996 and the GDP price index since then. The forecasts of inflation are calculated as: $f_{t+h,t} = 400 * \ln \left(\frac{P_{t+h}}{P_{t+h-1}} \right)$. For more information see Croushore (1993) or go to: <http://www.phil.frb.org/econ/spf/index.html>

²³The forecast for the current quarter is typically labeled forecast at horizon zero, $h = 0$, a convention that is followed in this paper. But from a theoretical point of view, this forecast should behave like an $MA(0)$ because it is the first forecast. Accordingly, the forecast labeled $h = 1$ should behave, if optimal, as an $MA(1)$, not as an $MA(0)$, because it is theoretically the second-step-ahead forecast.

²⁴The sample starts in 1968 rather than in 1965, as the Romers sample, because the quarterly data put together by the Philadelphia Fed starts the same date than the SPF data.

²⁵The bandwidth was chosen so that h lags were included to calculate the variance covariance matrix because under the null of forecast optimality the errors from the regressions should have autocorrelations up to order h different from zero. Newey-West method was chosen to avoid ending with a non positive definite matrix.

²⁶W.M. Martin Jr was the Chairman of the Federal Reserve until the first quarter of 1970, between February 1970 and January 1978 A. Burns was the Chairman, and G.W. Miller was the Chairman from March 1978 to August 1979. P. Volcker's period covered August 1979 to August 1987. Finally, A. Greenspan has been in charge since August 1987. On Chairmen's preferences about inflation see Romer and Romer

pattern will be exploited later in the paper.

With respect to serial correlation in the forecast errors, something not tested in other papers, the results with real-time data (Table 1) indicate that for horizons zero, one, and two there is evidence of serial correlation. When fully revised data is used (Table 2) the evidence is stronger, as all horizons but one show evidence of serial correlation. If the possibility of asymmetry in the Fed's loss function is ignored, these results point to the Fed's inefficient use of the information contained in its own past forecast errors.

Finally, the estimates of the coefficients associated with the encompassing tests show some evidence that the Federal Reserve inflation forecasts do not encompass those of the SPF consensus. With real-time data (Table 1) horizons one and four have estimates that are significantly different from zero, which is enough to reject the null of encompassing. The estimated coefficient for horizon zero indicates that the optimal combination assigns a weight of 0.21 to the SPF's forecasts and a weight of 0.79 to the Fed's. But the estimated coefficient for horizon four is negative which is difficult to explain. The fact that the estimate is different from zero implies that it contains information that the Fed can use. The fact that the estimated coefficient is negative indicates that the weight assigned to the Fed forecast is more than one, but that the SPF consensus is still worth looking at by the Fed because it explains a part of the forecast errors not explained by the Federal Reserve forecast. When revised data is used (Table 2) horizon zero has a significant coefficient of 0.35, which means that the optimal combination is to assign a weight of 0.65 to the Fed's forecasts and a weight of 0.35 to the SPF consensus. The overall picture is that with the full sample the SPF consensus seems to contain some information that the Federal Reserve does not have, in particular in the very short run.

Joint tests of rationality are also performed. These are Wald tests that test if all the coefficients are equal to zero. The tests reject the null at 10% for all horizons. The overall conclusion is that when the sample is extended to 1998 and asymmetries are not allowed Federal Reserve inflation forecasts seem irrational. But the rejections of the joint tests are mainly driven by the presence of serial correlation. This could be indicative of structural breaks not considered rather than of Fed's irrationality.

2.2.2 Structural Breaks: 1974-1975 and 1979-1980

To investigate the possibility of changes in the parameters of equation (6) the sample is split at each possible breakdate and the parameters of the model are estimated separately for each subsample.²⁷ Bai (1997) indicates that the OLS estimate of the break date is the date that minimizes the residual variance (sum of squared errors divided by sample size) as a function of the breakdate. Figure 3 plots the residual variance for horizons one to four using revised data.²⁸ Although it is not a formal test, the visual analysis is informative regarding the potential breakdates. The plots in Figure 3 show two well-defined minima. A global minimum for horizons one and two occurs around 1974 – 1975. A global minimum

(2004).

²⁷The sample is trimmed so that there are enough data points to estimate the first and last regressions.

²⁸Horizon zero is not used because it is not very informative as it has a dip from 1973 to 1981 with no clear minimum. The plots using real-time data are less informative (in the same sense as before) so are not presented. The formal tests presented below use both real-time and revised data.

for horizons three and four occurs around 1978 – 1980. This last period also coincides with local minima for horizons one and two. This evidence suggests that two structural breaks are present in the full sample.²⁹

To formally test for the presence of multiple structural breaks the procedure suggested by Bai and Perron (1998, 2003) is implemented. However, instead of using equation (6) the procedure is applied to regressions that have only a constant as a regressor. This is because convergence results are not available when there is a lagged dependent variable and serial correlation in the errors (Bai and Perron (2003)). The approach followed here is to test for multiple breaks in the mean with tests that permit serial correlation and heteroskedasticity in the errors. Allowance is made for up to three breaks and the trimming is fifteen percent of the sample. Different variances of the residuals across segments is also allowed. The results are presented in Tables 3 (real-time data) and 4 (revised data).³⁰ The conclusion from Table 3 is that there is a break in the first quarter of 1975 for $h = 0$, and a break in the third quarter of 1979 for $h = 4$. The conclusion from Table 4 is that there is a break around 1974 - 1975 for horizons zero, one, and two, and a break around 1979 - 1980 for horizons two, three, and four.³¹ The overall conclusion about parameter constancy for equation (6) is that there is evidence of two structural breaks, one around the beginning of 1975 and a second around the end of 1979.

From a statistical perspective, the first break marks the end of the era of “big forecast errors” (as high as five and a half percentage points; see Figures 1 and 2), so the tests are identifying a switch from large (positive) forecast errors to more moderate (positive) forecast errors. The second break identifies the change from systematically positive to systematically negative errors.

From an economic perspective both breaks coincide with negative supply shocks (Mishkin

²⁹The approach reported in this paper is to treat the breakdate as unknown, although the approach of taking the breakdate as known was also investigated. A Chow (1960) test was applied to the regression for each horizon and for each data set (real-time and revised). Results (not reported) indicate that if the breakdate is set at 1979:3, the time P. Volcker took office, there is strong evidence in favor of (the alternative hypothesis of) a break at that time. But the Chow test has two problems: (1) the breakdate is correlated with the data as it was picked by looking at the plot of the time series of the forecast errors, which can result in the test falsely indicating the presence of a break, and (2) it assumes that the variance for the two subsamples is the same, an assumption that is not a good description of the behavior of the forecast errors (figures 1 and 2).

³⁰ $UDmax$ is a test of the null hypothesis of no structural break against an unknown number of breaks given the upper bound of three breaks. The test maximizes an equal weighted version of the tests for each possible number of breaks. Bai and Perron (1998) indicate how to calculate the critical values. $SupF(l + 1|l)$ is a test for l versus $l + 1$ breaks. The test consists on applying $l + 1$ tests of the null of no structural change versus the alternative hypothesis of a single change. The way it is implemented is by taking the minimal value of the sum of squared residuals over all segments where an additional break is included ($l + 1$ breaks). If this minimal value is smaller than the one with l breaks then the test concludes with a rejection in favor of a model with $l + 1$ breaks (see Bai and Perron 1998). BIC (Bayesian Information Criterion) and Sequential refer to procedures to choose the number of breaks. BIC estimates the models with different number of breaks and selects the best model using the BIC criterion. Sequential is based on the sequential application of the $supF(l + 1|l)$ test (see Bai and Perron 2003). Finally, T_1 and T_2 are the estimated breakdates based on Bai and Perron’s procedure to find the global minimizer of the sum of squared residuals when two breaks are allowed (three subsamples).

³¹There is evidence of two breaks using horizons two and four. The second break using horizon four is estimated at the four quarter of 1985.

2001): In the 1973-1975 period the economy was hit by the first oil shock, a sharp increase in food prices due to a series of crop failures, and the termination of price controls, and during the 1979-1980 period the economy was hit again by crop failures and the second oil shock. But the second break also coincides with the appointment of Volcker as Chairman of the Federal Reserve. In the monetary policy literature the appointment of Volcker is considered as a change in the Fed's views towards inflation with less emphasis on controlling inflation in the pre-Volcker era than in the period since Volcker. For example, Romer and Romer (2004) review the narrative record of the Federal Reserve and find that key determinants of the monetary policy in the United States have been Chairmen's "... views about how the economy works and what monetary policy can accomplish." (Romer and Romer 2004, p. 130). Reviewing the Chairmen's views they also find that:

Well-tempered monetary policies of ... the 1980s and 1990s stemmed from the conviction that inflation has high costs and few benefits, ... In contrast, the profligate policies of the late 1960s and 1970s stemmed ... from a belief in a permanent trade-off between inflation and unemployment... (Romer and Romer 2004, p. 130).

Clarida, Gali and Gertler (2000) also support the idea that there is a significant difference in the way monetary policy was conducted pre- and post-Volcker. They find that the Fed let real interest rates decline as expected inflation rose before Volcker whereas it systematically raised real rates in response to higher expected inflation in the post-Volcker era. So, only the second break coincides with what is believed to be an endogenous change in preferences within the Federal Reserve, the producer and user of the forecasts.

2.2.3 Bias and Encompassing Before and After Volcker

To allow for the structural breaks, estimates of equation (6) are presented for three subsamples. The first covers from the beginning of the sample to the end of 1974. The second from the beginning of 1975 to the third quarter of 1979 and the third from the fourth quarter of 1979 to the end of the sample (the second quarter of 1998). The results are presented in Table 5 (real-time data) and 6 (revised data).

The samples pre-1979 have forecast errors with a significant positive mean for most horizons. All the coefficients but two, corresponding to horizons zero and one for the 1975-1979 period, are significantly different from zero when revised data is used. Real-time data shows only a few significant coefficients, but all of them are positive. The difference between the pre- and post-1975 period is a reduction in the magnitude of the bias for each horizon, but the qualitative results are the same across these two periods. This result contrast with the difference pre- and post-1979. The bias is significant post-1979 for all horizons and data sets, but the sign is negative. The negative sign corresponds to the Federal Reserve's systematic over-prediction of inflation. For example, a bias of -0.5 would correspond to the Federal Reserve systematically over-predicting inflation, on average, by half a percent: if the forecast for three-quarters-ahead is 3.0%, then the average realization of inflation is likely to be 2.5%, not 3.0%. When fully revised data is used the results are qualitatively the same,

but the magnitude of the bias is larger, with a bias as large as three quarters of a percent.³² So there is a systematic tendency to under-predict inflation before 1979 and a systematic tendency to over-predict it after 1979.

There is almost no evidence of serial correlation within samples. In fact, when real-time data is used only one coefficient is significantly different from zero. This indicates that the serial correlation found using the full sample is a reflection of not taking into account the structural breaks.

The results from the encompassing tests are very interesting. The dominance of the Fed is undermined with respect to the results obtained by the Romers. When data post-Volcker is considered, both real-time and revised data show that the SPF consensus has valuable information (from the Fed's point of view under quadratic loss) for the first two horizons. For the forecasts corresponding to horizon zero the estimated weights indicate that the optimal combination is to average the forecasts. This is a common result in the forecasting literature, but a new result with these data. Results pre-Volcker show that Fed forecasts encompass the SPF's, except for horizon zero when revised data is used.³³ That the SPF forecasts contain more information when the post-Volcker sample is used indicates learning over time by commercial forecasters. This result, although interesting, is not pursued further in this paper.

The results about encompassing using the post-Volcker sample show another very interesting aspect of the informational advantage of the Fed over the SPF. The weight associated with the SPF consensus is decreasing with the forecast horizon. Only the estimates for horizons zero and one are statistically significant, but the economic significance of the tendency is very important, as it points to the fact that the informational advantage of the Federal Reserve increases with the forecast horizon. Sims (2002) suggests that the main advantage of the Federal Reserve over commercial forecasters may be a better knowledge within the Fed of the timing of changes in the policy stance. The results presented here support Sims's suggestion, as one would expect knowledge about the monetary policy stance to be more important for longer horizons.

The results that Federal Reserve forecasts encompass the SPF consensus for some horizons but not for others can be interpreted as saying that commercial forecasters have a wider information set than the Fed's, at least for some horizons. But if Fed's information set is equal or wider than that of commercial forecasters, something plausible due to the resources devoted by the Fed to the task, and if the Fed's loss function is quadratic, then the result can also be interpreted as supporting the hypothesis that the Fed uses information inefficiently.

Finally, the joint tests clearly indicate rejection of the null hypothesis (rationality and quadratic loss) in each subsample, except for the period between 1975 to 1979 where the tests cannot reject for some horizons.³⁴ The fact that rationality is rejected (at least for some horizons) for the sample before 1979 provides evidence that the results presented in this paper differ from those of the Romers' even when the analysis is done using a subsample

³²It is interesting to notice that a comparison of the results using real-time versus revised data shows that there may be a small bias in the real time data that is corrected in the revisions. If this bias is indeed present, the Federal Reserve could immediately improve its forecasts by taking this information into account (i.e., taking into account both time series, the real-time series and the revised one).

