Banco de México

Documentos de Investigación

Banco de México

Working Papers

N° 2006-14

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

Carlos Capistrán
Banco de México

December 2006

La serie de Documentos de Investigación del Banco de México divulga resultados preliminares de trabajos de investigación económica realizados en el Banco de México con la finalidad de propiciar el intercambio y debate de ideas. El contenido de los Documentos de Investigación, así como las conclusiones que de ellos se derivan, son responsabilidad exclusiva de los autores y no reflejan necesariamente las del Banco de México.

The Working Papers series of Banco de México disseminates preliminary results of economic research conducted at Banco de México in order to promote the exchange and debate of ideas. The views and conclusions presented in the Working Papers are exclusively the responsibility of the authors and do not necessarily reflect those of Banco de México.

Documento de Investigación
2006-14

Working Paper
2006-14

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

Carlos Capistrán†
Banco de México

Abstract

Inflation forecasts of the Federal Reserve seem to have systematically under-predicted inflation from the fourth quarter of 1968 until Volcker's appointment as Chairman, and to systematically over-predict it afterwards until the second quarter of 1998. Furthermore, under quadratic loss, commercial forecasts seem to have information not contained in those forecasts. To investigate the cause of this apparent irrationality, this paper recovers the loss function implied by Federal Reserve's inflation forecasts. The results suggest 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's inflation forecasts are found to be rational.

Keywords: Inflation forecasts, Forecast evaluation, Monetary policy.
JEL Classification: C53, E52

Resumen

Los pronósticos de inflación de la Reserva Federal parecen haber sub-predicho la inflación sistemáticamente a partir del cuarto trimestre de 1968 hasta que Volcker fue nombrado Presidente, y posteriormente haberla sobre-predicho sistemáticamente hasta el segundo trimestre de 1998. Más aún, bajo pérdida cuadrática, pronósticos comerciales parecen tener información no contenida en los pronósticos de la Reserva Federal. Para investigar la causa de esta aparente irracionalidad, se recupera la función de pérdida implicada por los pronósticos de inflación de la Reserva Federal. Los resultados sugieren que el costo de tener inflación por arriba de un objetivo de inflación implícito fue mayor que el costo de tener inflación por abajo del mismo para el período a partir de Volcker, y que lo opuesto fue cierto para el periódo anterior a Volcker. Una vez que estas asimetrías son consideradas, se encuentra que los pronósticos de la Reserva Federal son racionales.

Palabras Clave: Pronósticos de inflación, Evaluación de pronósticos, Política monetaria.

*Special thanks to Graham Elliott, Clive W.J. Granger and Allan Timmermann for many useful conversations. I thank Daniel Chiquiar, Alejandro Díaz de León, Charles Elkan, Marjorie Flavin, James D. Hamilton, Nir Jaimovich, Ivana Komunjer, Craig R.M. McKenzie, George Monokroussos, Valerie A. Ramey, Christina D. Romer, Barbara Rossi, Kevin K. Sheppard, Raimundo Soto, Alberto Torres and many conference and seminar participants for important comments and suggestions.

† Dirección General de Investigación Económica. Email: ccapistran@banxico.org.mx.

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. 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, the first part of this paper shows that closer inspection of a data set that extends the one used by the Romers and by Sims indicates 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 seems to coincide with Paul Volcker's appointment as Chairman and that was previously overlooked. It is shown that the forecasts systematically under-predicted inflation before Volcker and systematically over-predicted it afterwards until the second quarter of 1998. Moreover, once this change in behavior is taken into account, Federal Reserve inflation forecasts do not seem to have efficiently incorporated 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 have missed information contained in the consensus forecast and in the spread across the surveyed forecasters. These results hold 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 percentage point for the sample since Volcker, is also economically significant. This bias would be typically considered to imply that the forecasts are irrational, however, this need not be true. Unbiasedness of rational forecasts 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.[3] In this case,

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

[2]The theory of Rational Expectations says that rational agents have expectations that are optimal forecasts (Mishkin, 1981).

[3]Other papers present evidence that the evaluation of forecasts depends on the loss function. Leitch and

the optimal forecast is the mean plus an optimal bias term.[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. But, does it make sense for a central bank to have symmetric preferences? Some authors have argued that it does not when referring to central banks that have a loss function that has as one of its arguments the divergence of inflation from an inflation target. 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. 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. 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). These papers indicate that for a central bank with an inflation target, inflation below the target is less costly than inflation above it.

This paper uses a simple model of an inflation targeting central bank with asymmetric preferences to reconcile the evidence of the apparent inefficient use of information on the part of the Federal Reserve. 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. In the situation just described, the actions taken by the central bank depend asymmetrically on the forecasts, and positive and negative forecasts

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.

[4]See Christoffersen and Diebold (1997), Granger (1969, 1999), and Zellner (1986).

errors do not have the same consequences.

The literature in psychology has shown that forecasters behave as to minimize a posible asymmetric loss function when they care about the accuracy of the forecasts and when they are able of adjusting these forecasts in a way that incorporates any consequences of their errors (Weber, 1994). It is clear that the Federal Reserve cares about the forecasts, and that the producers of the forecasts, the staff at the Board of Governors of the Federal Reserve System, are capable of adjusting the forecasts. The adjustment is likely to incorporate the consequences of the errors because the producers may be acting as agents of the Federal Open Market Committee. In this context, it is possible that the producers report the forecasts as if using the loss function of their client in response to strategic considerations. In fact, Ehrbeck and Waldmann (1996), Laster et al. (1999), and Ottaviani and Sorensen (2006) justify asymmetric loss functions for individual forecasters when they show that the main goal of the agents is to influence their clients' assessment of their forecasting ability.

To investigate if the empirical evidence is consistent with an asymmetric-cost Federal Reserve, this paper recovers the Federal Reserve's loss function as implied by its forecasts. The method used is to derive moment conditions 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 empirical results are that starting with Volcker's appointment as Chairman and until the end of the sample, the second quarter of 1998, the Federal Reserve's cost of under-prediction was 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, and is in line with the literature that has suggested that there is a significant difference in the way monetary policy was conducted pre- and post-Volcker (e.g., Clarida, Galí and Gertler, 2000; Romer and Romer, 2004). 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 since Volcker.

The paper proceeds as follows. In section 2, the empirical properties of the Federal Reserve forecast errors are analyzed using an original regression for forecast evaluation, and their biases and lack of encompassing of commercial forecasts under quadratic loss are documented. To rationalize this evidence, in section 3 the loss function implied by Federal Reserve inflation forecasts is recovered and evidence is found of asymmetric costs of

under- and over- inflation prediction. Once these costs are taken into account, using over-identification tests the Federal Reserve is found to be using information efficiently. Section 4 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 5 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.[5] 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 Bank of Philadelphia has put together a series of Green Book forecasts of inflation and output starting November 1965, 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.[6] This is convenient because FOMC's meetings have not always being as regular as they are today.[7] The analysis presented throughout this paper uses the data at the quarterly frequency.[8]

The Green Book contains forecasts for more that 50 variables. This paper uses inflation forecasts for the output deflator.[9] The forecast horizon varies from the current quarter to

[5] Reifschneider, et al. (1997) describe the role played by models in forecasting and the monetary policy process at the Federal Reserve. Sims (2002) analyzes both, Green Book forecasts and forecasts that are directly obtained from the econometric models.

