

Banco de México
Documentos de Investigación

Banco de México
Working Papers

N° 2010-11

**Forecast Revisions of Mexican Inflation and GDP
Growth**

Carlos Capistrán
Banco de México

Gabriel López-Moctezuma
Banco de México

October 2010

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 of the authors and do not necessarily reflect those of Banco de México.

Forecast Revisions of Mexican Inflation and GDP Growth*

Carlos Capistrán[†]
Banco de México

Gabriel López-Moctezuma[‡]
Banco de México

Abstract: We analyze forecasts of inflation and GDP growth contained in Banco de México's Survey of Professional Forecasters for the period 1995-2009. The forecasts are for the current and the following year, comprising an unbalanced three-dimensional panel with multiple individual forecasters, target years, and forecast horizons. The fixed-event nature of the forecasts enables us to examine efficiency by looking at the revision process. The panel structure allows us to control for aggregate shocks and to construct a measure of the news that impacted expectations in the period under study. The results suggest that respondents seem to rely for longer than appears to be optimal on their previous forecasts, and that they do not seem to use past information in an efficient manner. In turn, this means there are areas of opportunity to improve the accuracy of the forecasts, for instance, by taking into account the positive autocorrelation found in forecast revisions.

Keywords: Evaluating forecasts; Inflation forecasting; Macroeconomic forecasting; Panel data; Surveys.

JEL Classification: C23; C53; C83; E37.

Resumen: En este documento se analizan los pronósticos de inflación y de crecimiento del PIB de la Encuesta sobre Expectativas de los Especialistas en Economía del Sector Privado que lleva a cabo el Banco de México. El periodo analizado va de 1995 a 2009. Los pronósticos son para el año actual y el año siguiente, y forman un panel de tres dimensiones con múltiples pronosticadores, años a pronosticar y horizontes de pronóstico. Al ser pronósticos de eventos fijos, es posible examinar la eficiencia de los pronósticos analizando su proceso de revisión. La estructura de panel permite controlar por choques agregados y construir una medida de las noticias que tuvieron algún impacto sobre las expectativas durante el periodo estudiado. Los resultados sugieren que los encuestados parecen mantener sus pronósticos previos por más tiempo del que parece ser el óptimo, y que no parecen utilizar la información de una forma eficiente. Ello implica la existencia de áreas de oportunidad para mejorar la precisión de los pronósticos, por ejemplo, tomando en cuenta la autocorrelación positiva que existe en las revisiones de los mismos.

Palabras Clave: Evaluación de pronósticos; Pronósticos de inflación; Pronósticos macroeconómicos; Datos de panel; Encuestas.

*We thank Daniel Chiquiar, Nikita Perevalov, José Gonzalo Rangel and seminar participants at Banco de México and the XIV CEMLA Research Economists Meetings for helpful comments. The opinions in this paper correspond to the authors and do not necessarily reflect the point of view of the Banco de México.

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

[‡]Dirección General de Investigación Económica. Email: glopez@banxico.org.mx.

1 Introduction

It has been a common practice in several central banks and other institutions to design and collect surveys containing professionals' forecasts of macroeconomic and financial variables. These surveys can be particularly helpful to understand the expectation formation mechanism of private agents. Equally important, they can serve as an input in a number of decision-making processes, for instance monetary policy. Given their importance, it is natural to ask about their accuracy, and efficiency.

In this paper we analyze the forecasts of inflation and GDP growth supplied by the individual respondents to the Survey of Professional Forecasters (SPF in what follows) conducted by the Banco de México each month since 1995.¹ These forecasts comprise a three-dimensional panel dataset, with the additional dimension arising from the collection of forecasts at several horizons (Davies and Lahiri, 1995). The forecasts that we use are focused on the end-year outcome in the same year in which the survey is collected, and in the following year. The forecasts are revised in response to new information from one survey to the next, and eventually form a sequence of 24 forecasts before the respective outcome is known. This type of expectations is commonly referred to as fixed-event forecast. Following Isiklar and Lahiri (2007) we use this real-time survey data to analyze the evolution of these fixed-event forecasts over various horizons.

We study the extent to which information is incorporated into these forecasts. If it is efficiently incorporated, forecast errors and revisions should not be predictable (Nordhaus, 1987). When the information incorporated is publicly available data that might have been employed to construct previous forecasts, the forecasts are said to satisfy the property of strong efficiency. If the information used are past forecast errors or revisions, then the forecasts are said to satisfy the property of weak efficiency. Notice that the latter is a necessary condition for the former, but it is not sufficient.

The fixed-event nature of the forecasts enables us to examine forecast efficiency by looking at forecast revisions. Rationality tests based on forecast revisions are attractive because they are not sensitive to the data generating process or to data-revisions. This is important for Mexican data given that Mexico adopted an inflation targeting regime in 2001, with inflation apparently changing from a very persistent process to a stationary one around that year (Chiquiar et al., 2007).² In the case of GDP, its measurement had a major revision in 2008,

¹The survey's name is: Encuesta sobre Expectativas de los Especialistas en Economía del Sector Privado. Banco de México is México's central bank, with webpage <http://www.banxico.org.mx>

²Mexico took a first step towards inflation targeting in 1999 and consolidated it in 2001. Starting in 2003, a long-term inflation objective for the CPI was set at 3% with a variability interval of plus/minus 1%. See Banco de México's Monetary Program for 2003.

in which the base year was changed from 1993 to 2003.³ Two other advantages of using the fixed-event forecast in the SPF are that we have forecasts for reasonable long-horizons (up to twenty-four months), and that these are revised on a monthly basis.⁴

The panel structure of the data makes it possible to separate forecast errors into macroeconomic aggregate shocks and forecaster's specific idiosyncratic errors (Davies and Lahiri, 1995). The aggregate shocks constitute a measure of the news that impacted GDP growth and inflation expectations in the period under study. Following Davies and Lahiri (1995) and Boero et al. (2008b), we calculate a measure of the volatility of the aggregate shocks for the Mexican case. In addition, in order to take full advantage of the panel structure of the data, we consider pooling fixed-event forecasts across events and over individual respondents to deliver more powerful tests of forecast efficiency (Keane and Runkle, 1990; Clements, 1997).