³³I have no explanation for the significant coefficient at horizon two.

³⁴This last result could be due to the small number of observations in that period.

of their data.³⁵

3 Reconciling Evidence with Forecasts Under Asymmetric Loss

The results presented so far are tied to the assumption that the Federal Reserve has a symmetric forecasting loss function. That is, there is an implicit assumption that if the Fed's inflation forecast for four-quarters-ahead is 3%, the following two alternative events have the same costs for the Federal Reserve: That actual inflation turns out to be 4%, in which case the forecast error is negative, -1%, or that actual inflation turns out to be 2%, in which case the forecast error is positive, 1%. In both events, the magnitude of the error is the same, but the signs are different. Is it sensible to assume that for the Federal Reserve both events have the same costs?

Recent monetary policy literature suggests a negative answer to that question, indicating that it is likely that central banks have asymmetric preferences about inflation. Nobay and Peel (2003) employ an asymmetric loss function, the "linex" loss (described below), to model central bank preferences. The linex loss nests as a special case the quadratic loss, but in general allows for different marginal losses for errors of equal magnitudes but different signs. Ruge-Murcia (2003) also employs the linex loss function to model central bank's preferences. Using implications from his theoretical model he finds empirical evidence to support an asymmetric loss function for inflation using data on 21 OECD countries. He finds evidence of asymmetric costs for Canada, Sweden and the United Kingdom. For the rest of the countries, including the United States, he is not able to reject symmetric preferences.³⁶

The papers by Nobay and Peel (2003) and Ruge-Murcia (2000, 2003) use asymmetric costs to model the fact that for a prudent central bank inflation above the target is more costly than inflation below the target. When the target is explicit, control errors (inflation minus the target) can be used to test for asymmetric preferences (Ruge-Murcia (2003)). The problem for central banks with implicit inflation targets (like the Fed) is that control errors cannot be used to test for asymmetries. This paper suggests that in this case inflation forecast errors may be used, as a central bank pursuing an inflation target would set its optimal monetary policy so as to have the forecasts equal to the target (Svensson (1997)).

In this section the loss function implied by Fed's inflation forecasts is backed out using moment conditions derived from a model of an inflation targeting central bank. Evidence is found of asymmetric costs of under- and over-prediction, with a larger cost of under-prediction than of over-prediction for the post-Volcker era, and with the opposite result for the pre-Volcker period. Once asymmetric costs are taken into account the Fed is found to be rational and to efficiently incorporate the information contained in SPF inflation forecasts.

³⁵Rationality (joint test) is also rejected when a sample from the fourth quarter of 1979 to the fourth quarter of 1991 (the end of the Romers' sample) is used. The rejections are mainly driven by the (negative) bias.

³⁶However, Ruge-Murcia needs to impose that inflation follows a Gaussian distribution, and his empirical results may simply reflect the failure of the data to meet this assumption. He also uses linear approximations to his nonlinear theoretical model, which may further undermine the empirical results.

3.1 A Model of a Central Bank with an Asymmetric Loss

The economy is represented by a single equation for inflation:

$$\pi_{t+l} = as_t - bi_t + \varepsilon_{t+l}, \quad (7)$$

where $\pi_{t+l} \in \mathbb{R}$, $s_t \in \mathbb{R}^k$, $i_t \in \mathcal{I}$, \mathcal{I} a compact subset of \mathbb{R} , a and b are coefficients, and ε_{t+l} is a $N(0, \sigma_{t+l}^2)$ random variable independent to both i_t and s_t .³⁷ π_{t+l} is inflation l periods ahead, where l is the time between the policy action and its effect on the economy (the control lag). s_t are exogenous variables that affect inflation at t and that the central bank has no control of. The realization of π_{t+l} is not observable at t , whereas that of s_t is. i_t is the monetary policy instrument used by the central bank, for example a short-run interest rate. It is assumed that s_t and i_t are independent. ε_{t+l} represents variables that affect inflation and are not known at t . Equation (7) represents the effect of i_t on the distribution of inflation. For simplicity, the effect is assumed to be just a shift in the mean of inflation, and invariant to different choices of i_t .³⁸

Equation (7) can be seen as the reduced form for inflation derived from a new-Keynesian model that includes a Phillips curve and an aggregate demand (e.g., Svensson (1997); Woodford (2003)). Different assumptions about supply and demand would lead to different restrictions on the coefficients a and b , but this restrictions do not play a role for the empirical part of this paper. The reason is that the empirical analysis takes the forecasts of inflation as given, regardless of the model used by the central bank to produce them, and in that sense the model presented in this section is meant only to illustrate how to derive the moment conditions used in the empirical part. It will also be useful to illustrate the implications of different loss functions on equilibrium inflation once the model is solved.

The preferences of the central bank over the possible realizations of π_{t+l} are described by a loss function that indicates the costs associated with a particular realization of π_{t+l} and the central bank's inflation target, $\pi_{t+l,t}^T$, through the control error $ce_{t+l} = \pi_{t+l} - \pi_{t+l,t}^T$. The target is defined at t for $t+l$. The loss function will be denoted $L(\pi_{t+l} - \pi_{t+l,t}^T, \phi)$ where ϕ is a fixed parameter.³⁹ The loss functions used in this paper are the quadratic loss: $L(ce_{t+l}, \phi) = (ce_{t+l})^2$, which is differentiable everywhere and symmetric; linex loss: $L(ce_{t+l}, \phi) = \frac{1}{\phi^2} [\exp[\phi(ce_{t+l})] - \phi(ce_{t+l}) - 1]$ with $\phi \in \mathbb{R}$, which is differentiable everywhere and asymmetric (if $\phi > 0$ there are large losses from positive control errors); and asymmetric quadratic loss: $L(ce_{t+l}, \phi) = [\phi + (1 - 2\phi) 1_{(ce_{t+l} < 0)}] |ce_{t+l}|^2$ with $0 < \phi < 1$, which is differentiable almost everywhere and asymmetric (if $\phi > 0.5$ there are large losses from positive control errors).⁴⁰ The loss function indirectly depends on the central bank's

³⁷Gaussianity is assumed in the model for analytical tractability but, as will be made clear later in the paper, the empirical results are derived from moment conditions for which the assumption of Gaussianity is not necessary.

³⁸Reduced form equations that are very similar to the one used in this paper appear in Bernanke and Woodford (1997) and Deutsch and Granger (1992).

³⁹The class of loss functions that depends only on the error have the following basic properties (Granger 1999): (1) $L(0, \phi) = 0$; (2) $\min_{ce_{t+l}} L(ce_{t+l}, \phi) = 0$, so $L(ce_{t+l}, \phi) \geq 0$; and (3) $L(ce_{t+l}, \phi)$ is monotonic non-decreasing in $|ce_{t+l}|$. Additionally, the loss function may be symmetric ($L(-ce_{t+l}, \phi) = L(ce_{t+l}, \phi)$).

⁴⁰Granger (1999), Elliott and Timmermann (2004) and references therein discuss these and other loss functions in a forecasting context. For their use in a monetary policy context see Walsh (2003) for the

actions through the effect of i_t on π_{t+l} .

Two aspects of the loss function are worth highlighting. First, the inflation target is assumed to be time-varying. This is not common in the inflation targeting literature, but it appears to be a good approximation to describe the Federal Reserve’s behavior as argued by Gürkaynak, Sack, and Swanson (2005).⁴¹ In this paper, this assumption plays a critical role as will be apparent later on. Second, the loss function only has the inflation’s control error as an argument, whereas in the literature the loss typically depends on the divergence of inflation from a target, the divergence of output from its natural rate, and sometimes also on the interest rate (e.g., Woodford (2003)). This is also a critical assumption. The loss function used here is meant as a reduced form of a more involved loss function (perhaps derived from a model where the central bank’s loss function maximizes the welfare of the people in the economy) and the conclusions from the empirical part will be interpreted accordingly. The generalization to multivariate loss functions is left for future research.

In this environment the central bank chooses a policy action by minimizing expected loss conditional on all the information available at the time of the decision.⁴² Denote this information set by Ω_t for the decision taken at t . This set contains at least the current and past realizations of s_t , π_t , and i_t , as well as all the past and present inflation targets.

The optimal monetary policy action, i_t^* , solves:

$$\begin{aligned} \min_{i_t \in \mathcal{I}} E [L(\pi_{t+l} - \pi_{t+l,t}^T, \phi) \mid \Omega_t] \\ \text{s.t. } (7). \end{aligned} \tag{8}$$

The optimal action will be a function of the contents of the information set, the target, and the loss function.

Under quadratic loss, $L(ce_{t+l}, \phi) = (ce_{t+l})^2$, the optimal action is the one that satisfies the first order condition (FOC):

$$E[\pi_{t+l} \mid \Omega_t] \Big|_{i_t^*} = \pi_{t+l,t}^T. \tag{9}$$

For a central bank with a quadratic loss the mean summarizes the relevant information contained in the conditional density of inflation and a point forecast of the conditional mean is necessary and sufficient to solve the optimization problem. The conditional mean is a function of Ω_t , which contains i_t , so the optimal policy under quadratic loss is to set i_t so as to make the forecast equal to the target. Svensson (1997) calls this approach “inflation

quadratic and Ruge-Murcia (2003) for linex. Both linex and asymmetric quadratic nest the quadratic ($\phi \rightarrow 0$, $\phi = 0.5$ respectively).

⁴¹The paper by Gürkaynak, Sack and Swanson (2005) considers a specification in which the Fed’s long-run inflation target displays some dependence on past values of inflation. This permits the long-run level of inflation to vary over time.

⁴²The central bank’s objective, as usually modeled in the literature, is to choose a sequence of monetary policy actions so as to minimize the expected value of an infinite sum of discounted losses (Walsh (2003)). Under some conditions, and without loss of generality, the multi-period problem can be broken into a sequence of period-by-period problems (Svensson 1997). It is easy to verify that the model presented in this paper satisfies these conditions, as it implies a minimization with a convex objective function subject to a linear restriction.

forecast targeting”.⁴³ The assumption of a time-varying inflation target is important because otherwise condition (9) would imply constant optimal forecasts.

Assuming that the central bank knows the equation that determines inflation (the functional form, the coefficients a and b , and the process for ε_{t+l}) the optimal forecast (equal to the conditional mean) is:

$$f_{t+l,t}^* = as_t - bi_t^*. \quad (10)$$

Using this forecast and the FOC the optimal action is:

$$i_t^* = b^{-1} [as_t - \pi_{t+l,t}^T], \quad (11)$$

so that the optimal policy is to look at s_t and offset or reinforce its effect on inflation as needed to achieve the target. To solve for the equilibrium substitute the optimal monetary policy (11) on the equation for inflation (7):

$$\pi_{t+l} = \pi_{t+l,t}^T + \varepsilon_{t+l}. \quad (12)$$

In equilibrium, inflation dynamics are determined by the inflation target and the shocks that occur during the control lag. Because of the perfect information and control of the central bank, inflation will be (on average) on target.⁴⁴

Now suppose that the central bank has an asymmetric loss function. In this section the linex loss, $L(ce_{t+l}, \phi) = \frac{1}{\phi^2} [\exp[\phi(ce_{t+l})] - \phi(ce_{t+l}) - 1]$, will be used because of its analytical tractability.⁴⁵ Under linex loss the central bank’s FOC is:

$$E[\pi_{t+l} | \Omega_t]_{i_t^*} + \frac{\phi}{2} var[\pi_{t+l} | \Omega_t]_{i_t^*} = \pi_{t+l,t}^T. \quad (13)$$

In the model the monetary policy instrument does not affect the variance of inflation, so the optimal policy is to set the interest rate so as to make the expected value of inflation equal

⁴³Different loss functions imply different losses for a given outcome of π_{t+l} so different central banks will need forecasts of different summary statistics. When a forecaster knows the loss function of the decision maker, a forecast of the relevant summary statistic will be necessary and sufficient information for the decision maker. If the forecaster does not know the loss function, or if the forecast is to be used by different decision makers (with distinct losses), then a density forecast will be needed as shown in Diebold, Gunther, and Tay (1998), and Granger (1999).

⁴⁴Notice the identification problem that arises in this model (even with a constant target): If the central bank has been following the optimal policy, the available sample for inflation comes from the equilibrium equation, and for an econometrician the parameters a and b are unidentified in the typical sense that different values for them imply the same probability distribution for the available sample. In this extreme example the linear dependence of i_t on s_t introduces multicollinearity problems that leave the inflation target and the variance of ε_{t+l} as the only parameters that an econometrician can estimate. A consequence of this identification problem is that if in this economy a forecaster tries to estimate a forecasting model using s_t as an indicator for π_{t+l} , he would find that s_t is not an inflation indicator (i.e., a regression of π_{t+l} on s_t would give an estimated coefficient of zero). This point is made in Bernanke and Woodford (1997), although is not referred to as an identification problem.