[6] More 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

[7] The committee currently meets every six weeks.

[8] Romer and Romer (2000) uses 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.

[9] The series being forecasted is quarter-to-quarter (annualized) inflation from the level of the nominal output's price index. From 1965 to 1991, the index used was the price deflator implicit in the Gross National Product, from 1992 to the third quarter of 1996 it was the price deflator implicit in the Gross Domestic Product, and since then it has being the Gross Domestic Product's chain-weighted price index. All the series are seasonally adjusted.

as many as nine quarters ahead. In this paper 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.[10] 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 of may 1998 as fully revised data.[11] 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.[12] This paper has more to say about real-time versus revised data, but the discussion is postponed until section 3 where it can be framed in the context of the theoretical model.

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.[13] 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. The Romers reach their conclusion about rationality by estimating, for each forecaster and forecast horizon, 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

[10] See Croushore and Stark (2002) (with discussion).

[11] Real-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 (2002). Revised data is also taken from the real-time data base, and corresponds to the last vintage available in May 2004.

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

[13] The Romers' sample ends in 1991 because of the lag in the release of Green Book forecasts, and to avoid the change from the use of Gross National Product to Gross Domestic Product.

for period $t + h$. Then the Mincer-Zarnowitz regression is:

$$\pi_{t+h} = \alpha + \beta f_{t+h,t} + \varepsilon_{t+h}, \tag{1}$$

and a test of rationality is that $\alpha = 0$ and $\beta = 1$.[14] 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.

To fully understand what the Mincer-Zarnowitz regression tests, one can think of imposing $\beta = 1$ and then on substracting 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:

$$e_{t+h,t} = \alpha + \varepsilon_{t+h}. \tag{2}$$

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 it 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 estimated value of α to the forecasts improves them.

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 1991.[15] 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.

There are some advantages of using equation (2) instead of equation (1) when testing for unbiasedness. First, only one parameter has to be estimated. Second, equation (1) requires

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

[15] Orphanides (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.

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 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.[16] Regression (2) does not present this problem because the dependent variable is not persistent and the explanatory variable is just a constant.

However, some objections have emerged over the years about the use of the Mincer-Zarnowitz regression to test rationality. Granger and Newbold (1986) indicate that the regression is only testing a necessary condition for the optimality of the forecasts. 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. 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.

The Romers also show that Green Book forecasts dominate commercial forecasts of inflation. They show this by running forecast combination regressions pairwise with the Green Book forecasts in each regression. The regression is:

$$\pi_{t+h} = \alpha + \omega^F f^F_{t+h,t} + \omega^C f^C_{t+h,t} + \varepsilon_{t+h}, \tag{3}$$

where α is a constant, ω^F is the weight assigned to the Federal Reserve forecast (denoted by $f^F_{t+h,t}$) and ω^C is the weight assigned to the commercial forecast (denoted by $f^C_{t+h,t}$).[17] 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 optimal action would be to throw away the commercial forecasts.[18] Sims (2002) reaches

[16]See Cavanagh, Elliott and Stock (1995).

[17]Combination regressions like (3) first appeared in Granger and Ramanthan (1984), and were later used by Chong and Hendry (1986) to test what they called "forecast encompassing". According to Hendry and Chong, a forecast "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.

[18]Romer and Romer (2000) conclude that if both, the Fed and commercial forecasters are using all their

the same conclusion using a similar methodology.

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, imposing the restriction $\omega^F + \omega^C = 1$, can be used:

$$e^F_{t+h,t} = \alpha + \omega^C \left(f^C_{t+h,t} - f^F_{t+h,t} \right) + \varepsilon'_{t+h}. \tag{4}$$

An encompassing test is simply the test of $\omega^C = 0$. An alternative is to use the forecast $f^C_{t+h,t}$ as the explanatory variable in regression (4) instead of the error difference, 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 $\omega^F = \left(1 - \omega^C\right)$.[19]

2.2 Empirical Evidence

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, p. 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 $t+1$

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.

[19]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.

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. This property can be tested using the regression:

$$e_{t+h,t} = \gamma e_{t+h-j,t-j} + \varepsilon_{t+h}, \tag{5}$$

with $j \geq h$. The hypothesis of no serial correlation corresponds to $\gamma = 0$.

The dependent variable in equations (2), (4), and (5) is the same, which suggests that a single regression can be used to tests for unbiasedness, serial correlation, and encompassing. Such a regression is used to analyze Green Book inflation forecasts. The SPF consensus forecast of inflation is used as representative of commercial forecasts.[20] The regression is:

$$e^F_{t+h,t} = \alpha + \gamma e^F_{t+h-j,t-j} + \omega^C \left(f^C_{t+h,t} - f^F_{t+h,t} \right) + \varepsilon_{t+h}. \tag{6}$$

OLS is applied to the available sample (1968:4 to 1998:4) for each horizon, with $j = h+1$.[21] 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.[22] 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. But Figure 1 contains more

[20] The 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 being 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

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

[22] The bandwidth was chosen so that h lags were included to calculate the variance covariance matrix. Newey and West's method was chosen to avoid ending with a non–positive definite matrix.

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.[23] This pattern will be exploited later in the paper.

With respect to serial correlation in the forecast errors, 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. Under the assumption of quadratic loss, 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, under the maintained hypothesis the the Fed has a quadratic loss function, 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.[24] 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 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 appear to be irrational.

[23] W.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 was in charge since August 1987 until February 2006. On Chairmen's preferences about inflation see Romer and Romer (2004).

[24] 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.

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.[25] Bai (1997) indicates that the OLS estimate of the break date is the date that minimizes the residual variance as a function of the breakdate. Figure 3 plots the residual variance for horizons one to four using revised data.[26] 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 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.[27]

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).[28] 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.[29] The overall conclusion about parameter constancy is that there is evidence of

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

[26] Horizon zero is not used because is not very informative as it has a dip from 1973 to 1981 with no clear minimum. The formal tests presented below use both real-time and revised data.

[27] The 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 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.

[28] $UDmax$ is a test of the null hypothesis of no structural break against an unknown number of breaks given the upper bound of three breaks. $SupF(l+1|l)$ is a test for l versus $l+1$ breaks. 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 $supF(l+l|l)$ test. Finally, T_1 and T_2 are the estimated breakdates based on a procedure that finds the global minimizer of the sum of squared residuals when two breaks are allowed (see Bai and Perron (2003)).

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

two structural breaks, one around the beginning of 1975 and a second around the end of 1979.

From an economic perspective both breaks coincide with negative supply shocks: In the 1973-1975 period the economy was hit by the first oil shock, a sharp increase in food prices due to 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, Galí 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.

2.2.3 Bias and Encompassing Considering the Breaks

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 second quarter of 1998. The results are presented in Tables 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 percentage point. 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.[30] So there was 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 show that 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. That the SPF forecasts contain more information when the post-Volcker sample is used indicates learning over time by commercial forecasters.[31]

The results about encompassing using the sample since Volcker 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

[30] It 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 Fed 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).

[31] This result, although interesting, is not pursued further in this paper.

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 the 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.[32] The fact that rationality is rejected 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 of their data.

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 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%, or that actual inflation turns out to be 2%. 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, 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

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

the United States, he is not able to reject symmetric preferences.[33]

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.[34] 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 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 forecast equal to the target (Svensson, 1997).