We find that SPF respondents seem to start with a fixed value for the initial forecast, around 3.7% for inflation (during the inflation targeting period) and around 3.8% for GDP growth. Compared to the event's realization, these values tend to under-predict inflation and over-predict GDP. The forecasters appear to start incorporating news into the forecasts around 12 and 16 months before the realization of the end-of-the-year inflation and real GDP growth, respectively. We also find that forecasters tend to rely for longer than appears to be optimal on their previous forecasts when predicting both, annual inflation and GDP growth. Further results indicate inefficiencies in the use of information about the past evolution of monthly inflation and monthly measures of economic activity. All these inefficiencies suggest clear areas of opportunity to significantly improve the accuracy of the forecasts. For instance, the positive autocorrelation found in forecast revisions could be used to predict subsequent revisions.

With respect to the aggregate shocks, we find high volatility for both inflation and GDP growth around the crisis of 1994-1995. For the latter we also find high volatility around the recent global financial crisis. An important reduction in inflation uncertainty is evident after Banco de México implemented an inflation targeting framework. There is also a cluster of positive aggregate news to inflation around that time. Finally, GDP growth uncertainty appears to have a cyclical component.

The paper proceeds as follows. Section 2 provides information about the SPF and the data employed and presents an analysis of the information content of the forecasts. Section

³The revision was announced in April of 2008. It included, among other things, increasing the number of activities considered from 362 to 750. See <http://www.inegi.org.mx>

⁴Despite these advantages, the bulk of the forecast evaluation literature has focused on fixed-horizon forecasts, which are a series of predictions for different events at a fixed-horizon. See for example Mincer and Zarnowitz (1969) for early work on fixed-horizon forecasts, and Capistrán and López-Moctezuma (2010) for an evaluation of fixed-horizon predictions from Banco de México's SPF.

3 derives the efficiency tests used in this paper. Section 4 presents the results of the tests for forecast efficiency. The estimates of the aggregate shocks and their volatility appear in Section 5. Finally, Section 6 discusses the results and concludes.

2 Data

2.1 Banco de México's Survey of Professional Forecasters

Banco de México has conducted the SPF on a monthly basis since September 1994. Nowadays, the SPF covers around 20 macroeconomic variables related to investment, production, labor markets, public finance and international trade. In addition, the survey asks the professional forecasters for their views on some qualitative aspects of the Mexican economic environment.

The number of forecasters in the survey has varied over the years, although since the late 90s there have been approximately 30 regular respondents in each survey.⁵ The specialists who participate in the survey come mainly from commercial banks and other financial institutions (57%), followed by consulting firms (29%) and industrial and academic institutions (14%) in a smaller proportion. Their forecasts are gathered by mail on the second half of each month and the un-weighted mean (consensus forecast) is published monthly by the Banco de México in a detailed report that contains the evolution of these expectations.

In this paper we analyze the forecasts collected in the SPF that focus on predicting the current and following year of annual CPI inflation and the average GDP growth over the year, for a sample that begins as early as January (for inflation) and March (for GDP growth) of 1995 to December 2009. Ideally, from the frequency of the survey and the structure of the data we should be able to extract a sequence of 12 fixed-event forecasts for each forecaster for the target year 1995, 24 fixed-event forecasts for the target years 1996 to 2009 and 12 forecasts for the year 2010, with $h = 23, 22, \dots, 1, 0$; where horizon 23 corresponds to the forecast made in January of a given year for the outcome of the following year, and horizon 0 corresponds to the forecast released in December of a given year for the outcome of the current year. However, in practice, the sequences of individual fixed event forecasts take the form of an unbalanced panel as individual forecasters frequently enter and exit the survey (see Capistrán and Timmermann (2009b)).

To avoid the complications caused by long gaps in the data, our main analysis refers to respondents who reported more than 50% of the time (90 responses out of 180 survey issues since January 1995). With this threshold we are able to use information on 41 forecasters

⁵For the first year, there were around 15 respondents per month. The number of respondents rose steadily until 1998. From January 1998 to December 2009, the average number of respondents per month is 31.1.

(of a total of 78 forecasters) for both inflation and GDP growth forecasts.⁶

2.2 Descriptive Statistics

As a first step in the exploratory data analysis of the forecasting process in Mexican inflation and GDP growth, Figures 1 to 4 present a graphical summary of the cross-section mean and standard deviation of the panel forecasts and of the forecast revisions for each year under study, together with the available actual outcomes. In these figures, the sample mean is the common choice of a “consensus” forecast and the standard deviation across forecasters is a measure of the disagreement among respondents and could be viewed as a proxy of the uncertainty surrounding panel members responses about future outcomes (Harvey et al., 2001). It might be expected that variability in panel members’ forecasts decreases as they approach the outcome of interest. This might happen because as new information becomes available, each forecaster should provide more accurate estimates (Isiklar and Lahiri, 2007).

The general picture is that the evolution of the mean forecast seems to converge towards the actual outcome which is not surprising since for shorter horizons, much of the actual outcome has already been published, although the speed of convergence seems to vary across years and across variables being forecasted.

Similar to the majority of developed countries examined in Isiklar and Lahiri (2007), the initial average forecasts produced by the respondents of the SPF, that is the 24 months-ahead forecasts, seem to start from very similar initial points. This happens regardless of their final expectations and thus, of the actual value of the variables of interest, as can be seen in Table 1. Isiklar and Lahiri (2007) conjectured that the long term forecasts may be regarded as the unconditional mean of the time series of interest. Lahiri and Sheng (2008) find that almost all the variance of the initial forecasts can be explained by the prior beliefs of the forecasters. In the case of inflation forecasts, the initial expectations seem to change considerably from the years 1999 to 2002, which coincides with a strong disinflation effort and the transition towards an inflation targeting regime. From 2003 onwards, the longer-run expectations are located around 3.7%. This number is inside the variability interval announced by Banco de México around its inflation target, 3% plus/minus 1%, although it is close to its upper limit. Comparing this initial forecast with the average inflation at the end of the year for the period 2003-2009, 4.3%, there appears to be a systematic bias of the initial forecasts to under-predict inflation (the mean error for this period is 0.6%). In the case of GDP growth forecasts, the third column of Table 1 shows that initial forecasts start from a range between 3.3% and 4.6% for the whole period 2000-2009, with a mean value of 3.8%. Given that the

⁶Our results are robust to the use of a threshold of 30%, which leaves 44 forecasters.

average GDP growth in the period is 1.9%, the initial predictions of the forecasters seem to over-predict GDP growth, on average, by 1.9%.