⁴⁵This section assumes that inflation follows a Gaussian process and a linex loss is used as the asymmetric loss, but for the empirical application the assumption of Gaussianity is not needed and another asymmetric loss function will be used (the asymmetric quadratic). This is because linex is very tractable analytically (under Gaussianity) whereas the asymmetric quadratic is not.

to the inflation target minus a precautionary term (if ϕ is positive, meaning that inflation above the target is more costly than inflation below the target). The precautionary term depends on the degree of asymmetry of the central bank's objective function and on the variance of inflation. In this case the information contained in a measure of location is not enough information for the central bank. It also needs information about the dispersion of inflation. Everything else equal equation (13) implies that the interest rate chosen by a cautious central bank ($\phi > 0$) would be higher than that chosen by a symmetric (or neutral) central bank ($\phi \rightarrow 0$).

Under perfect information the optimal forecast is:

$$f_{t+l,t}^* = as_t - bi_t^* + \frac{\phi}{2}\sigma_{t+l,t}^2, \quad (14)$$

which is biased, with the bias given by $\frac{\phi}{2}\sigma_{t+l,t}^2$, where $\sigma_{t+l,t}^2$ is a forecast of the variance of inflation. The optimal monetary policy is:

$$i_t^* = b^{-1} [as_t - \pi_{t+l,t}^T] + \frac{\phi}{2}b^{-1}\sigma_{t+l,t}^2. \quad (15)$$

which leads to an equilibrium inflation of:

$$\pi_{t+l} = \pi_{t+l,t}^T - \frac{\phi}{2}\sigma_{t+l,t}^2 + \varepsilon_{t+l}. \quad (16)$$

An inflation targeting central bank with an asymmetric loss function will over-predict inflation if inflation above the target is more costly than inflation below it because is going to set the interest rate so that the expected value of inflation is below the target. The difference between the expected value and the target is a precautionary term that depends on the degree of asymmetry and the dispersion of inflation. Because the optimal forecast is equal to the target, there is also a difference between the optimal forecast and the expected value of inflation. This difference is an optimal forecasting bias.⁴⁶

3.2 Estimation of the Asymmetry Parameter

In this subsection the asymmetry parameter of the Federal Reserve is backed out from the orthogonality conditions derived from the central bank's model using the method proposed by Elliott, Komunjer and Timmermann (2005). The results are that under-prediction is four times as costly as over-prediction for the sample post-Volcker and one third of the cost for the pre-Volcker era, thus supporting the presence of asymmetric costs in both periods.

3.2.1 Derivation of Moment Conditions

Forecasts of inflation made by the Federal Reserve can be evaluated using the optimality conditions derived from the model presented above. From the general optimization problem

⁴⁶With the assumptions presented so far (including Gaussianity of the inflation process) the asymmetry parameter is identified and can be estimated using equation (16) as in Batchelor and Peel (1998). This is not the approach followed in this paper.

(8), the optimal monetary policy action satisfies the optimality condition (FOC):

$$E \left[L' (\pi_{t+l} - \pi_{t+l,t}^T, \phi) \mid \Omega_t \right] = 0, \quad (17)$$

where $L' (\pi_{t+l} - \pi_{t+l,t}^T, \phi)$ denotes the derivative of the loss function with respect to the control error.⁴⁷ Following Granger (1999) and Patton and Timmermann (2004), this derivative will be called the generalized error. It gives the change in total loss resulting from a one-unit change in the control error. Condition (17) implies that the optimal generalized error follows a martingale difference sequence with respect to the information set Ω_t (or conditional expectations if $l > 1$).⁴⁸ By orthogonality of martingale differences, for any finite random variable constructed from the contents of Ω_t , $v_t \subset \Omega_t$ (in fact for any finite function of a vector $v_t \subset \Omega_t$), the optimal generalized error satisfies the orthogonality condition:

$$E \left[v_t L' (\pi_{t+l} - \pi_{t+l,t}^T, \phi) \right] = 0. \quad (18)$$

The literature that suggests a decision-theoretic approach to forecast evaluation (Granger and Pesaran (2000); Pesaran and Timmermann (2004)) derives an orthogonality condition similar to (18). The difference is that the loss function in (18) depends on the control error $\pi_{t+l} - \pi_{t+l,t}^T$, whereas the loss function used in the forecasting literature depends on the forecasting error $\pi_{t+l} - f_{t+l,t}$. To link (18) with the condition used in the forecasting literature one can substitute the inflation target in (18) with the optimal forecast. By doing the substitution one gets:

$$E \left[v_t L' (\pi_{t+l} - f_{t+l,t}, \phi) \right] = 0 \quad (19)$$

as the relevant condition. The substitution is feasible because optimality implies that the optimal forecast equals the target, that is, the optimal action defined by (18) is the same as the one defined by (19).

The intuition for the substitution is the following: if an inflation targeting central bank cares more about inflation above the target than inflation below the target (asymmetric loss in control error space) then the substitution implies that for this central bank inflation above the forecast is more costly than inflation below the forecast (asymmetric loss in forecast error space). Asymmetry in the forecasting loss function induces the optimal action of systematically over-predict inflation, which in turn leads the central bank to avoid that under-prediction delays a necessary tightening of the monetary policy position. The bias induced by the forecasting asymmetric loss helps to achieve the objective of the control asymmetric loss of not having inflation systematically above the target.

The substitution depends on having a loss function that gives a FOC that permits solving for the inflation target. There are numerous cases where the FOC takes the required form. Some of them are: for quadratic loss $E^f [\pi_{t+l} \mid \Omega_t] = \pi_{t+l,t}^T$, where the superscript denotes the forecast, in this case of the conditional mean; for linex loss $E^f [\pi_{t+l} \mid \Omega_t] + \frac{\phi}{2} var^f [\pi_{t+l} \mid \Omega_t] =$

⁴⁷Regularity conditions are needed to interchange integral and derivation. In this paper those conditions are assumed to hold.

⁴⁸With a quadratic loss function, condition (17) implies that the optimal control error follows a martingale difference sequence (if $l = 1$).

$\pi_{t+l,t}^T$; and for asymmetric quadratic loss the forecast is the ϕ th expectile of the conditional density forecast of inflation.⁴⁹

If the loss function is known, condition (19) can be used to evaluate the optimality of a particular sequence of forecasts. The test consists on finding whether $L'(\pi_{t+l} - f_{t+l,t}, \phi)$ is uncorrelated to v_t , and power against alternative hypotheses is achieved by selecting the appropriate v_t . This is a generalization of the Mincer-Zarnowitz regression as discussed by Granger (1999) and Elliott, Komunjer and Timmermann (2004). If the loss function is not known, condition (19) and a sequence of forecasts can be used to estimate ϕ , provided it is identified, using v_t as instrument. Finally, one can also evaluate the sequence of forecasts conditioning on the estimated value of ϕ provided enough instruments are available. This is the approach suggested by Elliott, Komunjer and Timmermann (2004, 2005).

Condition (19) can shed light on the discussion about which data should be used as actual data for forecast evaluation: Revised data, allegedly more closely to the “true”, or real-time data. There are two places where actual data can be used. One is for the actual value inside the marginal loss that appears in condition (19). It is not clear what a central bank is forecasting, and arguments can be made both ways. The second place is in the central bank’s information set. The set contains past forecast errors and past values of inflation. The theory indicates that the content of the information set has to be known to the central bank at the moment at which the decision is made. Therefore, real-time data have to be used, as revised data are not in the information set at that time. This settles the discussion. To evaluate a central bank as a forecaster, and to learn from this exercise, real-time data have to be used. The rest of the paper continues to report some results with both data sets, but the reader should bear in mind that theory favours results using real-time data.

To use condition (19) with actual forecasts of inflation one should use forecasts at horizon l . But there are two problems. The first is that it is not clear what l is, although most researchers would agree that it is more than one year. The second problem is that the Green Book forecasts, as mentioned before, are made for horizons (h) zero to nine (in quarters), but unfortunately forecasts for more than four quarters ahead are not very common, so the longer horizon one can work with is a year. This horizon is the one more likely to contain information about the Fed’s preferences, although is also the one for which the sample is smaller. The rest of the paper will use h instead of l knowing that h , though all that is available, is not the best fit to what the theory indicates.

3.2.2 Asymmetric Quadratic Loss

Under quadratic loss the generalized error, $L'(\pi_{t+h} - f_{t+h,t}, \phi)$, is identical to the forecasting error, $\pi_{t+h} - f_{t+h,t}$. This is one of the reasons quadratic loss is so popular: it gives results that directly concern the errors and not a transformation of them. Under quadratic loss the forecast errors follow a martingale difference sequence (for $h = 1$), so that any variable in the information set of the forecaster has to be orthogonal (uncorrelated) to the forecast errors if the forecasts are optimal. Going back to equations (2), (4), (5), and (6) one can see the rationale of those equations more clearly. These equations assume quadratic loss and test rationality by using variables from the forecaster’s information set (in this case the Federal

⁴⁹See Newey and Powell (1987) on expectiles.

Reserve). In equation (6) v_t is a vector that contains a constant, past forecast errors, and forecast errors from the SPF consensus. When a researcher finds a significant correlation between a variable in the forecaster's information set and the forecast error one of two things can be happening. The first is that the forecaster is using a symmetric loss function to obtain the forecasts, but that she or he is not using the information in an efficient way (i.e., the forecasts are not optimal with respect to that particular variable). The second is that the forecaster is using an asymmetric loss function to obtain the forecasts, and then the variable in the information set has to be uncorrelated with a transformation of the error (i.e., the generalized error) but can be correlated with the error. In the latter the forecasts would be rational.

The problem from an empirical perspective when traditional tests are used (like regressions 1 to 6) is that the only information available to the researcher is the evidence of correlation between the forecast error and the variable in the information set. With that information is difficult for the researcher to distinguish between rejecting the hypothesis of rationality because the forecasts are irrational or rejecting the hypothesis of symmetric loss because the forecaster is actually using asymmetric loss. In formal terms, the researcher has little power to distinguish what is driving the rejection, irrationality or asymmetric loss. The argument is carefully explained in Elliott, Komunjer and Timmermann (2004). The results presented in section I suggest that if the Federal Reserve has a symmetric loss function it is not using available information efficiently. The alternative is that the information is being used efficiently, but that the Federal Reserve has an asymmetric loss function.

Under an asymmetric loss function the generalized error, not the forecast error, is the one that follows a martingale difference sequence (for $h = 1$). For example, when linex loss is employed (along with an assumption about Gaussianity of inflation) the optimality condition is $E[\pi_{t+h} - E^f[\pi_{t+h} | \Omega_t] - \frac{\phi}{2} \text{var}^f[\pi_{t+h} | \Omega_t] | \Omega_t] = 0$. In this case an optimal bias exists, $E[e_{t+h,t} | \Omega_t] = \frac{\phi}{2} \text{var}^f[\pi_{t+h} | \Omega_t]$, and it depends on the asymmetry parameter ϕ and on the second moment of inflation.⁵⁰ As mentioned before, with linex loss and Gaussianity of inflation the analytical results are very tractable. But the Gaussianity assumption is difficult to justify for empirical work. This paper uses another asymmetric loss function, the asymmetric quadratic, to get the empirical results. It has the disadvantage from a theoretical perspective that the results are not as analytically tractable as with linex, but the advantage in empirical work that there is no need to assume Gaussianity.

Asymmetric quadratic loss, also called quad-quad loss, was introduced before in this paper using control errors. In a forecasting context it is:

$$L(e_{t+h,t}, \phi) = \left[\phi + (1 - 2\phi) 1_{(e_{t+h,t} < 0)} \right] |e_{t+h,t}|^2, \quad (20)$$

with $0 < \phi < 1$. ϕ is the asymmetry parameter: $\phi = 0.5$ corresponds to symmetry, whereas $\phi > 0.5$ corresponds to under-prediction more costly than over-prediction and vice versa for $\phi < 0.5$. For instance, if $\phi = 0.8$ under-predictions are approximately four times as costly as over-predictions.⁵¹ An asymmetric quadratic loss is shown in Figure 4 for $\phi = 0.5$ and

⁵⁰See Christoffersen and Diebold (1997, 1998) for more on forecasting under linex loss.

⁵¹Appendix A contains a detailed derivation of asymmetric quadratic loss, and an explanation of the interpretation of the asymmetry parameter.

$\phi = 0.8$.

Under asymmetric quadratic loss orthogonality condition (19) is (algebra is in appendix A):

$$E [v_t (e_{t+h,t} - (1 - 2\phi) |e_{t+h,t}|)] = 0, \quad (21)$$

for $v_t \subset \Omega_t$. To simplify the manipulation of the expression, the following change of variable will be used: $\theta = (1 - 2\phi)$. So $\theta \in \mathbb{R}$, and $\theta = 0$ corresponds to symmetry.