3.1 A Model of a Central Bank with an Asymmetric Loss

The preferences of the central bank over the possible realizations of inflation h periods ahead, π_{t+h}, are described by a loss function that indicates the costs associated with a particular realization of π_{t+h} and the central bank's inflation target, $\pi^T_{t+h,t}$, through the control error $ce_{t+h} = \pi_{t+h} - \pi^T_{t+h,t}$. The target is defined at t for $t+h$. The loss function will be denoted $L\left(\pi_{t+h} - \pi^T_{t+h,t}, \phi\right)$ where ϕ is a fixed parameter. It is assumed that the loss function is convex and that $L(0, \phi) = 0$. The loss function indirectly depends on the central bank's actions through the effect of its monetary policy instrument, i_t, on π_{t+h}.

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).[35] Second, the loss function only has the inflation's control error as an argument, whereas it 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. The loss function used here is meant as a reduced form of a more involved loss function and the conclusions from the empirical part will be interpreted accordingly.

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.[36] Denote this

[33]However, 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.

[34]However, none consider the possibility that if the level of inflation is close to zero, then a prudent central bank may find inflation *below* the target more costly (i.e., the loss function may not only be a function of the control error (inflation minus the target) but also of the level of inflation).

[35]The 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.

[36]The 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. Under some

information set by Ω_t for the decision taken at t. This set contains at least the current and past realizations of π_t, and i_t, the models of the economy used by the central bank, as well as all the past and present inflation targets. The optimal monetary policy action, i_t^*, solves:

$$\min_{i_t \in \mathcal{I}} E\left[L\left(\pi_{t+h} - \pi_{t+h,t}^T, \phi\right) \mid \Omega_t\right] \tag{7}$$

The optimal action will be a function of the contents of the information set, the target, and the loss function. Under quadratic loss, $L\left(ce_{t+h}, \phi\right) = \left(ce_{t+h}\right)^2$, the optimal action is the one that satisfies the first order condition (FOC):

$$E\left[\pi_{t+h} \mid \Omega_t\right]\Big|_{i_t^*} = \pi_{t+h,t}^T. \tag{8}$$

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 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 forecast targeting".

Now suppose that the central bank has an asymmetric loss function that it is homogeneous (e.g., linex, asymmetric quadratic, asymmetric linear) and that inflation follows a location scale process. In this case, using results from Granger (1999), the central bank's FOC is:

$$E\left[\pi_{t+h} \mid \Omega_t\right]\Big|_{i_t^*} + \kappa\, var\left[\pi_{t+h} \mid \Omega_t\right]\Big|_{i_t^*} = \pi_{t+h,t}^T, \tag{9}$$

where κ is a function of the asymmetry parameter.[37] The optimal policy would be to set the interest rate so as to make the expected value of inflation equal to the inflation target minus a precautionary term (if the loss asymmetry is in the direction of inflation above the target being 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 (9) implies that the interest rate chosen by a cautious central bank would

conditions, and without loss of generality, the multi-period problem can be broken into a sequence of period-by-period problems (Svensson, 1997). The model presented in this paper satisfies these conditions.

[37] For example, assuming linex loss and that π follows a conditional Gaussian distribution (Christoffersen and Diebold, 1997), the FOC is: $E\left[\pi_{t+h} \mid \Omega_t\right]|_{i_t^*} + \frac{\phi}{2}\, var\left[\pi_{t+h} \mid \Omega_t\right]|_{i_t^*} = \pi_{t+h,t}^T$. In this case, $\kappa = \frac{\phi}{2}$. Under this loss function, as ϕ approaches zero, the loss approaches a quadratic function, and we obtain the result in equation (8).

be higher than that chosen by a symmetric central bank.

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 it will 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.[38]

3.2 Estimation of the Asymmetry Parameter

3.2.1 Derivation of Moment Conditions

From the general optimization problem (7), the optimal monetary policy action satisfies the optimality condition (FOC):

$$E\left[L'\left(\pi_{t+h}-\pi^{T}_{t+h,t},\phi\right)\mid\Omega_t\right]=0, \tag{10}$$

where $L'\left(\pi_{t+h}-\pi^{T}_{t+h,t},\phi\right)$ denotes the derivative of the loss function with respect to the control error.[39] Following Granger (1999) and Patton and Timmermann (Forthcoming), 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 (10) implies that the optimal generalized error follows a martingale difference sequence with respect to the information set Ω_t. 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:

$$E\left[v_t L'\left(\pi_{t+h}-\pi^{T}_{t+h,t},\phi\right)\right]=0. \tag{11}$$

The literature that suggests a decision-theoretic approach to forecast evaluation (Granger and Pesaran (2000)) derives an orthogonality condition similar to (11). The difference is that the loss function in (11) depends on the control error $\pi_{t+h}-\pi^{T}_{t+h,t}$, whereas the loss function used in the forecasting literature depends on the forecasting error $\pi_{t+h}-f_{t+h,t}$. To link (11) with the condition used in the forecasting literature one can substitute the inflation target

[38] The optimal degree of asymmetry, and consequently the bias, can be time-varying and dependant on past performance as measured, for example, by the level of inflation. However, none of these are considered in this paper because they do not seem necessary to explain what is observed in the sample under study.

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

with the optimal forecast. By doing the substitution one gets:

$$E\left[v_t L^{'}\left(\pi_{t+h} - f_{t+h,t}, \phi\right)\right] = 0 \tag{12}$$

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 (11) is the same as the one defined by (12) (Svensson, 1997).

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). This asymmetry in the forecasting loss function induces systematic over-prediction of inflation, which in turn leads the central bank to avoid the costly mistake of having an interest rate below the one required to keep inflation on or below the target. 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.

If the loss function is known, condition (12) can be used to evaluate the optimality of a particular sequence of forecasts. The test consists on finding whether $L^{'}\left(\pi_{t+l} - f_{t+l,t}, \phi\right)$ 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 Elliott, Komunjer and Timmermann (2006). If the loss function is not known, condition (12) 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 (2005, 2006).

Condition (12) 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 (12). It is not clear what a central bank is forecasting, and arguments can be made both ways, although some would argue that the realized losses depend on true inflation. The second place is in the central bank's information set. The set contains past forecast errors and past values of inflation. But 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. 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 favors results using real-time data.

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, so that any variable in the information set of the forecaster has to be orthogonal to the forecast errors if the forecasts are optimal. One can see that equations (2), (4), (5), and (6) assume quadratic loss and test rationality by using variables from the forecaster's information set (in this case the Federal Reserve). In equation (6) v_t is a vector that contains a constant, past forecast errors, and other forecasts. 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 it 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 low power to distinguish what is driving the rejection, irrationality or asymmetric loss. The argument is carefully explained in Elliott, Komunjer and Timmermann (2005). The results presented in section 2 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.

The asymmetric loss used in this paper is the asymmetric quadratic loss, also called quad-quad loss. In a forecasting context it is:

$$L(e_{t+h,t}, \phi) = \left[\phi + (1 - 2\phi) 1_{\left(e_{t+h,t} < 0\right)}\right] |e_{t+h,t}|^2, \tag{13}$$

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.[40] An asymmetric quadratic loss is shown in Figure 4 for $\phi = 0.5$ and $\phi = 0.8$.