Another relevant feature of the evidence presented above is that the sequences of consensus forecasts seem to move steadily up or down towards the actual outcomes. As we will see later in the paper, the theory of optimal forecasting predicts that, under certain conditions, these sequences should, if optimal, look more like a white-noise process, as every period optimal forecast revisions should be unpredictable. In the next section we develop the tools necessary to test if this pattern is compatible with the efficient use of information.

Finally, from Figures 1 to 4 it is also possible to see that for the majority of the events analyzed in the period 1999-2009 (in which the SPF includes 24 months-ahead forecasts), approximately for the first 6 to 8 months the cross-section average forecast does not seem to change considerably, which seems to indicate that the initial months do not seem to convey relevant news as to persuade forecasters to revise their predictions systematically. This evidence is also consistent with that reported by Isiklar and Lahiri (2007) for developed countries. Furthermore, Figure 5 shows the behavior of the Diebold and Kilian (2001) statistic to measure the predictability of a variable. This measure is $pred_{h,23} = 100 \left(1 - \frac{MSE_h}{MSE_{23}} \right)$, where MSE_h is the Mean Squared Error of the forecasts at horizon h (the MSE is taken across events). This statistic shows the evolution of the information content of the forecasts as measured by the decrease in MSE over that of the 24-month ahead forecast (the forecast at horizon 23, in our notation). For inflation, we see that predictability first drops, increases for about four months, then remains about the same for another six, and finally steadily increases for the next eleven months. The behavior of the forecasts for longer horizons is hard to rationalize, and most likely implies inefficient use of information, since the forecasts for horizon 23 appear to have the same or better predictive ability than forecasts up to horizon 12.⁷ The behavior of the shorter-horizon forecasts is consistent with rapid gains in predictability as the slope of the $pred_{h,23}$ measure is very steep. For GDP growth, we see that each additional month increases the information content of the forecasts, as the MSE improves steadily over the previous month. There is a change in the slope starting in the horizon 16, which indicates large gains in predictability starting about 4 quarters ahead of the realization of the event.

⁷See Capistrán (2007) for more on the multi-horizon properties of rational forecasts and a test of rationality based on the non-increasing property of MSEs as the forecast horizon decreases.

3 Efficiency Tests

Efficiency tests for fixed-event forecasts using forecast revisions were first introduced by Nordhaus (1987) who discusses the efficiency of forecast revisions for a single terminal event. Later on, Davies and Lahiri (1995) considered exploiting the panel data structure on survey forecasts under an econometric framework that renders possible the decomposition of forecasts errors into macroeconomic aggregate shocks and forecaster specific idiosyncratic errors.⁸ Forecast revisions of professional forecasters in developed economies have been extensively analyzed under this framework.⁹ However, in the case of developing economies, to our knowledge, there is no previous study analyzing the efficiency of fixed-event expectations in survey data.

First, we derive the properties of optimal fixed event forecasts under a general loss function. Then, we use the quadratic loss function to derive explicit tests that involve the forecast revisions. To facilitate the exposition, the optimal properties are derived for a representative forecaster. Later we show the advantages of having a panel of forecasters.

3.1 Optimal forecasts

We are interested in forecasting the outcome of an event that is going to be realized at period τ , y_τ , with information up to h periods before. The information set is denoted $\mathbf{I}_{\tau-h}$, and contains at least the realizations of the variable to be forecasted and possibly lagged values of other variables, up to period $\tau - h$.

The optimal forecast computed at period $\tau - h$ conditional on $\mathbf{I}_{\tau-h}$ is defined as

$$f_{\tau,h}^* \equiv \arg \min_{f_{\tau,h}} E [L(y_\tau - f_{\tau,h}) \mid \mathbf{I}_{\tau-h}],$$

where $L(\cdot)$ is the loss function, and $f_{\tau,h}$ is the forecast for τ made at $\tau - h$. The first order condition is

$$E [L'(y_\tau - f_{\tau,h}^*) \mid \mathbf{I}_{\tau-h}] = 0, \tag{1}$$

where $L'(\cdot)$ denotes the derivative of the loss function with respect to the predictor, $f_{\tau,h}$.¹⁰ Following Granger (1999) and Patton and Timmermann (2007), this derivative is called the generalized error. It gives the change in total loss resulting from a one-unit change in the

⁸See Davies et al. (2010) for a recent account of the literature analyzing three-dimensional panels of forecasts.

⁹For US data, see Nordhaus (1987) and Swindler and Ketcher (1990). Bakhshi et al. (2005), Clements (1997), Harvey et al. (2001), and Boero et al. (2008a, 2008b) analyze UK data. Dovern and Weisser (2008) use data for the G7 countries.

¹⁰We have assumed that integration and differentiation can be interchanged.

forecast. Under certain regularity conditions, the first order condition will be sufficient to get the optimal forecast.

The optimal forecast made at $\tau - (h + 1)$, that is, one period before h , has to satisfy

$$E [L'(y_\tau - f_{\tau,h+1}^*) | \mathbf{I}_{\tau-(h+1)}] = 0. \quad (2)$$

Assuming that the information set is a filtration, the first difference of the generalized error (times -1 , a simple normalization) satisfies

$$E [L'(y_\tau - f_{\tau,h+1}^*) - L'(y_\tau - f_{\tau,h}^*) | \mathbf{I}_{\tau-h}] = 0. \quad (3)$$

By the Law of Iterated Expectations, for any finite function of a random variable belonging to the information set, $\zeta_{\tau-h} \subset \mathbf{I}_{\tau-h}$, the first difference of the generalized error satisfies the orthogonality condition

$$E [(L'(y_\tau - f_{\tau,h+1}^*) - L'(y_\tau - f_{\tau,h}^*))\zeta_{\tau-h}] = \mathbf{0}. \quad (4)$$

Stinchcombe and White (1998) refer to ζ as the test function. For a given test function, we will construct a test that exploits the orthogonality condition.