Expression (21) can be cast in a regression setting. This is useful to understand what is the difference between a quadratic loss and an asymmetric quadratic loss. Start with the following orthogonality condition:

$$E [v_t (e_{t+h,t} - \theta |e_{t+h,t}| - v_t' \delta)] = 0. \quad (22)$$

Equation (22) is satisfied if the forecasts are optimal, if they were produced using an asymmetric quadratic loss function with parameter ϕ , if v_t is in the information set of the producer of the forecasts, and if $\delta = 0$. Equation (22) is in the typical form of a GMM orthogonality condition, and implies the following regression:

$$e_{t+h,t} = \theta |e_{t+h,t}| + v_t' \delta + \varepsilon_{t+h}, \quad (23)$$

where it is clear that there is an omitted variable problem in equations (1) to (6) if the producer of the forecasts is using an asymmetric quadratic loss with $\theta \neq 0$. The omitted variable is the absolute value of the errors. As mentioned before, under asymmetric quadratic loss the optimal forecast is the ϕ th expectile of inflation, which means that knowledge about the location of the distribution (e.g., the mean) is not enough to calculate the optimal forecast. To give an idea of the role the absolute value of the errors is playing notice that under Gaussianity of the forecast errors $E [|e_{t+h,t}|] = \sqrt{\frac{2}{\pi}} \text{var} (e_{t+h,t})$, where $\text{var} (e_{t+h,t})$ is the variance of the error, so the absolute value is a measure of the dispersion of the distribution. Under normality, asymmetric quadratic has an interpretation as intuitive as linex, that is, the omitted variable (and the optimal bias) depends on the degree of asymmetry (measured by ϕ) and the dispersion of the distribution (measured by the variance under linex and by the absolute value, related to the standard deviation, under asymmetric quadratic). With omitted variable bias in equations (1) to (6) both the estimated values of the coefficients (including the constant) and their associated standard errors are biased, invalidating hypothesis testing.

To test if an asymmetric loss function is a possibility, the presence of the variance under linex and of the absolute error under asymmetric quadratic suggest that a variable that measures the dispersion of inflation can be used as a proxy for the omitted term. The Survey of Professional Forecasters contains not only the consensus forecast, the variable that has been used so far in this paper, but information about the forecast of each of the forecasters that answered the survey. The number of forecasters change with each survey, but a measure of the dispersion of the forecasts has been used in the past as a measure of the variance of inflation (Zarnowitz and Braun (1992)) and as a measure of heterogeneity in inflation expectations (Mankiw, Reis, and Wolfers (2003)).⁵² For this paper the interquartile

⁵²See the discussion in Rich and Tracy (2003).

range across forecasters is calculated, and the regression:

$$e_{t+h,t} = \beta \text{inqr}_{t+h,t} + \varepsilon_{t+h} \quad (24)$$

is estimated for each horizon using OLS and Newey-West standard errors. The sample used is the post-Volcker sample. The results are presented in Table 7 using real-time data for actual values of inflation.

The results indicate that β is significantly different from zero for every horizon. Under the null hypothesis of symmetric loss and rationality, a test of $\beta = 0$ is testing if information about the dispersion of the forecasts is in the Fed's information set when producing the forecasts given that the Fed uses a symmetric loss. The evidence rejects this hypothesis. The alternative hypothesis is either that the Fed has symmetric loss but that it is not using information contained in the spread of forecasts from SPF, something that points again to the Fed's irrationality, or that the Fed has an asymmetric loss, and that the spread across forecasters is working as a proxy for the omitted variable in the regression.⁵³

If one believes that information in the spread of the forecasters from the SPF is part of the Federal Reserve's information set, then the results in Table 7 support the hypothesis that the Federal Reserve has an asymmetric loss. If this is the case, the estimate of β is an estimate of θ (compare equations (23) and (24) under the null of asymmetric loss and rationality), and an estimate of the asymmetry parameter can be recuperated using $\phi = \frac{1-\theta}{2}$. If one takes the value of $\theta = -0.5$, the estimated value for horizon four, then the estimated asymmetry parameter is 0.76, which implies that for the Federal Reserve since Volcker under-prediction is between three and four times as costly as over-prediction. This estimate is preliminary because it is obtained under the assumption that the interquartile range is in the Fed's information set and that it is used efficiently by the Fed, something that has to be tested, not assumed.

3.2.3 GMM Estimation and Tests for Symmetry

Orthogonality condition (21) can be used to estimate the asymmetry parameter ϕ using the Generalized Method of Moments (GMM) developed by Hansen (1982). For a consistent estimate of the asymmetry parameter only one instrument is needed because only one parameter has to be estimated. To guarantee that the variable used as instrument is in the Fed's information set a constant can be used as the instrument. If this is the case, the orthogonality condition is:

$$E [(e_{t+h,t} - \theta |e_{t+h,t}|)] = 0. \quad (25)$$

The intuition behind the estimation is simple. If the sample average of the errors is zero, then the estimate of θ would not be significantly different from zero, as the sample counterpart of orthogonality condition (25) would be satisfied only if θ is zero. If the sample average of the errors is not zero (i.e., if there is a bias), the value of θ is adjusted until the sample counterpart of the orthogonality condition is satisfied. Therefore the estimate

⁵³This does not imply that the interquartile range is a measure of the conditional variance of inflation, but rather than, in the absence of that variable, it captures some of its correlation with the forecasts errors. All that is needed for this is that the dispersion across forecasters and the dispersion of inflation be correlated.

of the asymmetry parameter is obtained by asking the question: what degree of asymmetry rationalizes the observed bias?

This method of estimation was originally proposed by Elliott, Komunjer and Timmermann (2004, 2005) using an instrumental variables estimator. They show the conditions under which the asymmetry parameter is identified for the case of an asymmetric quadratic loss function. The implication from what they find in their paper for this paper is that the most important assumption needed for identification is that the optimal parameter of the model used to produce the forecasts has to be inside the parameter space, so that the FOC used to derive the orthogonality condition is useful for finding the minimum. The assumption about the parameter being inside the parameter space guarantees that the FOC is necessary for the minimum, the fact that the loss function is convex indicates that the FOC is sufficient for the minimum.⁵⁴

Because orthogonality condition (21) must hold for every horizon ($h = 0$ to $h = 4$), there are two ways to estimate the asymmetry parameter. One is to estimate one parameter for each horizon. The second is to use all the horizons in a system. The second has the advantage of using the fact that the residuals in each of the implied regressions are correlated, giving a more efficient estimation. If the second strategy is used, one can further test the restriction that the asymmetry parameter is the same for all horizons. For this paper the second strategy is followed, and the results reported here include the restriction that the loss function is the same for all horizons (as well as tests of this restriction).⁵⁵ To clarify the estimation process, let $\mathbf{e}_t = [e_{t,t} - \theta |e_{t,t}|, \dots, e_{t+h,t} - \theta |e_{t+h,t}|]'$ be the $((h + 1) \times 1)$ vector containing the generalized errors. Notice the restriction that θ is the same for all horizons. With a constant as an instrument, $v_t = 1$, the sample counterpart of the orthogonality conditions can be expressed as the $((h + 1) \times 1)$ vector:

$$g_T = \frac{1}{T} \sum_{t=1}^T \mathbf{e}_t, \quad (26)$$

where T is the sample size. The GMM estimator $\hat{\theta}_T$ is the value of θ that minimizes the scalar $Q_T = [g_T' W_T g_T]$ where W_T is a positive definite weighting matrix which may be a function of the data. For all the estimations presented on the rest of the paper the inverse of the Newey-West (1987) estimate of the asymptotic variance of the sample mean of $\mathbf{e}_t \otimes v_t$ is used as the weighting matrix.

The estimation is done first for the post-Volcker sample using real-time data. The sample used is from the third quarter of 1979 to the second quarter of 1998 giving a total of 76 observations for each equation in the system.⁵⁶ The results of the estimation imply that

⁵⁴There are other technical conditions that have to be satisfied. The stochastic process of inflation has to be such that the expectations used in the orthogonality conditions exist, that at least two moments exist, and that there is not too much heterogeneity (Elliott, Komunjer and Timmermann (2005)). This paper assumes that the conditions are satisfied for the application at hand.

⁵⁵Estimation horizon by horizon was also done but is not reported. The results reported in the paper are a good summary of the results found horizon by horizon. The only detail that is worth mentioning is that the estimate of the asymmetry parameter has a slight tendency to increase with the forecast horizon.

⁵⁶Horizon zero is not used. The coefficient associated with this horizon was different than the coefficients associated with the other horizons (i.e., the Wald test of equality of coefficients rejected the null when horizon

the asymmetry parameter is $\hat{\phi} = 0.80$, with a standard error of 0.05, so that it is clearly statistically different from 0.5.⁵⁷ A p-value of 0.56 for the Wald test indicates that the restriction of the loss function being the same across horizons cannot be rejected.

The degree of asymmetry of the Federal Reserve is estimated to be around 0.8, which implies that for the Federal Reserve forecasts below inflation (under-predictions) are approximately four times as costly as forecasts above it (over-predictions). This estimate also implies that for the Federal Reserve inflation above the (implicit) inflation target is four times as costly than inflation below the target. Figure 4 plots the asymmetric loss implied by this estimate ($\phi = 0.8$) and compares it to the quadratic loss typically assumed in the literature ($\phi = 0.5$).

The technique can also be applied to the pre-Volcker sample. The problem is that the number of observations is 16 if all the horizons are used. If only horizons one and two are used, then 41 observations are available. With such a small number of observations the estimates are likely to be severely biased. Further, information about the longer horizons has to be thrown away, which casts further doubts on the estimates. The result with 41 observations and using real-time data for inflation is that $\hat{\phi} = 0.25$ with a standard error of 0.11. The estimate is significantly different from 0.5. This result imply that for the pre-Volcker Federal Reserve under-predictions are approximately one third as costly as over-predictions. The asymmetry is reversed, which implies that for the Federal Reserve pre-Volcker inflation *below* the target was about three times more costly than inflation above the target. The restriction that the loss function is the same across horizons cannot be rejected (p-value of 0.50).

A Wald test was used to investigate if there is a change (in the sample pre-Volcker) of the estimates of the asymmetry parameter before and after the break of 1974-1975. The statistic is 0.46. A chi-square with one degree of freedom gives a p-value of 0.49. There is strong evidence to conclude that the asymmetry parameter can be considered to be the same for the entire sample pre-Volcker despite the first break. The interpretation is that the first break was not caused by a change in Fed's preferences about inflation.

3.3 Testing Rationality Allowing for Asymmetric Costs

The results obtained in section I point toward irrationality of the Federal Reserve inflation forecasts under symmetric loss. But the evidence presented so far in section II supports that the Federal Reserve has an asymmetric loss function. Do inflation forecasts look irrational once asymmetric costs are taken into account?

The orthogonality condition (21) is satisfied for every $v_t \subset \Omega_t$. Only one parameter has to be estimated, so that if v_t is a vector, then one of the variables can be used to estimate the asymmetry parameter and the others can be used to test if the orthogonality condition holds for them, conditioning on the estimated value of the asymmetry parameter. In a GMM framework this can be done using Hansen's test (or J-test) of overidentifying restrictions (Hansen (1982)) with the advantage that GMM uses all the instruments for estimation and

zero was included).

⁵⁷The estimate of θ is -0.5932, with a standard error of 0.10. The standard error of $\hat{\phi}$ was obtained by using $var(\phi) = \frac{var(\theta)}{4}$. For the rest of the results only the estimate of ϕ will be reported although what is directly obtained from the regression is an estimate of θ .

testing. For estimation, it does this by searching for the value of θ that makes a linear combination of the sample counterparts of each orthogonality condition (from each element in the vector v_t) as close as possible to zero. Conditional on the estimated value of θ , a J-test tests if the linear combination is close enough to zero so as to believe that each of the orthogonality conditions is close enough to be satisfied in population. To clarify, let the dimension of v_t be $k \times 1$. Then the sample counterpart of the orthogonality conditions can be expressed as the $((h + 1) k \times 1)$ vector:

$$g_T = \frac{1}{T} \sum_{t=1}^T \mathbf{e}_t \otimes v_t, \quad (27)$$

where T is the sample size. Again, the GMM estimator $\hat{\theta}_T$ is the value of θ that minimizes the scalar $Q_T = [g_T' W_T g_T]$. Hansen's J test statistic is TQ_T and it converges in distribution to a $\chi_{(h+1)k-1}^2$. As before, the orthogonality conditions for all horizons are used in a system with the restriction that the asymmetry parameter is the same for all horizons.

For the post-Volcker sample, the bottom panel of Table 8 presents the results using real-time data and a constant and one extra variable as instruments. The instruments are the variables used before in the paper: errors lagged $h + 1$ periods, the SPF consensus forecast, and the SPF interquartile range across forecasters. The results indicate that, given the estimated asymmetry parameter (which is between 0.8 and 0.9 and significantly different from 0.5), Hansen's test cannot reject rationality. The Federal Reserve is using an asymmetric loss to produce the forecasts and is efficiently using all the information contained in the instruments. Results with revised data (not reported) lead to the same conclusion, but with an estimated asymmetry parameter around 0.9.