Under asymmetric quadratic loss orthogonality condition (12) is (algebra in Appendix):

$$E\left[v_t\left(e_{t+h,t} - (1-2\phi)\left|e_{t+h,t}\right|\right)\right] = 0, \tag{14}$$

for $v_t \subset \Omega_t$. Expression (14) 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:

$$E\left[v_t\left(e_{t+h,t} - (1-2\phi)\left|e_{t+h,t}\right| - v_t'\delta\right)\right] = 0. \tag{15}$$

Equation (15) 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 (15) is in the typical form of a GMM orthogonality condition, and implies the following regression:

$$e_{t+h,t} = (1-2\phi)\left|e_{t+h,t}\right| + v_t'\delta + \varepsilon_{t+h}, \tag{16}$$

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 $\phi \neq 0.5$. The omitted variable is the absolute value of the errors. Under asymmetric quadratic loss the optimal forecast is the ϕth expectile of inflation, which means that knowledge about the location of the distribution 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 $E\left[\left|e_{t+h,t}\right|\right] = \sqrt{\frac{2}{\pi} var\left(e_{t+h,t}\right)}$, where $var\left(e_{t+h,t}\right)$ is the variance of the error, so the absolute value is a measure of the dispersion of the distribution. Under normality, the asymmetric quadratic loss has the interpretation that the omitted variable and the optimal bias depend on the degree of asymmetry (measured by ϕ) and the dispersion of the distribution. 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.

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

To investigate if an asymmetric loss function is a possibility, the presence of the absolute error under asymmetric quadratic suggests 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, 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)). For this paper the interquartile range across forecasters is calculated, and the regression:

$$e_{t+h,t} = \beta inqr_{t+h,t} + \varepsilon_{t+h} \tag{17}$$

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 and asymmetric loss, and that the spread across forecasters is working as a proxy for the omitted variable in the regression.[41]

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 ϕ (compare equations (16) and (17) under the null of asymmetric loss and rationality), and an estimate of the asymmetry parameter can be recuperated. If one takes the value of $(1 - 2\phi) = -0.52$, the estimated value for horizon four, then the estimated asymmetry parameter is 0.76, which implies that for the Federal Reserve since Volcker took office and until the second quarter of 1998 under-prediction was 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.

[41]This 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.

3.2.3 GMM Estimation and Tests for Symmetry

The orthogonality condition (14) 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:

$$E\left[\left(e_{t+h,t} - (1-2\phi)\left|e_{t+h,t}\right|\right)\right] = 0. \tag{18}$$

The intuition behind the estimation is simple. If the sample average of the errors is zero, then the estimate of ϕ would not be significantly different from 0.5, as the sample counterpart of orthogonality condition (18) would be satisfied only if $\phi = 0.5$. If the sample average of the errors is not zero (i.e., if there is a bias), the value of ϕ is adjusted until the sample counterpart of the orthogonality condition is satisfied. Therefore the estimate 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 (2005, 2006) 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 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.[42]

Because orthogonality condition (14) 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 other is to use all the horizons in a system. The latter 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

[42] There 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 this conditions are satisfied.

the same for all horizons (as well as tests of this restriction).[43] To clarify the estimation process, let $\mathbf{e}_t = [e_{t,t} - (1-2\phi)\,|e_{t,t}|\,, ..., e_{t+h,t} - (1-2\phi)\,|e_{t+h,t}|]'$ be the $((h+1)\times 1)$ vector containing the generalized errors. Notice the restriction that ϕ 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:

$$g_T = \frac{1}{T}\sum_{t=1}^{T} \mathbf{e}_t, \tag{19}$$

where T is the sample size. The GMM estimator $\widehat{\phi}_T$ is the value of ϕ 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.[44] The results of the estimation imply that the asymmetry parameter is $\widehat{\phi} = 0.80$, with a standard error of 0.05, so that it is clearly statistically different from 0.5. 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, between the third quarter of 1979 and the second of 1998, the Federal Reserve under-predictions of inflation were approximately four times as costly as 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

[43] Estimation 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.

[44] Horizon 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 zero was included).

and using real-time data for inflation is that $\widehat{\phi} = 0.25$ with a standard error of 0.11. The estimate is significantly different from 0.5. This result implies that for the pre-Volcker Federal Reserve, under-predictions are approximately one third as costly as over-predictions. The asymmetry turns over, 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 no strong evidence against the conclusion 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 evidence presented in section 2 points toward irrationality of the Federal Reserve inflation forecasts if symmetric loss is assumed. But the results so far in section 3 show that the evidence can be rationalized by an asymmetric loss function. The rest of this section tests the rationality of inflation forecasts allowing for asymmetric costs by using over-identification tests.

The orthogonality condition (14) 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 (J-test) of overidentifying restrictions (Hansen (1982)) with the advantage that GMM uses all the instruments for estimation and testing. For estimation, it does this by searching for the value of ϕ 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 ϕ, 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:

$$g_T = \frac{1}{T} \sum_{t=1}^{T} \mathbf{e}_t \otimes v_t, \tag{20}$$

where T is the sample size. The GMM estimator $\widehat{\phi}_T$ is the value of ϕ 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 $\chi^2_{(h+1)k-1}$. 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. Thus the evidence supports a Federal Reserve that used an asymmetric loss to produce its forecasts and that, once these asymmetries are taken into account, efficiently used 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 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 (17)) are also similar to the estimates shown in Table 8. Notice that equation (17) does not include a constant, so that power is not being obtained simply by the presence of it.

Another possible concern could be that the explanation offered here may be explaining too much, in the sense that rationality cannot be rejected. This amounts as to say that the overidentification tests have no power against the alternative hypothesis of irrationality. To investigate this possibility, Capistrán (2005) 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 no strong evidence against the loss function being the same across horizons.

4 Implications and Alternative Explanations

4.1 Implications of Asymmetric Loss

The explanation given in this paper for the evidence about the Federal Reserve's apparent irrationality is that the Fed has an asymmetric loss function over inflation forecast errors. This explanation has some implications for the normative loss function typically postulated for the Federal Reserve, for the way inflation behaves in equilibrium, and for the use researches can give to the Green Book forecasts of inflation.

The Federal Reserve is directed by law to promote maximum employment and stable prices, and for this reason its loss function is typically modelled as a function of inflation deviations from a target and deviations of output from potential. In contrast, in this paper we have used a loss function that omits the output component. This loss is meant as a reduced form of the former, and the implication of the finding of asymmetry is that only functions that can be mapped into asymmetric losses over inflation are consistent with the data. Examples may be functions that are symmetric but not quadratic, such as the one proposed by Orphanides and Wieland (2000) that considers "inflation zone targeting", or functions that include precautionary terms, such as the one considered by Cukierman and Gerlach (2003) that considers a function where the Fed is more concerned about downward deviations of output from its potential than about upward deviations.

The second 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 until the end of the sample in 1998) the fear of having inflation above the target will induce the central bank

to maintain inflation below it. 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 estimated here for the pre-Volcker Fed.

The third implication of asymmetric loss is that higher moments of inflation, such as the variance, enter the process for the mean of inflation in equilibrium. In the model under asymmetric loss, equilibrium inflation will follow a GARCH-in-mean process induced by the central bank's choice of monetary policy. If the assumption about scale-location of ε is relaxed, then other moments are likely to be important. In the case of an asymmetric quadratic loss function one would expect the ϕth expectile to matter for equilibrium inflation.