3.2 Quadratic loss

3.2.1 A representative forecaster

Under quadratic loss, $L(y_\tau - f_{\tau,h}) = \frac{1}{2}(y_\tau - f_{\tau,h})^2$, the generalized error $L'(y_\tau - f_{\tau,h})$ is identical to the forecasting error $y_\tau - f_{\tau,h}$ and hence, equation (1) becomes

$$E [(y_\tau - f_{\tau,h}^*) | \mathbf{I}_{\tau-h}] = 0, \quad (5)$$

and we obtain the traditional result that under quadratic loss the optimal forecast is the expected value of the variable of interest conditional on the forecaster's information set (Granger and Newbold, 1986). Under this loss function, equation (4) becomes

$$E [(f_{\tau,h}^* - f_{\tau,h+1}^*) \zeta_{\tau-h}] = \mathbf{0}, \quad (6)$$

therefore, under optimality, the forecast revision, $v_{\tau,h} \equiv f_{\tau,h} - f_{\tau,h+1}$, satisfies

$$E [(v_{\tau,h}^*) \zeta_{\tau-h}] = \mathbf{0}. \quad (7)$$

From equation (7) we can obtain the usual properties of optimal revisions (as in Nordhaus (1987)):

1. $E [v_{\tau,h}^*] = 0$, which is obtained with $\zeta_{\tau-h} = 1$.
2. $E [(v_{\tau,h}^*)(v_{\tau,h+j}^*)] = 0$, for all $j > 1$, which is obtained when $\zeta_{\tau-h} = v_{\tau,h+j}^*$.
3. $E [(v_{\tau,h}^*)(x_{\tau-(h+1)} - x_{\tau-(h+2)})] = 0$, which is obtained when $\zeta_{\tau-h} = x_{\tau-(h+1)} - x_{\tau-(h+2)}$.
In this case $x_{\tau-(h+1)}$ denotes the most recently observed value of a variable included in the forecaster's information set, thus, $x_{\tau-(h+1)} - x_{\tau-(h+2)}$ approximates the period-by-period news in this variable.

A series of revisions that satisfy the first two properties are said to be weakly efficient. Property 1 is known as the unbiasedness condition, and it implies that there should not be a systematic bias in the revisions. Property 2 implies that the revisions should be white noise. Hence, if the revisions are correlated, it would be evidence of inefficient use of information under the assumption of quadratic loss. Property 3 is associated with the concept of strong efficiency as the forecaster's information set includes additional variables other than the past forecast revisions. From the point of view of the researcher, testing this property has the disadvantage, with respect to the first two, that it requires the extra assumption that the variable $x_{\tau-(h+1)}$ actually belonged to the forecaster's information set at the time the forecast was computed.

Property 2, has been disputed by Davies and Lahiri (1995), who postulate that efficient forecast revision sequences should behave as a first-order moving average, with a negative first-order autocorrelation coefficient, but with all the other autocorrelations equal to zero. However, under our derivation we obtain Nordhaus' (1987) result that optimal forecast revisions (under quadratic loss), should behave as a zero-order moving average.

3.2.2 A panel of N forecasters

In a panel data setting with N forecasters, we can think of the optimal revision for each forecaster as a component term that can be divided into an aggregate shock, $u_{\tau,h}$ (common across forecasters), and an idiosyncratic (white-noise) error, $\varepsilon_{i,\tau,h}$ (specific to each forecaster). Thus the optimal forecast revision can be written as

$$v_{i,\tau,h}^* = u_{\tau,h} + \varepsilon_{i,\tau,h}, \tag{8}$$

as long as $E[u_{\tau,h} + \varepsilon_{i,\tau,h}] = 0$, so that property 1 is still satisfied. In this context, an estimate of the aggregate shock at each time period is readily available as

$$\hat{u}_{\tau,h} = N^{-1} \sum_{i=1}^N v_{i,\tau,h}^* \quad (9)$$

Hence, an estimate of the systematic portion of the forecast errors can be extracted from equation (9) assuming: *i*) a homogenous quadratic loss across forecasters; and, *ii*) that each forecaster included to extract the aggregate shock ought to be rational.

The error structure of equation (8) implies a three-way panel with multiple individuals, forecast horizons, and events of interest. Davies and Lahiri (1995) have an expression similar to (8) but they add a bias term and a $MA(1)$ structure in the revision process. These additional terms do not appear in our expression because we have derived it under optimality and using a quadratic loss.¹¹

3.3 Tests based on forecast revisions

Following the three-dimensional panel structure of our data, we base our analysis of the properties of efficient forecast revisions, under quadratic loss, on a regression framework that facilitates the pooling of the forecast revision sequences across different target years and over individual respondents. The general regression for testing efficiency with the forecast revision as the dependent variable can be written as

$$\begin{aligned} v_{i,\tau,h} &= \zeta'_{\tau-h} \boldsymbol{\alpha} + \varepsilon_{i,\tau,h}, \\ H_0 &: \boldsymbol{\alpha} = \mathbf{0} \\ H_a &: \boldsymbol{\alpha} \neq \mathbf{0}, \end{aligned} \quad (10)$$

where $\zeta_{\tau-h}$ is a $k \times 1$ vector of variables contained in the information set $\mathbf{I}_{\tau-h}$, and the null and alternative hypotheses involve k parameters each. Forecasters are indexed by $i = 1, \dots, N$, target years by $\tau = 1, \dots, T$, and forecast horizons by $h = 1, \dots, H$. For estimation, the data is sorted first by forecaster, then by target year, and lastly by forecast horizon. Under the joint hypothesis of forecast efficiency and quadratic loss, the optimal revision is uncorrelated with the variables contained in each forecaster's information set, $\zeta_{\tau-h}$, and hence all the

¹¹In particular, Davies and Lahiri (1995) include a constant bias term. They suggest that it can account for the presence of systematic biases that can emerge even under optimality, for instance under an asymmetric loss function. We do not include it because we have derived the properties of optimal revisions under a symmetric (quadratic) loss. In any case, under an asymmetric loss, the bias would typically be time-varying (e.g., Clements (1997)).

components of α should equal zero. Notice that we are assuming the same quadratic loss function for all the forecasters (i.e., $\alpha_i = \alpha \forall i$).