One possible concern is that of weak instruments. In this context weak instruments refers to weak identification. Weak identification occurs if $E[\mathbf{e}_t \otimes v_t]$ is close to zero for $\phi \neq \phi_0$, where ϕ_0 denotes the parameter used to produce the optimal forecasts. According to Stock, Wright and Yogo (2002) if identification is weak then GMM estimates can be sensitive to the addition of instruments, so that if this occurs in an empirical application it can be indicative of weak identification. As can be seen in Table 8 the estimates do not change much when different instruments are used, which reversing Stock, Wright and Yogo's argument can be considered evidence of strong identification. If this is true it also implies that the tests for symmetry and rationality have good power (relative to the case of weak identification). Another evidence that the instruments are not weak is that the preliminary estimates of the asymmetry parameter obtained from the use of the spread across forecasters (equation (24)) are also similar to the estimates shown in Table 8. Notice that equation (24) does not include a constant, so that power is not obtained simply by the presence of it.

Another possible concern for some readers is that the explanation offered here may be explaining too much, in the sense that rationality cannot be rejected for any variable. This amounts as to say that the overidentification tests have no power against the alternative hypothesis of irrationality. To investigate this possibility, appendix B contains a Monte Carlo experiment in which the same method is able to correctly reject rationality once asymmetric loss is allowed for.

Rationality for the pre-Volcker sample is also tested. But the power of the tests is seriously

undermined because the number of observations is very small, so these results have to be taken with less confidence. The results using real-time data are presented in the upper panel of Table 8. They indicate that once asymmetric costs are taken into account the forecasts are rational. The estimates of the asymmetry parameter with different instruments are between 0.16 and 0.25. Results with revised data (not reported) lead to the same conclusion, but with an asymmetry parameter between 0.19 and 0.32.

The Wald tests that appear in Table 8 test the restriction that the loss function is the same across horizons. There is strong evidence that the loss function is the same across horizons.

4 Implications and Alternative Explanations

4.1 Implications of Asymmetric Loss

Section II provides empirical evidence that an asymmetric loss function is sensible for the Federal Reserve. The loss function could either be in control-error space (as used in the monetary policy literature) or in forecast-error space (as used in the forecasting literature). That the Federal Reserve has an asymmetric loss function is the explanation given in this paper for the evidence presented in section I about the Fed’s apparent irrationality. This explanation has some implications for the way inflation behaves in equilibrium and for the use researchers can give to the Green Book forecasts of inflation.

The first implication of asymmetric loss is that equilibrium inflation will not be on target (on average) as there exists an optimal bias induced by the asymmetric costs. Ruge-Murcia (2003) shows that in a model with an asymmetric loss function around an inflation target certainty equivalence no longer holds and therefore the expected marginal loss is nonlinear in the control error. The implication is that inflation can be on average below or above the target (a bias with respect to the target exists) depending on the type of asymmetry (the sign of ϕ). If inflation above the target is more costly for the central bank than inflation below the target (the case of the Federal Reserve since Volcker) the fear of having inflation above the target will induce the central bank to maintain inflation below the target. In the model this is reflected in the fact that an asymmetric central bank has a higher interest rate (everything else equal) than a symmetric central bank, because the asymmetric bank is setting the expected value of inflation to be below the target. Nobay and Peel (2003) named this phenomena deflationary bias. The reverse would be true for a central bank with preferences such as those of the pre-Volcker Fed.

The second implication of asymmetric loss is that higher moments of inflation, like the variance, are going to enter the process for the mean of inflation in equilibrium. In the model under linex loss, equilibrium inflation (16) follows a GARCH-in-mean process induced by the central bank’s choice of monetary policy.⁵⁸ If the assumption about Gaussianity is relaxed, then other moments are likely to be important. In the case of an asymmetric quadratic loss function one would expect the ϕ th expectile to matter for equilibrium inflation.

The third implication of asymmetric loss is not for the economy but for researchers that would like to work with Green Book inflation forecasts. If the Green Book forecasts were

⁵⁸See Ruge-Murcia (2003) and Engle, Lilien, and Robins (1987).

produced under a quadratic loss they would be the Fed’s expected value of inflation given its information set (as in (10), see Cochrane (2004)). But if they are produced, as it seems to be the case, by using an asymmetric loss function then they are not the expected value of inflation, but the expected value plus a bias term (as in (14)). To obtain the expectation one has to correct or de-bias the forecasts. From the results presented so far one can calculate an average factor that is useful as a rule-of-thumb to correct the forecasts. For the post-Volcker sample, the factor that seems appropriate is -0.5 for real-time data and -0.6 for fully revised data.⁵⁹ For example, if a Green Book forecast predicts inflation to be 3.0% four-quarters-ahead, then a good proxy for the Fed’s expected value of inflation four-quarters-ahead is 2.5%.⁶⁰

Finally, once asymmetric costs have been taken into account the implication that the Federal Reserve forecasts are rational is that the Fed can be modeled as having rational expectations, and that as Orphanides (2002) suggested, the actions taken by the Fed, even pre-Volcker, were optimal given their information.

4.2 Alternative Explanations

Other theories have been put forward to explain some of the empirical findings documented in this paper.

One argument to explain the bias in inflation forecasts is based on the Phillips curve theory and under- or over-estimation of the NAIRU (Nonaccelerating inflation rate of unemployment). Primiceri (2003) and Orphanides (2002) document under-estimation of the NAIRU during the sixties and seventies. Meyer (2004) writes about over-estimation of the NAIRU during the nineties. He relates that during the nineties the increase in productivity in the United States caused a decline of the NAIRU, but that the data was slow in showing the change in productivity, and therefore the Federal Reserve was expecting inflation to rise due to the low unemployment (though to be below the NAIRU) but that the rise never happened. This explanation certainly can be used to explain part of the bias, but the result documented in this paper is too systematic to be explained by an error in the estimation of the NAIRU. If the Fed does not like to over-predict inflation, then simply looking at past errors is enough to give a factor that corrects the bias.

A related argument is that of learning. In this case the forecasts would appear irrational while the Fed learns about a key aspect of the economy, for example the persistence of inflation (Primiceri (2003)). But this argument cannot explain the sudden change in the sign of the bias in 1979 nor can it account for the duration of it (20 years in the sample post-Volcker). A simple OLS learning mechanism is helpful to explain why. Suppose the parameter that is not known is the mean of inflation. At each point in time the Fed would estimate the mean with the available observations. If the first observation is far above the true mean of inflation and if the forecast of inflation is just the mean then the forecast would over-predict inflation for a while, but eventually the estimate will converge to the true value and the bias would disappear.

⁵⁹The factors are obtained from the estimated constants from Tables 5 and 6.

⁶⁰The bias depends on moments higher than the mean, and in the case of inflation these moments are likely to be time-varying (e.g., the variance), so the bias is likely to be time-varying. A more formal method to de-bias the forecasts has to take this time-varying component into account.

As another explanation one can think of a reversion to the mean mechanism. This because inflation was under-predicted when its level was high (the seventies) and over-predicted when its level was low (the nineties). In this case a symmetric loss function that depends not only on the forecast error but also on the level of inflation could be used to model the Federal Reserve's preferences. But a closer inspection of Figures 5 and 6 reveals that the period from 1979 to 1983 had an inflation level above 5% and a systematic *over-prediction* of inflation, invalidating the use of a level-dependent loss function as a way to model Federal Reserve's preferences.

One explanation of why there is information in the SPF consensus forecasts in the short run under symmetric costs is likely to be that the extra information by the Fed is about the timing of the monetary policy, something valuable for long horizons (Sims (2002)). This is a very interesting explanation and needs further research. Under symmetric loss there is no explanation (other than distraction) to explain why the information in the dispersion across forecasters from the SPF is not taken into account by the Federal Reserve.

5 Conclusion

This paper documents two facts about Federal Reserve inflation forecasts. The first is that there was a systematic under-prediction of inflation during the sixties and the seventies and a systematic over-prediction of inflation during the eighties and nineties. This change in behavior coincides with Volcker's appointment as Chairman in 1979. The second is that under quadratic loss the Federal Reserve is not efficiently using information contained in the consensus forecast of inflation and in the dispersion across forecasters from the SPF.

The immediate conclusion derived from these facts is that the Federal Reserve is not using information efficiently to forecast inflation and, therefore, to take monetary policy decisions. But this paper presents evidence to support the alternative explanation that the Federal Reserve has asymmetric costs of under- and over-prediction and that when allowance is made for these costs it uses information efficiently.

Thus Federal Reserve inflation forecasts are rational and incorporate the information contained in forecasts from the SPF, as Romer and Romer (2000) pointed out, but only after taking into account that the Federal Reserve has been cautious about inflation since Volcker, and was less careful about inflation before him.

The estimated degree of asymmetry is quite high, but so is the bias found in the forecasts. The empirical results indicate that the size of the bias post-Volcker (the results more confidently estimated) is consistent with the Federal Reserve's seeing inflation above an implicit target as four times more costly than inflation below it. Further research, perhaps using structural models, is needed to investigate to what extent the Federal Reserve of this period was overly cautious.

Among the theories in the literature that can be used to give alternative explanations to these findings the predominant one seems to be the learning theory. Under this theory, the Federal Reserve would have biased forecasts or appear to be using information inefficiently because of its lack of knowledge of some key aspect of the economy. But this theory cannot account, among other things, for the twenty-year duration of the post-Volcker period bias, as it implies that any bias found in the forecasts would tend to evaporate over time.

To further reveal what its forecasts tell about the Federal Reserve, future research should incorporate output in the loss function and evaluate other forecasts contained in the Green Book. The approach used in this paper could also be applied to other forecasts, for example revenue forecasts from the Congressional Budget Office (Auerbach (1999)), to reveal what they tell about other forecasters and also to investigate the extent to which they include a precautionary term such as the one found in Federal Reserve inflation forecasts.

Appendix A. Mathematical Derivations

To derive the asymmetric quadratic loss, start with a piecewise asymmetric loss function:

$$L(e_{t+h,t}) = \begin{cases} aL(e_{t+h,t}) & e_{t+h,t} > 0 \\ 0 & e_{t+h,t} = 0 \\ bL(e_{t+h,t}) & e_{t+h,t} < 0 \end{cases} \quad (28)$$

where $a, b > 0$. If $L(e_{t+h,t}) = |e_{t+h,t}|^p$ then this is the family of asymmetric functions defined in Elliott, Komunjer, and Timmermann (2004, 2005). With $p = 2$, this is the asymmetric quadratic loss, with a giving the weight attached to positive errors (under-prediction) and b giving the weight attached to negative errors (over-prediction). $a = b$ gives symmetry in the sense that errors of the same magnitude but different signs receive the same weight. The loss is not differentiable at zero, but it is continuous.

Define the asymmetry parameter as $\phi = \frac{a}{a+b}$, so that $0 < \phi < 1$. Then the asymmetric quadratic loss function can be written as:

$$L(e_{t+h,t}) = (a+b) \left[\phi + (1-2\phi) 1_{(e_{t+h,t} < 0)} \right] |e_{t+h,t}|^2, \quad (29)$$

where $1_{(e_{t+h,t} < 0)}$ is the indicator function that equals one if the error is negative and zero if it is positive. This loss function is homogeneous, so that the first factor $(a+b)$ is just a scale factor and can be normalized to one. This normalization gives equation (20).

The interpretation of the asymmetry parameter is clear. $\phi = 0.5$ gives $a = b$ so that it corresponds to symmetry. Further, after some algebra one can get: $\frac{a}{b} = \frac{\phi}{1-\phi}$. For example, if $\phi = 0.8$, then $\frac{a}{b} = 4$, so that positive errors are weighted four times more than negative ones (are four times as costly). With $\phi = 0.2$, $\frac{a}{b} = 0.25$ so that positive errors have one fourth of the cost of negative errors.