The fourth implication of asymmetric loss is not for the economy but for researchers working with Green Book inflation forecasts. If the Green Book forecasts were produced under a quadratic loss they would be the Fed's expected value of inflation given its information set. 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. 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 Volcker-Greenspan sample, from the third quarter of 1979 to the second of 1998, the factor that seems appropriate is -0.5 for real-time data and -0.6 for fully revised data.[45] 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%.[46]

Finally, once asymmetric costs have being taken into account the implication of rationality is that the Fed can be modelled as having rational expectations, and that the actions taken by the Fed, even pre-Volcker, were optimal given their preferences and information.

4.2 Alternative Explanations

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

The first is that the Green Book forecasts are not forecasts but projections based on

[45] The factors are obtained from the estimated constants in Tables 5 and 6.

[46] The 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.

an assumed instrument path, and that most of the time the assumed path was one of no-change. Even if this were the case, this cannot account for the almost twenty year bias from the third quarter of 1979 until the end of the sample. If the forecasts were always assuming a no-change path for the instrument, then they will have over-predicted inflation when policy was tightened, and under-predicted when policy was loosened. But the path was not always assumed constant. Indeed, Reifschneider et al. (1997) report that, although the point of departure for the Fed staff most often is an assumption that the instrument, the federal funds rate, will remain unchanged over the forecast horizon, if the unchanged instrument path assumption is "at odds with the stated objectives of most policymakers"then another path is assumed. Furthermore, Reifschneider et al. (1997) also report that the forecasts are judgemental and that the forecasting process involve several intercept corrections by members of the staff, so that the possibility exists that even if the forecasts coming out from the model were based on a constant path, the forecasts that appeared in the Green Book may have been closer to forecasts that reflected the most likely path for the instrument.

Another argument to explain the bias in inflation forecasts is based on the Phillips curve theory and under- or over-estimation of the NAIRU (Nonnaccelerating inflation rate of unemployment). Primiceri (2006) 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, 2006). 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 since Volcker). A simple OLS learning mechanism is helpful to explain why. Suppose the parameter that is not know 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.

As another explanation one can think of a reversion to the mean mechanism. This because apparently 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 Federal Reserve's preferences. But 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. However, there are other reasons to believe that a loss function that is asymmetric and depends on the level of inflation may be useful. This is because for a cautious central bank inflation *below* the target could become *more* costly if the level of inflation is close to zero (deflation scare), as in this case the main instrument of monetary policy is at risk of being rendered useless. In this context, a cautious central bank may have an asymmetry against low inflation when inflation is close to zero, but an asymmetry against high inflation when inflation is safely above zero. This paper does not employ a level dependent loss because deflation scares did not occur during the sample period.

In a similar vein, it could be that the Federal Reserve was indeed irrational an produced its forecast using adaptive expectations. If this was the case, the forecasts would have under-predicted inflation when inflation had an increasing trend (as before Volcker) and would have over-predicted it when inflation had a decreasing trend (as was the case since Volcker until about 1998). The data appears to be consistent with this explanation (Figure 6). In addition, Reifschneider et al. (1997) indicate that at least up to 1997 the models used to support the forecasting process at the Board of Governors indeed used adaptive expectations.[47] The problem with this explanation is that it implies that the Fed did not systematically evaluated its own inflation forecasts errors. Hanson and Whitehorn (2006) explore this alternative explanation in detail.

Finally, another possible explanation is that since private forecast errors show a pattern of bias similar to that of the Fed's forecasts, the bias can not be related to individual incentives but has to be explained by systemic factors. However, although it is true that the mean across forecasters shows a similar pattern of bias, this is not true for individual forecasters. Capistrán and Timmermann (2006) analyze individual forecasters from the SPF since 1968 and show that there are forecasters that systematically over-predict inflation and forecasters that systematically under-predict it.

[47]Two different models were used as reference models during the sample, the MPS from the late 1960s until the beginning of 1996 and the FRB/US model from mid-1996 onwards.

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 was 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 would be, if one is willing to sustain symmetric loss, that the Federal Reserve was 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 had asymmetric costs of under- and over-prediction and that, when allowance is made for these costs, it seems to have used information efficiently.

Thus, the Federal Reserve inflation forecasts analyzed here seem to be rational and to incorporate the information contained in forecasts from the SPF, as Romer and Romer (2000) pointed out, but only if asymmetries in the loss function are taken into account. The Federal Reserve appears to have been cautious about inflation since Volcker and until the second quarter of 1998, and appears to have been less worried about it before him.

The estimated degree of asymmetry is high, but so is the bias found in the forecasts. The empirical results indicate that the size of the bias in the sample since 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. To further reveal what its forecasts tell about the Federal Reserve, future research should incorporate output in the loss function and evaluate other Green Book forecasts.

References

[1] Andrews, Donald W.K. "Heteroscedasticity and Autocorrelation Consistent Covariance Matrix Estimation." *Econometrica*, 1991, 59, pp. 817-858.

[2] Bai, Jushan. "Estimation of a Change Point in Multiple Regression Models." *The Review of Economics and Statistics*, 1997, 79, pp. 551-563.

[3] Bai, Jushan and Pierre Perron. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica*, 1998, 66, pp. 47-78.

[4] Bai, Jushan and Pierre Perron. "Computation and Analysis of Multiple Structural Change Models." *Journal of Applied Econometrics*, 2003, 18, pp. 1-22.

[5] Bernanke, Ben S. and Michael Woodford. "Inflation Forecasts and Monetary Policy." *Journal of Money, Credit and Banking*, 1997, 24, pp. 653-684.

[6] Blinder, Alan S. *Central Banking in Theory and Practice*. Cambridge: MIT Press, 1998.

[7] Capistrán, Carlos. "Bias in Federal Reserve Inflation Forecasts: Is the Federal Reserve Irrational or Just Cautious?" UCSD Discussion Paper: 2005-05.

[8] Capistrán, Carlos and Allan Timmermann. "Disagreement and Biases in Inflation Expectations." Banco de México Working Paper: 2006-07.

[9] Cavanagh, Christopher L., Graham Elliott and James H. Stock. "Inference in Models with Nearly Nonstationary Regressors." *Econometric Theory*, 1995, 11, pp. 1131-1141.

[10] Chong, Yock Y. and David F. Hendry. "Econometric Evaluation of Linear Macroeconomic Models." *Review of Economic Studies*, 1986, 53, pp. 671-690.

[11] Chow, Gregory C. "Test of Equality Between Sets of Coefficients in Two Linear Regressions." *Econometrica*, 1960, 28, pp. 591-605.

[12] Christoffersen, Peter F. and Francis X. Diebold. "Optimal Prediction Under Asymmetric Loss." *Econometric Theory*, 1997, 13, pp. 808-817.

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

[14] Croushore, Dean. "Introducing: The Survey of Professional Forecasters." Federal Reserve Bank of Philadelphia *Business Review*, November/December, 1993, pp. 3-13.

[15] Croushore, Dean and Thomas Stark."Forecasting with a Real-Time Data Set for Macroeconomists." *Journal of Macroeconomics,* 2002, 24, pp. 507-531.

[16] Cukierman, Alex and Stefan Gerlach."The Inflation Bias Revisited: Theory and Some Evidence." *The Manchester School,* 2003, 71, pp. 541-565.

[17] Elliott, Graham, Ivana Komunjer and Allan Timmermann. “Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss?” Manuscript, UCSD, 2006.