Using equation (10) as reference, we can perform a test of weak efficiency incorporating properties 1 and 2 of optimal revisions in a single regression with $\zeta_{\tau-h} = [1, v_{\tau, h+j}]$. In this sense, we are testing both unbiasedness and lack of autocorrelation in the forecast revision process. The regression is

$$v_{i,\tau,h} = \alpha_0 + \sum_{j=1}^p \alpha_j v_{i,\tau,h+j} + \epsilon_{1i,\tau,h}, \quad (11)$$

$$H_0 : \alpha_0 = \alpha_1 = \dots = \alpha_p = 0.$$

In addition, as can be seen from property 3, optimal revisions incorporate publicly available information in an efficient manner and therefore, they are revised exclusively in response to unexpected shocks. Thus, a test of strong efficiency can be performed by regressing the forecast revision on candidate variables contained in each respondent's information set. Since the information set might differ across individuals and is not observable, any variable employed for evaluation purposes is rather arbitrary (Thomas, 1999). In order to avoid using variables not available to the respondents at the time the forecasts were made, we use the monthly observed changes in the target variables as independent variables. That is, in the case of inflation, we use monthly inflation as the independent variable. For GDP growth, as actual values of GDP are not published on a monthly basis, we employ the indicator of global economic activity (IGAE), which is a monthly timely indicator of economic activity. The strong efficiency regression is

$$v_{i,\tau,h} = \alpha'_1 (x_{\tau-(h+1)} - x_{\tau-(h+2)}) + \epsilon_{2i,\tau,h}, \quad (12)$$

$$H_0 : \alpha'_1 = 0$$

$$H_a : \alpha'_1 \neq 0.$$

Notice that we can express the revisions in a vector, Υ , as follows

$$\Upsilon' = (v_{1,1,H-1}, \dots, v_{1,1,1}, v_{1,2,H-1}, \dots, v_{1,2,1}, \dots, v_{1,T,H-1}, \dots, v_{1,T,1}, v_{2,1,H-1}, \dots, v_{N,T,H-1}).$$

3.4 Econometric considerations

We can provide consistent estimates of the relevant coefficients of equations (11) and (12) for each individual respondent and for the pooled data by ordinary least squares (OLS). However, under the null hypothesis of efficiency, the error terms $\epsilon_{1i,\tau,h}$ and $\epsilon_{2i,\tau,h}$ will have a

special structure. Following Isiklar et al. (2006), since the forecasts span up to a 24-month period with monthly revisions, the structure of the variance-covariance matrix of the error vector, $\mathbf{\Omega}$, should take into account: (i) different error variances across forecasters; and (ii) correlations of contemporaneous revisions across individuals for the same target year. In addition, since in any given month forecasters will revise their predictions for two consecutive years: (iii) contemporaneous revisions for consecutive target years for each forecaster; and, (iv) across forecasters.

The elements of $\mathbf{\Omega}$ can then be estimated from the OLS residuals $\hat{\epsilon}_{1i,\tau,h}$ (or $\hat{\epsilon}_{2i,\tau,h}$ in its case) by subtracting means and averaging across horizons and target years. We assume that all other covariances among $\epsilon_{1i,\tau,h}$ and $\epsilon_{2i,\tau,h}$ are zero. Using the general efficiency regression (10):

- For the error variance of forecaster i , σ_i^2 , we estimate:

$$\hat{\sigma}_i^2 = \frac{1}{TH} \sum_{h=1}^H \sum_{t=1}^T \hat{\epsilon}_{i,\tau,h}^2.$$

- The covariances across respondents, γ_{ij} , are computed as:

$$\hat{\gamma}_{ij} = \frac{1}{TH} \sum_{h=1}^H \sum_{t=1}^T \hat{\epsilon}_{i,\tau,h} \hat{\epsilon}_{j,\tau,h}, \forall i \neq j.$$

- Contemporaneous covariances for consecutive target years for each forecaster are estimated as:

$$\hat{\omega}_i = \frac{1}{(T-1)\tilde{H}} \sum_{h=1}^{\tilde{H}} \sum_{t=1}^{T-1} \hat{\epsilon}_{i,\tau,h} \hat{\epsilon}_{i,\tau+1,h+12}, \text{ where } \tilde{H} < 12.$$

- The contemporaneous covariances for consecutive target years across forecasters, s_{ij} , are computed as follows:

$$\hat{s}_{ij} = \frac{1}{2(T-1)\tilde{H}} \sum_{h=1}^{\tilde{H}} \sum_{t=1}^{T-1} \hat{\epsilon}_{i,\tau,h} \hat{\epsilon}_{j,\tau+1,h+12} + \hat{\epsilon}_{j,\tau,h} \hat{\epsilon}_{i,\tau+1,h+12}, \text{ where } \tilde{H} < 12.$$

The number of diagonal elements estimated, $\hat{\sigma}_i^2$, and of individual covariances for the same target, $\hat{\omega}_i$, equal the number of forecasters, N . The number of estimated covariances across forecasters for the same target, $\hat{\gamma}_{ij}$, and for consecutive targets, \hat{s}_{ij} , is $(N \times (N - 1))/2$.

In order to obtain consistent standard errors for the coefficients of equations (11) and (12) we calculate the general method of moments (GMM) covariance estimator given by

$(\zeta'\zeta)^{-1}\zeta'\widehat{\Omega}\zeta(\zeta'\zeta)^{-1}$. The consistent GMM standard errors are then obtained as the squared root of the diagonal elements of this covariance estimator.

4 Empirical Results from Efficiency Tests

In this section we present the empirical results of the weak and strong-efficiency tests developed above. We focus on the pooled results containing the information of the panel data structure and, when possible, discuss separate individual results obtained by using the relevant portion of the estimated variance-covariance matrix. These results should be taken with caution because of the low power of the tests when applied to small samples and because the gaps in the individual time series are particularly problematic to estimate autocorrelations in the forecast revision process.

4.1 Weak efficiency

The results of weak efficiency tests applied to the pooled forecast revisions are shown in Table 2 along with consistent standard errors obtained from the GMM covariance estimator. The second column presents the results for the weak-efficiency test with $p = 2$. That is, we are testing if forecast revisions at period $\tau - h$ are correlated with past revisions at periods $\tau - (h + 1)$ and $\tau - (h + 2)$.

First, these results indicate that for annual inflation forecast revisions, SPF's respondents are, on average, unbiased, which is consistent with property 1 of an optimal forecast revision. However, given the initial bias to under-predict of the long-horizon forecasts documented in subsection 2.2, the lack of a systematic pattern in the revision process implies that the term structure of inflation forecast is likely to remain biased.

Second, the results for annual inflation show that the first-order autocorrelation coefficient in regression (11) is slightly positive but not statistically different from zero at conventional significance levels, and that the second-order autocorrelation coefficient is significantly different from zero at the 1% level, suggesting an inefficient use of available information. The fact that α_2 is estimated to be positive can be interpreted as evidence of over-smoothing. That is, forecasters seem to maintain their initial forecasts for too long, which could be related to "anchoring", the common human tendency to rely too heavily on one piece of information when making decisions (Tversky and Kahneman, 1974). In addition, as can be seen from the second column of Table 2, even under Davies and Lahiri's (1995) framework, the inflation forecasts appear to be inefficient because the second-order autocorrelation is statistically positive.