To obtain orthogonality condition (21) one needs to solve the following problem:

$$\min_{f_{t+h,t}} E [L(e_{t+h,t}) | \Omega_t], \quad (30)$$

using asymmetric quadratic loss. The first order condition (necessary and sufficient due to the convexity of the loss function) is:

$$\frac{\partial}{\partial f_{t+h,t}} E \left[\left[\phi + (1-2\phi) 1_{(e_{t+h,t}^* < 0)} \right] |e_{t+h,t}^*|^2 | \Omega_t \right] = 0, \quad (31)$$

where the asterisk, $*$, denotes optimality. The loss function is not differentiable at zero, but because of the continuity of the function the derivative can be taken using ‘‘Dirac’’ Delta δ . Provided that integral and differentiation operators can be interchanged (which is assumed

in this paper), the derivative is:

$$E_t \left[\begin{array}{c} -2\phi \left(1 - (2) 1_{(e_{t+h,t}^* < 0)}\right) |e_{t+h,t}^*| + \\ (1 - 2\phi) \frac{\partial}{\partial f_{t+h,t}} 1_{(e_{t+h,t}^* < 0)} |e_{t+h,t}^*|^2 - \\ 2(1 - 2\phi) 1_{(e_{t+h,t}^* < 0)} \left(1 - (2) 1_{(e_{t+h,t}^* < 0)}\right) |e_{t+h,t}^*| \end{array} \right] = 0, \quad (32)$$

where E_t denotes the expectation conditional on Ω_t . Using “Dirac” Delta $\delta(\cdot)$ one gets:

$$E_t \left[\begin{array}{c} 2 \left(1_{(e_{t+h,t}^* < 0)} - \phi\right) |e_{t+h,t}^*| - \\ (1 - 2\phi) |e_{t+h,t}^*|^2 \delta(e_{t+h,t}^*) \end{array} \right] = 0, \quad (33)$$

which can be further simplified to:

$$E_t \left[\left(1_{(e_{t+h,t}^* < 0)} - \phi\right) |e_{t+h,t}^*| \right] = 0. \quad (34)$$

The last expression indicates that the optimal forecast is the ϕ th expectile of the expected distribution of the variable of interest given the information set.⁶¹

From the last expression one can see that the orthogonality condition is:

$$E \left[v_t \left(1_{(e_{t+h,t}^* < 0)} - \phi\right) |e_{t+h,t}^*| \right] = 0. \quad (35)$$

Expression (21) is the same as this last expression, except that the following algebraic change is applied to (34):

$$\begin{aligned} -2 \left(1_{(e_{t+h,t}^* < 0)} - \phi\right) |e_{t+h,t}^*| &= 2\phi |e_{t+h,t}^*| - (2) 1_{(e_{t+h,t}^* < 0)} |e_{t+h,t}^*| \\ &= 2\phi |e_{t+h,t}^*| - [|e_{t+h,t}^*| - e_{t+h,t}^*] \\ &= (2\phi - 1) |e_{t+h,t}^*| + e_{t+h,t}^* \\ &= e_{t+h,t}^* - (1 - 2\phi) |e_{t+h,t}^*|. \end{aligned}$$

⁶¹Expectile as defined in Newey and Powell (1987).

Appendix B. Monte Carlo Experiment

The null hypothesis of interest is that the forecasts are rational. To test it, this paper uses a J-test that allows for the true value of the asymmetry parameter, ϕ_0 , to be unknown. This appendix contains a Monte Carlo experiment that investigates the power of this J-test.

Intuitively, the power of the test arises from the existence of more than one moment conditions when only one parameter has to be estimated. The moment conditions arise from different forecast horizons and different instruments.⁶² For each moment condition, an estimate of the asymmetry parameter is obtained. This estimate is answering the question: What asymmetry parameter rationalizes the observed sequence of forecasts? If the asymmetry parameter implied by all moment conditions is the same (taking into account sampling variation), then the test statistic is small and the test fails to reject the null hypothesis. The conclusion in this case is that the forecasts are rational and that the degree of asymmetry is the common value estimated. However, if there is no common value, the test statistic becomes large and the test rejects the null of rationality. The conclusion in this case is that the information contained in the instruments was not used efficiently in the production of the forecasts.

The Setup

For each of M Monte Carlo repetitions, T observations are generated from the following bivariate normal distribution:

$$\begin{pmatrix} y_{t+1} \\ z_t \end{pmatrix} \sim N \left(\begin{pmatrix} \mu_y \\ \mu_z \end{pmatrix}, \begin{pmatrix} \sigma_y^2 & \sigma_{yz} \\ \sigma_{zy} & \sigma_z^2 \end{pmatrix} \right), \quad (36)$$

where y_{t+1} and z_t are scalars. With the simulated data, T actual values for the variable of interest are obtained from the following conditional distribution:

$$y_{t+1}|z_t \sim N(\mu_y - B\mu_z + Bz_t, \sigma_y^2 - B^2\sigma_z^2), \quad (37)$$

where $B = (\sigma_z^2)^{-1}\sigma_{yz}$. Notice that the conditional distribution presents dynamics only in the mean. Next, T forecasts are generated from:

$$y_{t+1,t} = F_{t+1,t}^{-1}(\phi_0), \quad (38)$$

where $0 < \phi_0 < 1$ is the asymmetry parameter of the loss function and:

$$F_{t+1,t} \equiv N(\mu_y - B\mu_z + bBz_t, \sigma_y^2 - B^2\sigma_z^2), \quad (39)$$

where b is the bias term. When $b = 1$ there is no bias and the forecasts are the optimal forecasts (compare 37 and 39).

To calculate the J-statistic, an estimate of ϕ_0 is needed. Following Elliott, Komunjer and

⁶²In the paper, the instruments are a constant, past forecasts errors, and past forecasts from the SPF (the consensus and the spread across forecasters).

Timmermann (2004) the estimate is the one that solves:

$$\widehat{\phi}_T \equiv \arg \min \widehat{g}'_T \widehat{S}^{-1} \widehat{g}_T, \quad (40)$$

where

$$\widehat{g}_T = \frac{1}{T} \sum_{t=1}^T v_t \left(\widehat{\phi}_T - 1_{(e_{t+1,t} < 0)} \right) |e_{t+1,t}|^{p-1} \quad (41)$$

and $\widehat{S} = \frac{1}{T} \sum_{t=1}^T v_t v_t' \left(1_{(e_{t+1,t} < 0)} - \widehat{\phi}_T \right)^p |e_{t+1,t}|^p$.⁶³ v_t is a $k \times 1$ vector of instruments, $e_{t+1,t} \equiv y_{t+1} - y_{t+1,t}$, 1 is an indicator function, and $p = 1$ for lin-lin and $p = 2$ for quad-quad.⁶⁴ The J-statistic is calculated as:

$$J_T \equiv T \widehat{g}'_T \widehat{S}^{-1} \widehat{g}_T, \quad (42)$$

under the null of rationality, J_T is distributed as χ^2_{k-1} .

The power is calculated as:

$$power = \frac{1}{M} \sum_{m=1}^M 1_{(J_T > \chi^2_{k-1, 0.95})}, \quad (43)$$

so that the size of the test is set at 5%.

The Experiment

The number of Monte Carlo repetitions is $M = 5,000$. The values of the parameters are:

$$\begin{pmatrix} y_{t+1} \\ z_t \end{pmatrix} \sim N \left(\begin{pmatrix} 4 \\ 1 \end{pmatrix}, \begin{pmatrix} 2.5 & 0 \\ 0 & 0.5 \end{pmatrix} \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix} \begin{pmatrix} 2.5 & 0 \\ 0 & 0.5 \end{pmatrix} \right), \quad (44)$$

where the means, standard deviations and correlation are chosen to fit the data: Real-time inflation for the actual values, and the interquartile range across SPF's forecasters for the instrument. The other parameters are set at $\phi_0 = 0.2, 0.5, 0.8$ and $p = 1$ (lin-lin). The results are presented in Figures 7 ($T = 100$) and 8 ($T = 500$).

Comments

In general the experiment shows that the test has power to reject the null of rationality when the forecasts are irrational given a certain degree of asymmetry. However, two comments are in order.

First, when $T = 100$ (Figure 7) there is no power when the asymmetry parameter is close to either 0 or 1. But this only happens on one side of the bias. For instance, when $b > 1$ and z_t is positive (as it is most of the times in the experiment) the bias shifts the location of the conditional distribution to the right. In this case, there is no power when $\phi_0 = 0.8$. The

⁶³A $k \times k$ identity matrix is used for the first stage.

⁶⁴Patton and Timmermann (2004) show that when the process generating the variable of interest is a location scale process there is a one-to-one mapping between quantiles and expectiles.

reason is that *both* instruments rationalize the bias by estimating ϕ_0 to be very close to one (because of small sample bias), and since both moment conditions lead to the same estimate the test fails to reject the null. This is clearly a small sample problem, as can be seen when $T = 500$ (Figure 8). In the latter, the instruments are able to more precisely estimate the degree of asymmetry needed to rationalize the bias, but since the degree needed is different across instruments, the test is able to correctly reject the null. This small sample problem does not occur if in the same experiment one uses a z with a zero mean, because in this case the bias (with $\phi_0 = 0.8$) does not shift the conditional distribution to the right and the degree of asymmetry that rationalizes the bias is not close to one. So this problem can be avoided by having a large sample or by using different instruments. The latter is the approach followed in the paper.

Second, notice that the power curves for $\phi_0 = 0.2$ and 0.8 are mostly below the curve for $\phi = 0.5$. This indicates that the power is higher the smaller the asymmetry. But as is clear from Figure 8, when the bias is large the power is one for the degrees of asymmetry investigated.

Appendix C. Tables

Table 1: Rationality Tests for Federal Reserve Inflation Forecasts Under Quadratic Loss Using Real-Time Data. Equation is: $e_{t+h,t}^F = \alpha + \gamma e_{t-1,t-h-1}^F + \omega^C (e_{t+h,t}^F - e_{t+h,t}^C) + \varepsilon_{t+h}$

Forecast horizon	α	γ	ω^C	p-value	Sample	N ^a
0	-0.08 (0.09)	0.23** (0.10)	0.21* (0.12)	0.00	68:4–98:2	119
1	-0.07 (0.14)	0.22* (0.13)	-0.20 (0.19)	0.07	69:1–98:2	118
2	-0.01 (0.17)	0.37** (0.18)	-0.12 (0.27)	0.04	69:3–98:2	116
3	-0.32* (0.17)	0.11 (0.10)	-0.03 (0.15)	0.01	74:3–98:2	96
4	-0.22 (0.21)	0.07 (0.12)	-0.37** (0.17)	0.01	75:3–98:2	92

Source: Data from the Federal Reserve Bank of Philadelphia

Notes: e^F is the forecast error from Green Book forecasts of inflation, e^C is the forecast error from the median of SPF forecasts. The actual value of inflation is taken from the second revision available from the real-time database. t and h index the date and horizon respectively. OLS estimates. In parentheses robust standard errors using Newey-West with h lags. The p-value is for the test of the null hypothesis that the three parameters associated with the coefficients are equal to zero (Wald test with three df).

^a After adjusting endpoints.

* $p < 0.10$. ** $p < 0.05$.

Table 2: Rationality Tests for Federal Reserve Inflation Forecasts Under Quadratic Loss Using Revised Data. Equation as in Table 1

Forecast horizon	α	γ	ω^C	p-value	Sample	N ^a
0	-0.10 (0.09)	0.16* (0.08)	0.35** (0.10)	0.00	68:4–98:2	119
1	-0.08 (0.11)	0.38** (0.10)	-0.03 (0.18)	0.00	69:1–98:2	118
2	-0.05 (0.15)	0.44** (0.17)	-0.05 (0.19)	0.00	69:3–98:2	116
3	-0.40** (0.16)	0.21** (0.09)	0.04 (0.16)	0.00	74:3–98:2	96
4	-0.36 (0.23)	0.10 (0.15)	-0.27 (0.18)	0.00	75:3–98:2	92

Source: Data from the Federal Reserve Bank of Philadelphia

Notes: As in Table 1, except that the actual value of inflation is taken from the last vintage available from the real-time database as of May 2004.

^a After adjusting endpoints.

* $p < 0.10$. ** $p < 0.05$.

Table 3: Tests for Multiple Structural Changes in the Mean of Federal Reserve Inflation Forecast Errors Using Real-time Data

Forecast horizon	Specifications			Tests	
	Sample		N	UDmax	SupF(2/1)
0	66:1	98:2	130	7.79*	3.64
1	68:3	98:2	120	5.90	1.08
2	68:4	98:2	119	5.04	1.83
3	73:3	98:2	100	5.59	4.14
4	74:2	98:2	97	8.68*	7.01
Numbers of Breaks Selected			Estimates with Two Breaks		
	BIC	Sequential	T ₁	T ₂	
0	1	1	75:1	82:4	
1	1	0	74:3	79:2	
2	1	0	74:2	79:4	
3	1	0	79:3	86:1	
4	2	1	79:3	85:4	

Source: Data from the Federal Reserve Bank of Philadelphia. The program used is available from Professor Perron's web page: <http://econ.bu.edu/perron/code.html>

Notes: The supF tests and sequential selection of the number of breaks are constructed using heteroskedasticity and autocorrelation consistent covariance matrices using a quadratic kernel with automatic bandwidth selection following Andrews (1991). A size of 10% is used for the sequential tests.

* $p < 0.10$. ** $p < 0.05$.