[18] Elliott, Graham, Ivana Komunjer and Allan Timmermann. “Estimation and Testing of Forecast Rationality under Flexible Loss.” *The Review of Economics Studies.*, 2005, 72, pp. 1107-1125.

[19] Ehrbeck, T. and R. Waldmann. “Why are professional forecasters biased? agency versus behavioral explanations.” *Quarterly Journal of Economics*, 1996, 111, pp. 21-40.

[20] Fair, Ray C. and Robert J. Shiller. “The Informational Content of Ex Ante Forecasts.” *Review of Economics and Statistics,* 1989, 71, pp. 325–331.

[21] Granger, Clive W.J. “Prediction with a Generalized Cost Function.” *Operational Research,* 1969, 20, pp. 199-207.

[22] Granger, Clive W.J. “Outline of Forecast Theory using Generalized Cost Functions.” *Spanish Economic Review*, 1999, 1, pp. 161-173.

[23] Granger, Clive W.J. and Paul Newbold. *Forecasting Economic Time Series.* 2nd Edition. New York: Academic Press. 1986.

[24] Granger, Clive W.J. and Hashem M. Pesaran. “Economic and Statistical Measures of Forecast Accuracy.” *Journal of Forecasting*, 2000, 19, pp. 537-560.

[25] Granger, Clive W.J. and Ramu Ramanathan. “Improved Methods of Combining Forecasts.” *Journal of Forecasting*, 1984, 3, pp. 197-204.

[26] Gürkaynak, Refet S., Bryan Sack and Eric Swanson. “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models.” *American Economic Review,* 2005, 95, pp. 425-436.

[27] Hansen, Lars P. “Large Sample Properties of Generalized Method of Moments Estimators.” *Econometrica*, 1982, 50, pp. 1029-1054.

[28] Hanson, Michael S. and Jayson Whitehorn. “Reconsidering the Optimality of Federal Reserve Forecasts.” *Manuscript, Wesleyan University*, 2006.

[29] Keane, Michael P. and David E. Runkle. “Testing the Rationality of Price Forecasts: New Evidence From Panel Data.” *American Economic Review*, 1990, 80, pp. 714-735.

[30] Laster, D., P. Bennett and I.S. Geoum. “Rational Bias in Macroeconomic Forecasts.” *Quarterly Journal of Economics*, 1999, 114, pp. 293-318.

[31] Leitch, Gordon and J. Ernest Tanner. “Economic Forecast Evaluation: Profits versus the Conventional Error Measures.” *American Economic Review,* 1991, 81, pp. 580-590.

[32] Mankiw, Gregory N., Ricardo Reis and Justin Wolfers. “Disagreement about Inflation Expectations.” in: M.Gertler and K. Rogoff (eds.) *NBER Macroeconomics Annual 2003.* Cambridge: The MIT Press. 2003.

[33] Meyer, Laurence. *A Term at the Fed: An Insider's View.* New York: HarperBusiness. 2004.

[34] Mincer, Jacob and Victor Zarnowitz. "The Evaluation of Economic Forecasts." in: J. Mincer (ed.) *Economic Forecasts and Expectations.* New York: National Bureau of Economic Research. 1969.

[35] Mishkin, Frederic S. "Are Market Forecasts Rational?" *American Economic Review,* 1981, 71, pp. 295-306.

[36] Newey, Withney and Kenneth West. "A Simple Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 1987, 55, pp. 703-708.

[37] Nobay, Robert A. and David A. Peel."Optimal Discretionary Monetary Policy in a Model of Asymmetric Central Bank Preferences." *The Economic Journal*, 2003, 113, pp. 657-665.

[38] Orphanides, Athanasios. "Monetary Policy Rules and the Great Inflation." *American Economic Review Papers and Proceedings*, 2002, 92, pp. 115-120.

[39] Orphanides, Athanasios and V. Wieland. "Inflation Zone Targeting." *European Economic Review*, 2000, 44, pp. 1351-1387.

[40] Ottaviani, M and P.N. Sorensen. "The Strategy of Professional Forecasting." *Journal of Financial Economics*, 2006, 81, pp. 441-466.

[41] Patton, Andrew and Allan Timmermann."Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity." *Forthcoming in the Journal of Econometrics.*

[42] Primiceri, Giorgio E. "Why Inflation Rose and Fell: Policymaker's Beliefs and US Postwar Stabilization Policy." *Quarterly Journal of Economics*, 2006, 1221, pp. 867-902.

[43] Reifschneider, David, L., Savid, J. Stockton and David W. Wilcox. "Econometric Models and the Monetary Policy Process." *Carnegie-Rochester Conference Series on Public Policy*, 1997, 47, pp. 1-37.

[44] Romer, Christina D. and David H. Romer. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 2000, 90, pp. 429-457.

[45] Romer, Christina D. and David H. Romer. "Choosing the Federal Reserve Chair: Lessons from History." *Journal of Economic Perspectives*, 2004, 18, pp. 129-162.

[46] Ruge-Murcia, Francisco J. "Uncovering Financial Markets' Beliefs about Inflation Targets." *Journal of Applied Econometrics*, 1999, 15, pp. 483-512.

[47] Ruge-Murcia, Francisco J. "Inflation Targeting Under Asymmetric Preferences." *Journal of Money, Credit and Banking*, 2003, 35, pp. 763-785.

[48] Sims, Christopher A. "The Role of Models and Probabilities in the Monetary Policy Process." *Brookings Papers on Economic Activity,* 2002, 2, pp. 1-62.

[49] Stock, James H., James H. Wright and Motohiro Yogo. "A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments." *Journal of Business and Economic Statistics*, 2002, 20, pp. 518-529.

[50] Svensson, Lars E.O. "Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets." *European Economic Review*, 1997, 41, pp. 1111-1146.

[51] Weber, E.U. "From Subjective Probabilities to decision weights: The Effect of Asymmetric Loss functions on the Evaluation of Uncertain Outcomes and Events." *Psychological Bulletin*, 1994, 115, pp. 228-242.

[52] Zarnowitz, Victor and Phillip Braun. "Twenty-two Years of the NBER-ASA Quarterly Economic Outlook Surveys: Aspects and Comparisons of Forecasting Performance." *NBER working paper 3965*, 1992.

[53] Zellner, Arnold. "Biased Predictors, Rationality and the Evaluation of Forecasts." *Economics Letters*, 1986, 21, pp. 45-48.

Appendix A. Mathematical Derivations

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

$$L\left(e_{t+h,t}\right)=\begin{cases} aL\left(e_{t+h,t}\right) & e_{t+h,t}>0 \\ 0 & e_{t+h,t}=0 \\ bL\left(e_{t+h,t}\right) & e_{t+h,t}<0 \end{cases} \tag{21}$$

where $a,b>0$. If $L\left(e_{t+h,t}\right)=\left|e_{t+h,t}\right|^{p}$ then this is the family of asymmetric functions defined in Elliott, Komunjer, and Timmermann (2005, 2006). 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:

$$L\left(e_{t+h,t}\right)=(a+b)\left[\phi+(1-2\phi)\,1_{\left(e_{t+h,t}<0\right)}\right]\left|e_{t+h,t}\right|^{2}, \tag{22}$$

where $1_{\left(e_{t+h,t}<0\right)}$ 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 (13).