Consistent with the panel regression, the individual results for annual inflation reject the weak-efficiency hypothesis for 59% of the individual respondents. In particular, although only 6 of the 41 forecasters analyzed seemed to bias their inflation forecasts, 54% of them do not appear to incorporate at least one of the two last forecast revisions in their expectation formation process.

The panel results for GDP growth forecasts revisions follow the same pattern as those for inflation. As in the previous case, we are able to reject the weak-efficiency hypothesis as well; although in this case, we also find a negative and significant constant bias, which imply that, on average, these forecasters were systematically decreasing their forecasts for annual GDP growth. This is consistent with the result documented in subsection 2.2. that the initial, long-horizon, forecasts tend to over-predict GDP growth. Hence, the systematic bias in the revisions may be an attempt to reduce the bias in shorter-horizon forecasts. At the individual level, we find that 76% of the forecasters examined seem to have failed the weak-efficiency tests, and thus have displayed inefficiencies in incorporating relevant information from their own past GDP growth forecasts. Particularly, 12% of the forecasters under study showed biased GDP growth revisions and 76% of them reveal the presence of autocorrelation in their forecast revision sequences.

4.2 Strong efficiency

The results of testing the strong efficiency property with consistent GMM standard errors for both, inflation and GDP growth, are presented in Table 3. We include a constant in the regression to allow for the presence of systematic bias in the forecast revision process. We also include a $MA(1)$ term in the error process to capture any presence of autocorrelation that might bias the estimated coefficients.

The aggregate results for the inflation forecasts reject the strong efficiency hypothesis at the 1% significance level. The positive sign of the estimated coefficient indicates that increases in past monthly inflation are associated with a positive monthly revision in the annual forecasts. In particular, a 1% increase in monthly inflation news induces an upward change in the monthly revision process of 0.27%. This evidence suggests, in principle, that annual inflation forecasts could be improved by using information contained in the past evolution of monthly inflation. At the individual level, the previous result is less apparent as there are only 11 (out of 41) forecasters for whom strong efficiency can be rejected at least at the 10% level.

In contrast with the pooled results for annual inflation forecasts, strong efficiency cannot be rejected at conventional significance levels for GDP growth forecasts according to

the estimated coefficients presented in the third column of Table 3.¹² At the individual level, we find evidence to reject the hypothesis of strong efficiency for about 7% of the forecasters. These individual forecasters could improve their monthly predictions by following more closely the trajectory of timely indicators of economic activity.

Although at the aggregate level the forecasters analyzed seem to efficiently incorporate the monthly news of the IGAE leading indicator, the validity of the strong efficiency hypothesis for GDP growth forecasts should be further investigated taking into account other candidate variables available to the forecasters when producing their forecasts.

5 Estimates of the Monthly Aggregate Shocks

In this section we provide estimates of the aggregate shocks that might have affected SPF's respondents homogeneously at each period in time. It is important to note that, since the estimates obtained through equation (9) can only be interpreted as pure aggregate shocks under the joint hypothesis of optimality and quadratic loss, to extract the perceived impact of monthly aggregate shocks we only employ those individuals who pass the weak efficiency tests at the 10% significance level.

Time series estimates of the impact of aggregate macroeconomic shocks affecting the expectation formation process of annual inflation and GDP growth forecasts are calculated following equation (9). Those corresponding to inflation are plotted in Figure 6, while those corresponding to GDP growth can be found in Figure 7. In each case, we decided to report shocks affecting current year and following year forecasts separately to distinguish whether current news affect short and medium-term expectations in a different fashion. Positive values of these shocks imply an increase in the value of the target variable, while negative values refer to a decrease in the value of the target variable. In this sense, positive shocks can be interpreted as “bad news” for inflation forecasts and “good news” for GDP growth forecasts, and viceversa. The dates on the horizontal axis give the month of the survey that collected the revised forecast, indicating the timing of the new information that may have influenced the revision. In both figures, the left panel plots the aggregate shocks and the right panel plots the volatility of those shocks, obtained as the interquartile range of forecast revisions across respondents. The larger the volatility of the aggregate shock in a given month, the larger is the disagreement among forecasters as to the effect of that month's news on the target variable. This could be seen as a proxy of the overall uncertainty of

¹²Note that the constant is significant at the 1% level, and hence, the forecasts can not be labeled efficient. This explains the apparent contradiction that forecasts for GDP growth seem to fail weak efficiency tests but seem to pass strong efficiency ones.

forecasters regarding the impact of news.

The left panel of Figure 6 shows that the largest shocks to inflation expectations in the sample under study happened in February, March, April, and November of 1995, where the respondents of the SPF revised, on average, their annual inflation forecasts around 7%, 22%, 4.4%, and 2.7%, respectively, a cumulative revision that accounts for 70% of the annual inflation level that year, situated around 52%. These “bad news” shocks to aggregated inflation expectations can be regarded as a consequence of the Tequila crisis that started on December of 1994, in which the sudden exchange rate devaluation triggered a new surge of inflation and a considerable output drop during 1995 as internal adjustments were needed in order to absorb the impact of the rundown of international reserves and the posterior suspension of access to external savings (Gil-Díaz and Carstens, 1996). In order of magnitude, the next important shock to inflation expectations occurred in August of 1998, when SPF’s respondents revised their inflation expectations upwards in 1.56% and 2.24% for current and following year, respectively. As in the case of the economic crisis of 1995, since the sign of this shock is positive, it can be interpreted as bad news for inflation (i.e., an increase in inflation). This shock could be related to the effects on Latin American countries of the financial crises originated in Asia and Russia in mid-1997. It can also be seen that between 1999 and 2001 the aggregate shocks are, on average, negative, which imply good news for inflation. These “good news” shocks coincide with a series of disinflation efforts exerted by Banco de México and with the transition to an inflation-targeting regime that culminated with its implementation in 2001.

With respect to the volatility of the aggregate shocks to expected inflation presented in the right panel of Figure 6, it is noteworthy the high volatility before 2001, both in periods of bad and good news, and the low volatility since inflation-targeting was implemented. In fact, the average interquartile range of current year aggregate shocks from February 1995 to January 2001 is 0.84%; whereas it is 0.25%, on average, from February 2001 until July 2009. In addition, a small increase in the volatility can be perceived at the end of the sample, clearly related to the uncertainty surrounding the global financial crisis that started at the end of 2007. In addition, when we distinguish between the aggregate shocks for current year from those available for the following year, we can see that, in the majority of the sample, short-term expectations (i.e., those for current year) are more sensitive to a given shock than medium-term expectations (i.e., those for the following year).