Table 4: Tests for Multiple Structural Changes in the Mean of Federal Reserve Inflation Forecast Errors Using Revised Data

Forecast horizon	Specifications			Tests	
	Sample		N	UDmax	SupF(2/1)
0	66:1	98:2	130	58.02**	2.05
1	68:3	98:2	120	45.67**	4.32
2	68:4	98:2	119	22.76*	7.78*
3	73:3	98:2	100	8.29*	1.42
4	74:2	98:2	97	28.46**	7.65*
Numbers of Breaks Selected			Estimates with Two Breaks		
	BIC	Sequential	T ₁	T ₂	
0	1	1	75:1	80:3	
1	2	1	74:3	79:2	
2	2	1	74:3	79:1	
3	1	1	80:2	85:4	
4	2	1	80:1	85:4	

Source: Data from the Federal Reserve Bank of Philadelphia. The program used is available from Professor Perron's web page: <http://econ.bu.edu/perron/code.html>

Notes: As in Table 3 except that the actual value of inflation is taken from the last vintage available from the real-time database as of May 2004.

* $p < 0.10$. ** $p < 0.05$.

Table 5: Rationality Tests for Federal Reserve Inflation Forecasts Under Quadratic Loss Using Real-time Data and Subsamples. Equation is: $e_{t+h,t}^F = \alpha + \gamma e_{t-1,t-h-1}^F + \omega^C (e_{t+h,t}^F - e_{t+h,t}^C) + \varepsilon_{t+h}$

Forecast horizon	α	γ	ω^C	p-value	Sample	N ^a
Pre-1975						
0	0.34* (0.21)	0.41** (0.20)	-0.02 (0.18)	0.06	68:4-74:4	25
1	0.72* (0.42)	0.30 (0.22)	-0.30 (0.47)	0.00	69:1-74:4	24
2	0.68 (0.57)	0.18 (0.43)	-1.19 (0.79)	0.00	69:3-74:4	22
1975-1979						
0	-0.29 (0.27)	-0.12 (0.15)	0.28 (0.41)	0.48	75:1-79:2	18
1	-0.05 (0.41)	-0.03 (0.17)	-0.14 (0.49)	0.97	75:1-79:2	18
2	0.42 (0.43)	0.13 (0.13)	0.91** (0.42)	0.02	75:1-79:2	18
3	0.43 (0.39)	0.14 (0.12)	-0.42 (0.70)	0.27	75:1-79:2	18
4	0.89** (0.42)	0.09 (0.07)	-0.46 (0.56)	0.00	75:3-79:2	16
Post-1979						
0	-0.31** (0.08)	0.01 (0.12)	0.51** (0.10)	0.00	79:3-98:2	76
1	-0.53** (0.13)	-0.15 (0.12)	0.26** (0.11)	0.00	79:3-98:3	76
2	-0.45** (0.11)	0.11 (0.08)	0.11 (0.11)	0.00	79:3-98:2	76
3	-0.52** (0.14)	0.07 (0.13)	0.09 (0.15)	0.00	79:3-98:2	76
4	-0.56** (0.19)	-0.02 (0.12)	-0.19 (0.20)	0.00	79:3-98:2	76

Source: Data from the Federal Reserve Bank of Philadelphia.

Notes: As in Table 1.

^a After adjusting endpoints.

* $p < 0.10$. ** $p < 0.05$.

Table 6: Rationality Tests for Federal Reserve Inflation Forecasts Under Quadratic Loss Using Revised Data and Subsamples. Equation is: $e_{t+h,t}^F = \alpha + \gamma e_{t-1,t-h-1}^F + \omega^C (e_{t+h,t}^F - e_{t+h,t}^C) + \varepsilon_{t+h}$

Forecast horizon	α	γ	ω^C	p-value	Sample	N ^a
Pre-1975						
0	1.01** (0.19)	-0.24 (0.15)	0.17 (0.12)	0.00	68:4-74:4	25
1	1.70** (0.32)	-0.17* (0.10)	0.05 (0.29)	0.00	69:1-74:4	24
2	1.72** (0.6)	0.00 (0.39)	-0.09 (0.54)	0.00	69:3-74:4	22
1975-1979						
0	0.05 (0.25)	-0.37** (0.17)	0.53* (0.29)	0.01	75:1-79:2	18
1	0.18 (0.28)	0.12 (0.21)	0.04 (0.37)	0.72	75:1-79:2	18
2	0.70* (0.35)	-0.07 (0.07)	1.08** (0.27)	0.00	75:1-79:2	18
3	0.60** (0.25)	0.10 (0.11)	0.07 (0.31)	0.06	75:1-79:2	18
4	0.96** (0.24)	-0.03 (0.12)	-0.17 (0.28)	0.00	75:3-79:2	16
Post-1979						
0	-0.51** (0.10)	-0.01 (0.12)	0.58** (0.11)	0.00	79:3-98:2	76
1	-0.71** (0.10)	-0.09 (0.10)	0.32** (0.13)	0.00	79:3-98:3	76
2	-0.62** (0.11)	0.07 (0.13)	0.03 (0.17)	0.00	79:3-98:2	76
3	-0.64** (0.14)	0.18* (0.10)	0.14 (0.20)	0.00	79:3-98:2	76
4	-0.76** (0.24)	-0.10 (0.16)	-0.20 (0.23)	0.00	79:3-98:2	76

Source: Data from the Federal Reserve Bank of Philadelphia.

Notes: As in Table 2.

^a After adjusting endpoints.

* $p < 0.10$. ** $p < 0.05$.

Table 7: Testing Federal Reserve’s Use of the Spread Across Forecasters from SPF Using Real-time Data and the Sample Since P. Volcker.

Equation is: $e_{t+h,t} = \beta inqr_{t+h,t} + \varepsilon_{t+h}$

Forecast horizon	β	ϕ^a
0	-0.25** (0.11)	0.62** (0.06)
1	-0.39** (0.11)	0.69** (0.06)
2	-0.44** (0.11)	0.72** (0.05)
3	-0.48** (0.12)	0.74** (0.06)
4	-0.52** (0.15)	0.76** (0.07)

Source: Data from the Federal Reserve Bank of Philadelphia.

Notes: The sample is from the third quarter of 1979 to the second quarter of 1998 (76 observations). e denotes the forecast error from Green Book forecasts of inflation. $inqr$ denotes the interquartile range across forecasters from the SPF. The actual value of inflation is taken from the second revision available from the real-time database from the Philadelphia Fed. t and h index the date and horizon of the forecast respectively. OLS estimates. Numbers in parentheses are robust standard errors calculated using Newey-West procedure with number of lags equal to h .

^a The null hypothesis for the t-tests is $\phi = 0.5$.

* $p < 0.10$. ** $p < 0.05$.

Table 8: Rationality Tests for Federal Reserve Inflation Forecasts Under Asymmetric Quadratic Loss Using Real-time Data

Instruments	ϕ	J-stat	p-value	Wald test	p-value
Pre-Volcker					
Constant	0.25** (0.11)	-	-	0.45	0.50
C + lagged error	0.16** (0.07)	1.94	0.86	0.35	0.85
C + SPF median	0.21** (0.09)	2.06	0.84	1.07	0.30
C + SPF inqr	0.24** (0.09)	5.14	0.40	1.14	0.28
Since-Volcker					
Constant	0.80** (0.05)	-	-	2.06	0.72
C + lagged error	0.90** (0.02)	8.63	0.28	7.29	0.50
C + SPF median	0.82** (0.02)	9.26	0.23	4.12	0.84
C + SPF inqr	0.87** (0.02)	9.21	0.24	3.11	0.92

Source: Data from the Federal Reserve Bank of Philadelphia.

Notes: System GMM estimates imposing the restriction that the asymmetry parameter is the same across horizons. Numbers in parentheses are robust standard errors calculated using Newey-West procedure with 5 lags. Horizons: (1) Pre-Volcker only horizons one and two are used (two equations in the system); (2) For the sample since Volcker horizons one to four are used (four equations in the system). Samples: (1) For the pre-Volcker period the sample used goes from 1969:1 to 1979:2, except when the instrument used is lagged errors for which the sample starts 1969:3; (2) For the period since Volcker the sample used goes from 1979:3 to 1998:2. Instruments: (1) the lagged error is the forecast error of the Green Book Forecasts lagged $(h+1)$ quarters, where h is the forecast horizon; (2) SPF median is the consensus forecasts formed used the median across forecasters from the SPF; (3) SPF inqr denotes the interquartile range across forecasters from the SPF. The actual value of inflation is taken from the second revision available from the real-time database from the Philadelphia Fed. J-stat is the value of Hansen's test statistic used to test the over-identifying restrictions, for the p-value a chi-squared with five df is used for the pre Volcker period and with 19 df for the post Volcker period. The null for the Wald tests is that the asymmetry parameter is the same across horizons, for the p-value a chi-squared with one df is used for the pre Volcker period and with three df for the period since Volcker.

* $p < 0.10$. ** $p < 0.05$.

Appendix D. Figures




Figure 1: One- and Four-step-ahead Green Book Forecasts Errors with Revised Data

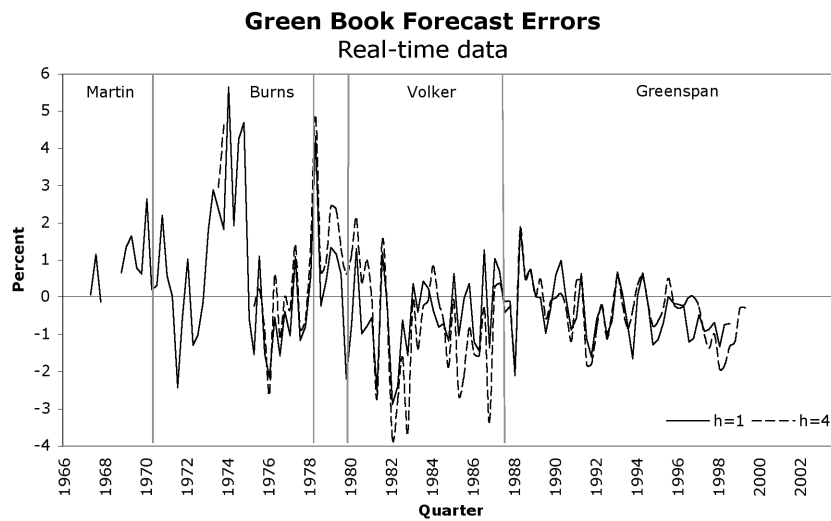


Figure 2: One- and Four-step-ahead Green Book Forecasts Errors with Real-time Data

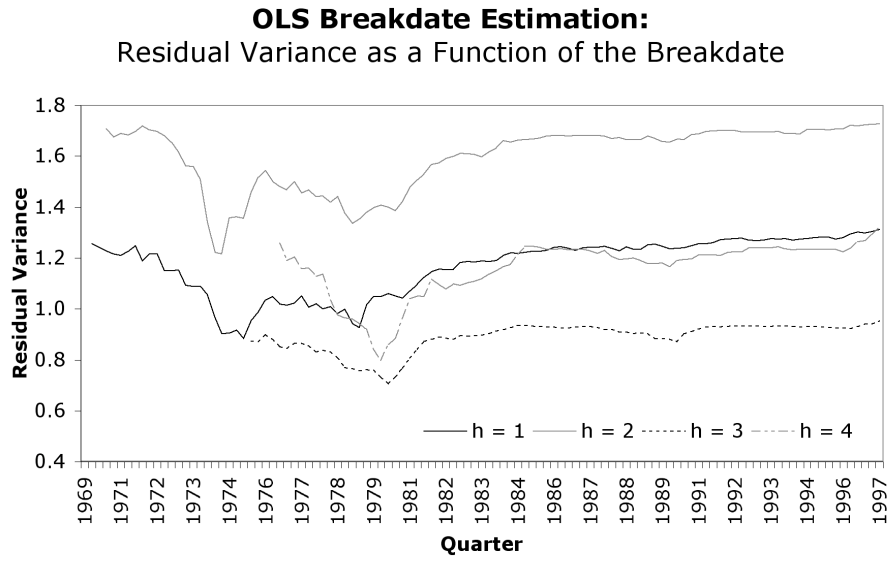


Figure 3: OLS Breakdate Estimation (single break)