The interpretation of the asymmetry parameter is as follows: $\phi=0.5$ gives $a=b$ so that it corresponds to symmetry, further, after some algebra one can get: $\frac{a}{b}=\frac{\phi}{1-\phi}$ and, 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 time as costly).

To obtain orthogonality condition (14) one needs to solve the following problem:

$$\min_{f_{t+h,t}} E\left[L\left(e_{t+h,t}\right)|\Omega_{t}\right], \tag{23}$$

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

$$\frac{\partial}{\partial f_{t+h,t}}E\left[\left[\phi+(1-2\phi)\,1_{\left(e^{*}_{t+h,t}<0\right)}\right]\left|e^{*}_{t+h,t}\right|^{2}|\Omega_{t}\right]=0, \tag{24}$$

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:

$$E_{t}\left[\begin{array}{c} -2\phi\left(1-(2)\,1_{\left(e^{*}_{t+h,t}<0\right)}\right)\left|e^{*}_{t+h,t}\right|+ \\ (1-2\phi)\,\frac{\partial}{\partial f_{t+h,t}}1_{\left(e^{*}_{t+h,t}<0\right)}\left|e^{*}_{t+h,t}\right|^{2}- \\ 2\,(1-2\phi)\,1_{\left(e^{*}_{t+h,t}<0\right)}\left(1-(2)\,1_{\left(e^{*}_{t+h,t}<0\right)}\right)\left|e^{*}_{t+h,t}\right| \end{array}\right]=0, \tag{25}$$

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

$$E_t \left[\begin{array}{c} 2\left(1_{\left(e^*_{t+h,t}<0\right)} - \phi\right) \left|e^*_{t+h,t}\right| - \\ (1-2\phi) \left|e^*_{t+h,t}\right|^2 \delta\left(e^*_{t+h,t}\right) \end{array} \right] = 0, \tag{26}$$

which can be further simplified to:

$$E_t \left[\left(1_{\left(e^*_{t+h,t}<0\right)} - \phi\right) \left|e^*_{t+h,t}\right| \right] = 0. \tag{27}$$

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

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

$$E \left[v_t \left(1_{\left(e^*_{t+h,t}<0\right)} - \phi\right) \left|e^*_{t+h,t}\right| \right] = 0. \tag{28}$$

Expression (14) is the same as this last expression, except that the following algebraic change is applied to (27):

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

Appendix B. Tables

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

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

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	0.16*	0.35**	0.00	68:4–98:2	119
	(0.09)	(0.08)	(0.10)			
1	-0.08	0.38**	-0.03	0.00	69:1–98:2	118
	(0.11)	(0.10)	(0.18)			
2	-0.05	0.44**	-0.05	0.00	69:3–98:2	116
	(0.15)	(0.17)	(0.19)			
3	-0.40**	0.21**	0.04	0.00	74:3–98:2	96
	(0.16)	(0.09)	(0.16)			
4	-0.36	0.10	-0.27	0.00	75:3–98:2	92
	(0.23)	(0.15)	(0.18)			

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: $e^F_{t+h,t} = \alpha + \gamma e^F_{t-1,t-h-1} + \omega^C(f^C_{t+h,t} - f^F_{t+h,t}) + \varepsilon_{t+h}$

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

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: $e^F_{t+h,t} = \alpha + \gamma e^F_{t-1,t-h-1} + \omega^C(f^C_{t+h,t} - f^F_{t+h,t}) + \varepsilon_{t+h}$

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

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**	0.62**
	(0.11)	(0.06)
1	-0.39**	0.69**
	(0.11)	(0.06)
2	-0.44**	0.72**
	(0.11)	(0.05)
3	-0.48**	0.74**
	(0.12)	(0.06)
4	-0.52**	0.76**
	(0.15)	(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**	-	-	0.45	0.50
	(0.11)				
C + lagged error	0.16**	1.94	0.86	0.35	0.85
	(0.07)				
C + SPF median	0.21**	2.06	0.84	1.07	0.30
	(0.09)				
C + SPF inqr	0.24**	5.14	0.40	1.14	0.28
	(0.09)				
		Since–Volcker			
Constant	0.80**	-	-	2.06	0.72
	(0.05)				
C + lagged error	0.90**	8.63	0.28	7.29	0.50
	(0.02)				
C + SPF median	0.82**	9.26	0.23	4.12	0.84
	(0.02)				
C + SPF inqr	0.87**	9.21	0.24	3.11	0.92
	(0.02)				

Source: Data from the Federal Reserve Bank of Philadelphia.

Notes: System GMM estimates imposing the restriction than 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 C. Figures

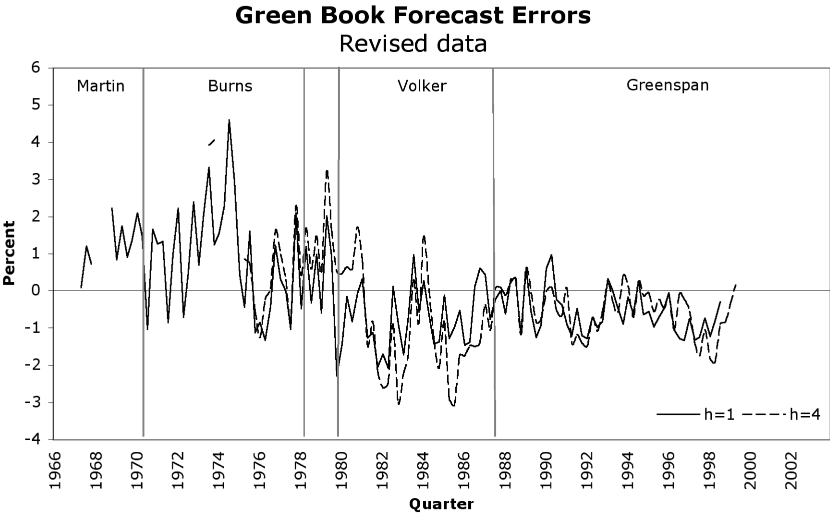

Figure 1: One- and Four-step-ahead Green Book Forecasts Errors with Revised Data

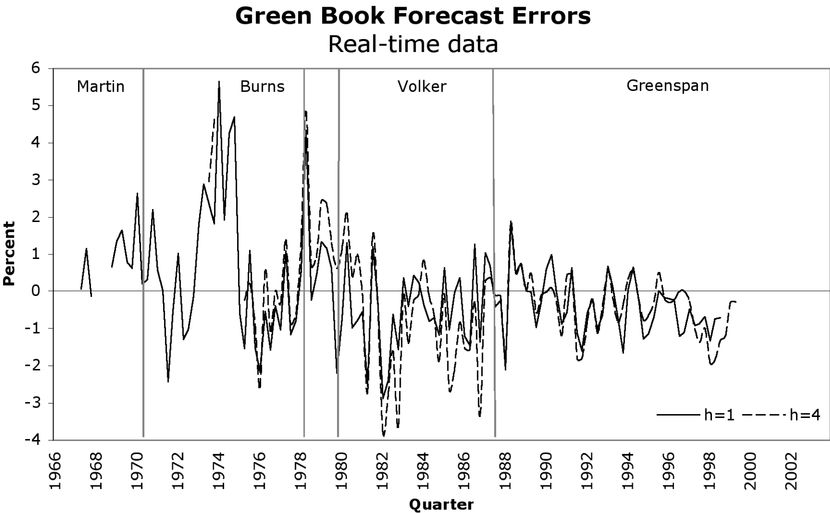

Figure 2: One- and Four-step-ahead Green Book Forecasts Errors with Real-time Data