The left panel of Figure 6 displays the sequence of monthly aggregate shocks available for GDP growth expectations for current and following year forecasts. Consistent with the effects of the Tequila crisis on inflation expectations, the estimates of aggregate shocks to GDP growth forecasts also indicate one extreme negative (“bad news”) shock on August 1995

of around 0.90%. In the same line, a series of “bad news” shocks to GDP growth expectations can be found along 1998. These are related, as in the case of inflation expectations, to the effect of the Asian and Russian crises. A bad news shock on GDP growth forecast revisions of 1.2% can also be found in October 2001, which coincides with the recession in Mexico between 2001 and 2002, which is in turn related to the global economic downturn of those same years, particularly, in the United States. In the last part of the sample one can see a series of increasing bad news shocks appearing since January 2008 until the first half of 2009 related to the recent global financial crisis and the recession in the United States that started in December 2007. The volatility of the aggregate shocks to GDP growth forecast revisions indicate high volatility for most of the sample, except for a relatively brief period of stability between 2004 and 2006. As a consequence of the most recent recession, the volatility of the aggregate shocks to GDP growth forecasts has been increasing steadily in the last part of the sample.

6 Discussion

One important difference between survey forecasts with respect to pure model-based forecasts is that they typically include the subjective views (i.e., extra-sample information) of the experts that participate in the survey. This extra information, or expertise, may or may not help to forecast more precisely than other forecasting methods, given a certain evaluation metric.

In this paper we have evaluated the fixed-event forecasts of inflation and GDP growth from the respondents to the Banco de México’s Survey of Professional Forecasters. As in other countries, the forecasters seem to start giving a forecast close to what they consider to be the unconditional mean of the variable of interest 24 months before the realization of the actual value. After the initial forecasts, news seem to start to be incorporated around 12 months before the realization of the actual end-of-the-year inflation, and about 16 months before in the case of GDP growth.

We exploit the fixed-event forecast structure of our data to study the efficiency of the forecast revision process. Tests of weak efficiency based on pooled estimations show evidence of over-smoothing (or anchoring) of the forecasts, in the sense that the revisions exhibit positive serial correlation. Hence, forecasters tend to hold to their views for longer than appears to be optimal. With respect to strong efficiency, forecasters seem to incorporate macroeconomic news as reflected on the global indicator of economic activity in the case of forecasts of GDP growth. In contrast, forecast revisions of inflation do not seem to systematically incorporate the information in past monthly inflation.

One important issue to point out is that the empirical assessment of efficiency implemented in this paper assumes that forecasters only care for the deviations of their predictions with respect to the actual value of the target variable. However, professional forecasters might have other considerations when computing their forecasts as the forecasts they produce matter for decisions on monetary and fiscal policy, portfolio allocations, wage negotiations, etc. For example, Batchelor and Dua (1992) provide a plausible explanation of the forecasting “smoothness” displayed by perfectly rational forecasters in the sense that forecasters are likely to take into account that their clients might mistrust forecasters who make frequent revisions to forecasts and that in this sense, a smooth forecast evolution implies stability, with stable forecasts not being overly sensitive to new information.

The different incentives faced by the forecasters can be consistent as well with other loss functions attaching different weights to the costs of over and under-predicting the variable of interest. These asymmetric loss functions have been studied previously in the literature (Elliott et al., 2005; Patton and Timmermann, 2007; Capistrán, 2008). In fact, the assumption of the loss function is crucial in determining whether the forecasters analyzed can be labeled efficient or not. As Patton and Timmermann (2007) demonstrated, none of the properties traditionally associated with optimal forecasts under a quadratic loss hold with asymmetric loss. The same conclusion arises when we deal with the properties of optimal revisions. Additionally, it can be shown that under certain asymmetric loss functions, the forecast revision process includes an optimal bias depending on the degree of asymmetry and on a measure of the uncertainty surrounding the target variable.

Hence, the inefficiencies found in the forecasts from the SPF are only so with respect to a particular cost function, the quadratic one. However, these results do not imply that the forecasts can not be useful to a decision maker with a cost function of this type. As shown in Bentancor and Pincheira (2010) and in Capistrán and Timmermann (2009b), a user of the forecasts can take advantage of systematic inefficiencies, such as the bias and the excess of autocorrelation, to build new and more accurate forecasts. In this manner, the empirical results presented in this paper suggest that some gain may be achieved in the actual forecasts of the SPF from putting more weight in variables already at the disposal of the professional forecasters. Specifically, one important lesson of this analysis is that even the past performance of the consensus forecast could be employed to refine predictions.

In addition to the efficiency tests, we also estimated a measure of the aggregate monthly shocks that have affected inflation and GDP growth. These are shocks that had an impact on the forecasts, producing forecast errors, but for which the forecasters should not be held accountable for. The estimated shocks appear to be related to specific periods of macroeconomic uncertainty. For instance, for inflation we find a declining impact of the

news the longer the forecast horizon, and a marked effect of the latest disinflation efforts (as good news shocks) and of inflation targeting (as reduced volatility). The aggregate shocks to expected GDP growth show a clear effect of the business cycle, in particular of downturns, and high volatility overall.

Future research should consider the extent to which the results found here can be rationalized by other loss functions, as in Capistrán and Timmermann (2009a). In the same vein, it is necessary to see what portion of the macroeconomic aggregate shocks provided in this document are driven by the assumption of quadratic loss. Finally, another path for future research could consider the different dynamics in the data generating process governing both inflation and GDP growth for the period under study. In particular, it may be important to consider the possible changes in the expectation formation process of professional forecasters in Mexico given the presence of structural breaks in the economy such as the implementation of an inflation targeting regime in the year 2001.