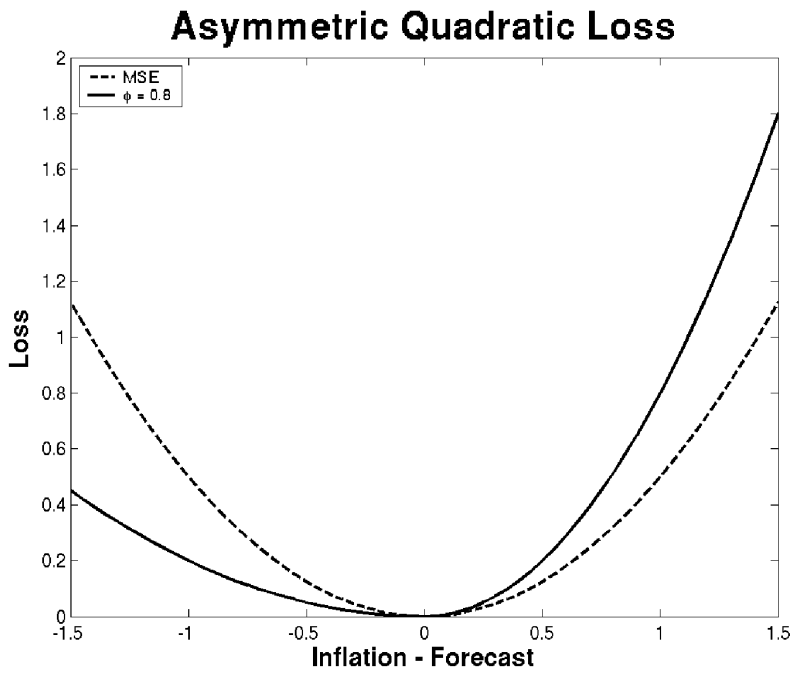


Figure 4: Symmetric and Asymmetric Quadratic Loss in Forecast Error Space




Figure 5: Inflation and One-step-ahead Green Book Forecasts with Revised Data

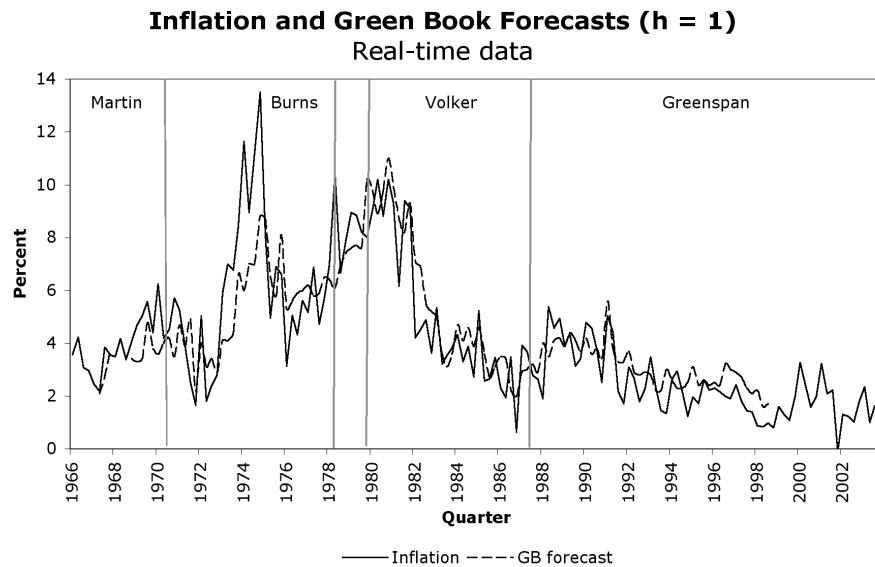


Figure 6: Inflation and One-step-ahead Green Book Forecasts with Revised Data




Figure 7: $T = 100$ and $M = 5000$




Figure 8: $T = 500$ and $M = 5000$

References

- [1] Andrews, D.W.K., 1991, "Heteroscedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59, 817-858.
- [2] Auerbach, A.J., 1999, "On the Performance and Use of Government Revenue Forecasts," *National Tax Journal*, December, 767-782.
- [3] Bai, J., and P. Perron, 1998, "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, 66, 47-78.
- [4] Bai, J. and P. Perron, 2003, "Computation and Analysis of Multiple Structural Change Models," *Journal of Applied Econometrics*, 18, 1-22.
- [5] Batchelor, R. and D.A. Peel, 1998, "Rationality Testing Under Asymmetric Loss," *Economics Letters*, 61, 49-54.
- [6] Bernanke, B.S., and M. Woodford, 1997, "Inflation Forecasts and Monetary Policy," *Journal of Money, Credit and Banking*, 24, 653-684.
- [7] Blinder, A.S., 1998, *Central Banking in Theory and Practice*. Cambridge: MIT Press.
- [8] Cavanagh, C.L., G. Elliott, and J.H. Stock, 1995, "Inference in Models with Nearly Nonstationary Regressors," *Econometric Theory*, 11, 1131-1141.
- [9] Chong, Y.Y., and D.F. Hendry, 1986, "Econometric Evaluation of Linear Macroeconomic Models," *Review of Economic Studies*, 53, 671-690. Reprinted in: C.W.J. Granger (ed.), 1990, *Modelling Economic Series*, Oxford: Clarendon Press.
- [10] Chow, G.C., 1960, "Test of Equality Between Sets of Coefficients in Two Linear Regressions," *Econometrica*, 28, 591-605.
- [11] Christoffersen, P.F., and F.X. Diebold, 1997, "Optimal Prediction Under Asymmetric Loss," *Econometric Theory*, 13, 808-817.
- [12] Clarida, R., J. Gali, and M. Gertler, 2000, "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, February, 147-180.
- [13] Clements, M.P., and D.F. Hendry, 1998, *Forecasting Economic Time Series*. United Kingdom: Cambridge University Press.
- [14] Croushore, D., 1993, "Introducing: The Survey of Professional Forecasters," Federal Reserve Bank of Philadelphia *Business Review*, November/December, 3-13.
- [15] Croushore, D. and T. Stark, 2000, "A Funny Thing Happened On the Way to the Data Bank: A Real-Time Data Set for Macroeconomists," Federal Reserve Bank of Philadelphia *Business Review*, September/October, 15-27.

- [16] Croushore, D. and T. Stark, 2002, "Forecasting with a Real-Time Data Set for Macroeconomists," *Journal of Macroeconomics*, 24, 507-531.
- [17] Cochrane, J.H., 2004, Comments on "A New Measure of Monetary Shocks: Deviations and Implications" by C.D. Romer and D.H. Romer, *NBER EFG meeting*, July.
- [18] Deutsch, M., and C.W.J. Granger, 1992, "Comments on the Evaluation of Policy Models," *Journal of Policy Modelling*, 14, 497-516.
- [19] Diebold, F.X., T.A. Gunther, and A.S. Tay, 1998. "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39, 863-883.
- [20] Elliott, G., I. Komunjer, and A. Timmermann, 2004, "Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?," *manuscript, UCSD*.
- [21] Elliott, G., I. Komunjer, and A. Timmermann, 2005, "Estimation and Testing of Forecast Rationality under Flexible Loss," *Forthcoming, Review of Economics Studies*.
- [22] Elliott, G. and A. Timmermann, 2004, "Optimal Forecast Combinations Under General Loss Functions and Forecast Error Distributions," *Journal of Econometrics*, 122, 47-79.
- [23] Engle, R.F., D.M. Lilien, and R.P. Robins, 1987, "Estimating Time Varying Risk Premia in the Term Structure: The ARCH-M Model," *Econometrica*, 55, 391-407.
- [24] Fair, R.C. and R.J. Shiller, 1989, "The Informational Content of Ex Ante Forecasts," *Review of Economics and Statistics*, 71, 325-331.
- [25] Gavin, W.T. and R.J. Mandal, 2001, "Forecasting Inflation and Growth: Do Private Forecasts Match those of Policy-makers?," *Business Economics*, 36, 13-20.
- [26] Granger, C.W.J., 1969, "Prediction with a Generalized Cost Function," *Operational Research*, 20, 199-207. Reprinted in: E. Ghysels, N.R. Swanson and M.W. Watson (eds.), 2001, *Essays in Econometrics: Collected Papers of Clive W.J. Granger*, Vol. I, Cambridge: Cambridge University Press.
- [27] Granger, C.W.J., 1999, "Outline of Forecast Theory using Generalized Cost Functions," *Spanish Economic Review*, 1, 161-173.
- [28] Granger, C.W.J. and P. Newbold, 1986, *Forecasting Economic Time Series*. 2nd Edition. New York: Academic Press.
- [29] Granger, C.W.J. and H.M. Pesaran, 2000, "Economic and Statistical Measures of Forecast Accuracy," *Journal of Forecasting*, 19, 537-560.
- [30] Granger, C.W.J. and R. Ramanathan, 1984, "Improved Methods of Combining Forecasts," *Journal of Forecasting*, 3, 197-204.
- [31] Greenspan, A., 2003, *Risk and Uncertainty in Monetary Policy*. Remarks at the Meetings of the American Economic Association, San Diego, California.

- [32] Gürkaynak, R.S., B. Sack, and E. Swanson, 2005, “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review*, 95, 425-436.
- [33] Hamilton, J.D., 1994, *Time Series Analysis*. Princeton: Princeton University Press.
- [34] Hansen, L., 1982, “Large Sample Properties of Generalized Method of Moments Estimators,” *Econometrica*, 50, 1029-1054.
- [35] Keane, M.P. and D.E. Runkle, 1990, “Testing the Rationality of Price Forecasts: New Evidence From Panel Data,” *American Economic Review*, 80, 714-735.
- [36] Leitch, G. and J.E. Tanner, 1991, “Economic Forecast Evaluation: Profits versus the Conventional Error Measures,” *American Economic Review*, 81, 580-590.
- [37] Mankiw, G.N., R. Reis, and J. Wolfers, 2003, “Disagreement about Inflation Expectations,” in: M.Gertler and K. Rogoff (eds.) *NBER Macroeconomics Annual 2003*. Cambridge: The MIT Press.
- [38] Meyer, L., 2004, *A Term at the Fed: An Insider’s View*. New York: HarperBusiness.
- [39] Mincer, J. and V. Zarnowitz, 1969, “The Evaluation of Economic Forecasts,” in: J. Mincer (ed.) *Economic Forecasts and Expectations*. New York: National Bureau of Economic Research.
- [40] Mishkin, F.S., 1981, “Are Market Forecasts Rational?” *American Economic Review*, 71, 295-306.
- [41] Mishkin, F.S., 2001, *The Economics of Money, Banking and Financial Markets*, 6th Edition. Reading: Addison-Wesley.
- [42] Muth, J.F., 1961, “Rational Expectations and the Theory of Price Movements,” *Econometrica*, 29, 315-335.
- [43] Newey, W., and J. Powell, 1987, “Asymmetric Least Squares Estimation and Testing,” *Econometrica*, 55, 819-847.
- [44] Newey, W. and K. West, 1987, “A Simple Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703-708.
- [45] Nobay, R.A., and D.A. Peel, 2003, “Optimal Discretionary Monetary Policy in a Model of Asymmetric Central Bank Preferences,” *The Economic Journal*, 113, 657-665.
- [46] Orphanides, A., 2002, “Monetary Policy Rules and the Great Inflation,” *American Economic Review Papers and Proceedings*, 92, 115-120.
- [47] Patton, A. and A. Timmermann, 2004, “Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity.” *Manuscript, UCSD*.

- [48] Primiceri, G.E., 2003, "Why Inflation Rose and Fell: Policymaker's Beliefs and US Postwar Stabilization Policy," *Manuscript, Princeton University*.
- [49] Rich, R. and J. Tracy, 2003, "Modeling Uncertainty: Predictive Accuracy as a Proxy for Predictive Confidence," *Federal Reserve Bank of New York Staff Reports*, no. 161.
- [50] Romer, C.D. and D.H. Romer, 2000, "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review*, 90, 429-457.
- [51] Romer, C.D. and D.H. Romer, 2004, "Choosing the Federal Reserve Chair: Lessons from History," *Journal of Economic Perspectives*, 18, 129-162.
- [52] Ruge-Murcia, F.J., 2000, "Uncovering Financial Markets' Beliefs about Inflation Targets," *Journal of Applied Econometrics*, 15, 483-512.
- [53] Ruge-Murcia, F.J., 2003, "Inflation Targeting Under Asymmetric Preferences," *Journal of Money, Credit and Banking*, 35, 763-785.
- [54] Sims, C.A., 2002, "The Role of Models and Probabilities in the Monetary Policy Process," *Brookings Papers on Economic Activity*, 2, 1-62.
- [55] Stekler, H.O., 2002, "The Rationality and Efficiency of Individuals' Forecasts," in M.P. Clements and D.F. Hendry (eds.), *A Companion to Economic Forecasting*, 222-240. Massachusetts: Blackwell Publishers Inc.
- [56] Stock, J.H., J.H. Wright, and M. Yogo, 2002, "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments," *Journal of Business and Economic Statistics*, 20, 518-529.
- [57] Svensson, L.E.O., 1997, "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets," *European Economic Review*, 41, 1111-1146.
- [58] Varian, H., 1974, "A Bayesian Approach to Real Estate Assessment," in S.E. Fienberg and A. Zellner (eds.), *Studies in Bayesian Econometrics and Statistics in Honor of L.F. Savage*, 195-208. Amsterdam: North-Holland.
- [59] Walsh, C.E., 2003, *Monetary Theory and Policy*, 2nd Edition. Cambridge: MIT Press.
- [60] Woodford, M., 2003, *Interest and Prices. Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.
- [61] Zarnowitz, V., 1985, "Rational Expectations and Macroeconomic Forecasts," *Journal of Business and Economic Statistics*, 3, 293-311.
- [62] Zarnowitz, V. and P. Braun, 1992, "Twenty-two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance," *NBER working paper 3965*.
- [63] Zellner, A., 1986, "Biased Predictors, Rationality and the Evaluation of Forecasts," *Economics Letters*, 21, 45-48.