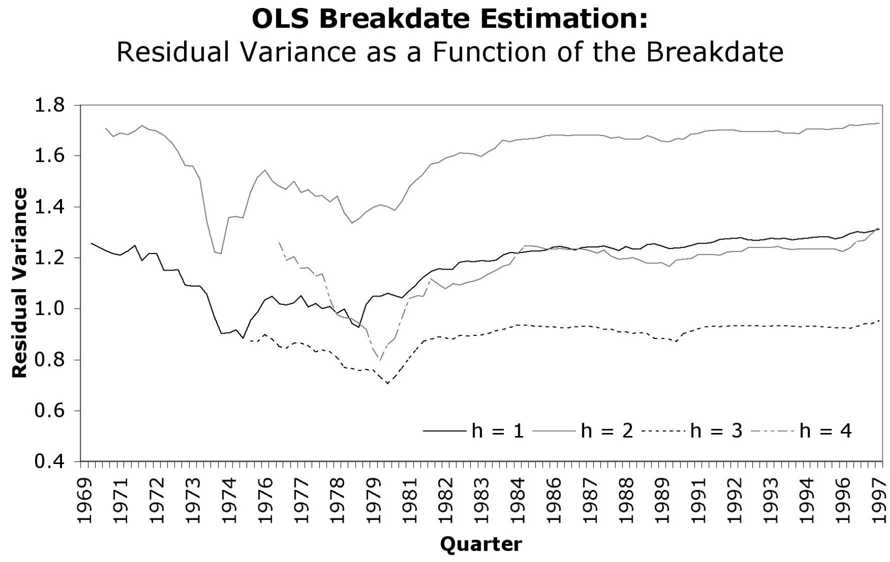

Figure 3: OLS Breakdate Estimation (single break)

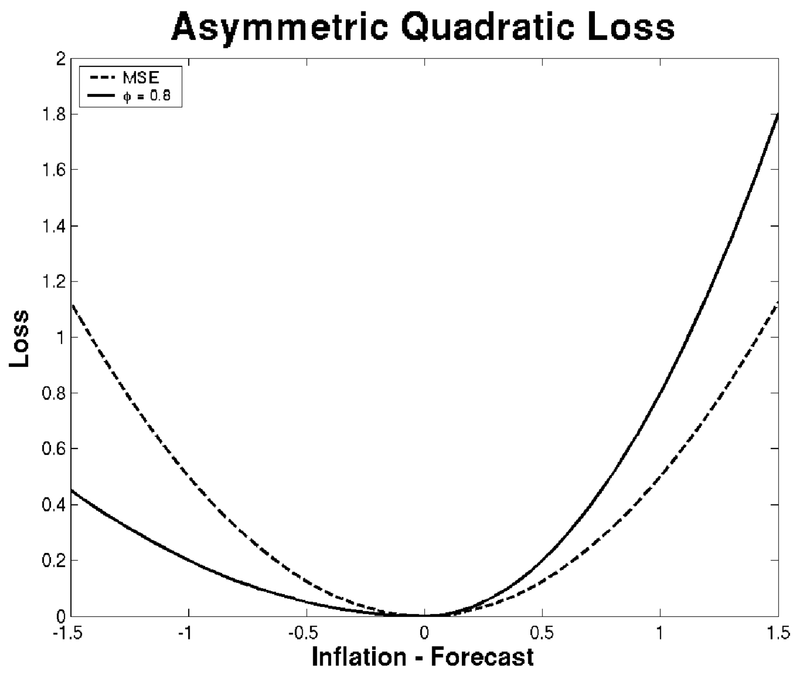

Figure 4: Symmetric and Asymmetric Quadratic Loss in Forecast Error Space

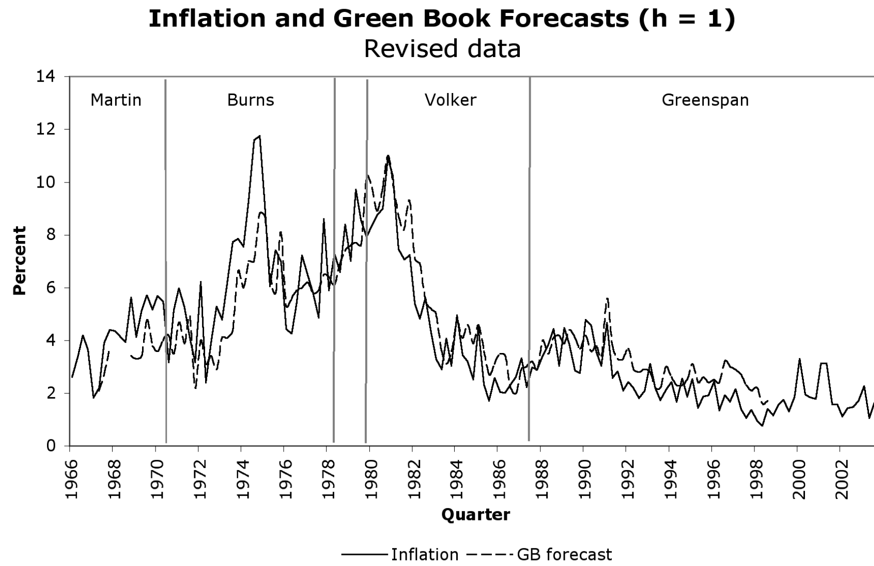

Figure 5: Inflation and One-step-ahead Green Book Forecasts with Revised Data

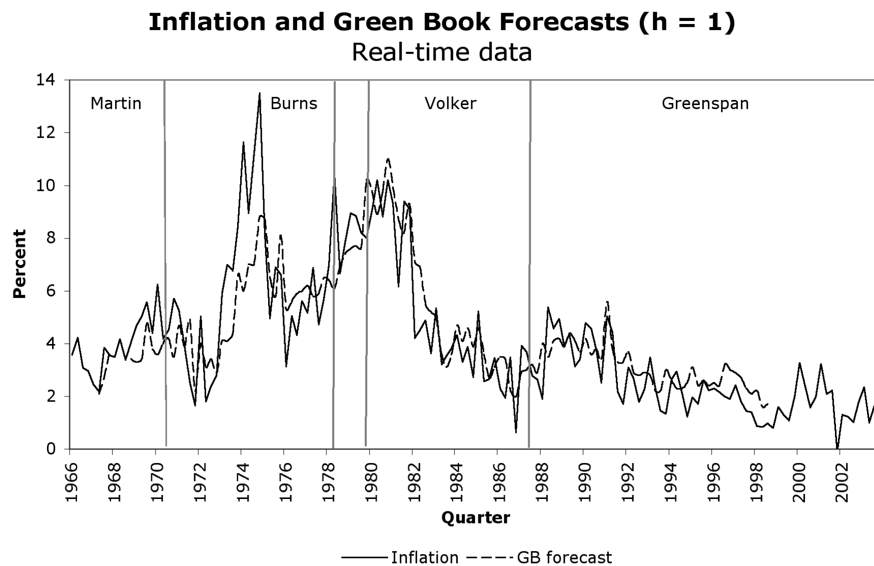

Figure 6: Inflation and One-step-ahead Green Book Forecasts with Revised Data