References

- [1] Bakhshi, H., G. Kapetanios and T. Yates. 2005. Rational expectations and fixed-event forecasts: An application to UK inflation. *Empirical Economics*, 30, 539-553.
- [2] Batchelor, R. and P. Dua. 1992. Conservatism and consensus-seeking among economic forecasters. *Journal of Forecasting*, 11, 169-181.
- [3] Bentancor, A. and P. Pincheira. 2010. Predicción de errores de proyección de inflación en Chile. *El Trimestre Económico*, LXXVII, 305, 129-154.
- [4] Boero, G., J. Smith and K.F. Wallis. 2008a. Evaluating a three-dimensional panel of point forecasts: The Bank of England Survey of External Forecasters. *International Journal of Forecasting*, 24, 354-367.
- [5] Boero, G., J. Smith and K.F. Wallis. 2008b. Here is the news: Forecast revisions in the Bank of England survey of external forecasters. *National Institute Economic Review*, 203, 68-77.
- [6] Chiquiar, D., A. E. Noriega and M. Ramos-Francia. 2007. A time series approach to test a change in inflation persistence: The Mexican experience. Banco de México Working Paper 2007-01, forthcoming in *Applied Economics*.
- [7] Capistrán, C. 2007. Optimality tests for multi-horizon forecasts. Working paper Banco de México 2007-14.
- [8] Capistrán, C. 2008. Bias in Federal Reserve inflation forecasts: Is the Federal Reserve irrational or just cautious?. *Journal of Monetary Economics*, 55, 1415-1427.

- [9] Capistrán, C. and G. López-Moctezuma. 2010. Las expectativas macroeconómicas de los especialistas: Una evaluación de pronósticos de corto plazo en México. *El Trimestre Económico*, LXXVII, 306, 275-312.
- [10] Capistrán, C. and A. Timmermann. 2009a. Disagreement and biases in inflation expectations. *Journal of Money, Credit and Banking*, 41, 365-396.
- [11] Capistrán, C. and A. Timmermann. 2009b. Forecast combination with entry and exit of experts. *Journal of Business and Economic Statistics*, 27, 428-440.
- [12] Clements, M.P. 1997. Evaluating the rationality of fixed-event forecasts. *Journal of Forecasting*, 16, 225-239.
- [13] Davies, A. and K. Lahiri. 1995. A new framework for analyzing survey forecasts using three-dimensional panel data. *Journal of Econometrics*, 68, 205-227.
- [14] Davies, A., K. Lahiri and X. Sheng. 2010. Analyzing three-dimensional panel data forecasts. Forthcoming in M. Clements and D. Hendry (eds.) *Oxford Handbook of Economic Forecasting*. Oxford University Press.
- [15] Diebold, F. X., and L. Kilian. 2001. Measuring predictability: theory and macroeconomic applications. *Journal of Applied Econometrics*, 16, 657-669.
- [16] Dovern, J. and J. Weisser. 2008. Are they really rational? Assessing professional macroeconomic forecasts from the G7-Countries. Kiel Institute for the World Economy. Working Paper No.1447.
- [17] Elliott, G., I. Komunjer and A. Timmermann. 2005. Estimation and testing of forecast rationality under flexible loss. *Review of Economic Studies*, 72, 1107-1125.
- [18] Gil-Díaz, F. and A. Carstens. 1996. One year of solitude: Some pilgrim tales about Mexico's 1994-1995 crisis. *The American Economic Review*, 86(2), 164-169.
- [19] Granger, C.W.J. 1999. Outline of forecast theory using generalized cost functions. *Spanish Economic Review*, 1(2), 161-173.
- [20] Granger, C.W.J. and P. Newbold. 1986. *Forecasting Economic Time Series*, 2nd ed. Orlando, FL: Academic Press.
- [21] Harvey, D.I., S.J. Leybourne and P. Newbold. 2001. Analysis of a panel of UK macroeconomic forecasts. *Econometrics Journal*, 4, S37-S55.
- [22] Isiklar, G. and K. Lahiri. 2007. How far ahead can we forecast? Evidence from cross-country surveys. *International Journal of Forecasting*, 23, 167-187.
- [23] Isiklar, G., K. Lahiri, and P. Loungani. 2006. How quickly do forecasters incorporate news? Evidence from cross-country surveys. *Journal of Applied Econometrics*, 21, 703-725.

- [24] Keane, M. and D. Runkle. 1990. Testing the rationality of price forecasts: New evidence from panel data. *The American Economic Review*, 80, 714-735.
- [25] Lahiri, K. and X. Sheng. 2008. Evolution of forecast disagreement in a Bayesian learning model. *Journal of Econometrics*, 144, 325-340.
- [26] Mincer, J. and V. Zarnowitz. 1969. The evaluation of economic forecasts in j. Mincer (ed.), *Economic Forecasts and Expectations*, National Bureau of Economic Research, New York.
- [27] Nordhaus, W.D. 1987. Forecasting efficiency: concepts and applications. *The Review of Economics and Statistics*, 69(4), 667-674.
- [28] Patton, A.J. and A. Timmermann. 2007. Properties of optimal forecasts under asymmetric loss and nonlinearity. *Journal of Econometrics*, 140, 884-918.
- [29] Stinchcombe, M.B. and H. White. 1998. Consistent specification testing with nuisance parameters present only under the alternative. *Econometric Theory*, 14, 295-325.
- [30] Tversky, A. and D. Kahneman. 1974. Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1130.
- [31] Thomas, L. 1999. Survey measures of expected U.S. inflation. *The Journal of Economic Perspectives*, 13(4), 125-144.
- [32] Swidler, S. and D. Ketcher. 1990. Economic forecasts, rationality, and the processing of new information over time. *Journal of Money, Credit and Banking*, 22, 65-76.

Table 1: Initial Forecasts *vs.* Actual Values, Inflation and GDP Growth.

Event	Inflation		GDP Growth	
	Initial Forecast	Actual Value	Initial Forecast	Actual Value
1999	10.70	12.32	-	-
2000	13.78	8.96	3.49	6.63
2001	9.18	4.40	3.92	-0.14
2002	6.37	5.70	4.56	0.82
2003	3.88	3.98	4.19	1.72
2004	3.77	5.19	4.16	4.04
2005	3.70	3.33	3.63	3.19
2006	3.99	4.05	3.64	4.95
2007	3.60	3.76	3.34	3.33
2008	3.54	6.53	3.69	1.52
2009	3.50	3.57	3.63	-6.54
Mean**	3.71	4.34	3.83	1.95
Standard Deviation**	0.18	1.13	0.38	3.60

*The initial forecast is the cross-section mean forecast made at horizon 23

(i.e., 24 months-ahead forecast).

**The mean and standard deviations include the period 2003-2009 for inflation forecasts and 2000-2009 for GDP growth forecasts.

Table 2: Evaluation of Weak Efficiency, Inflation and GDP Forecast Revisions, 1995-2010.

Variables	Inflation	GDP
α	0.008 (0.018)	-0.047** (0.019)
$u_{\tau,h+1,i}$	0.006 (0.020)	0.002 (0.027)
$u_{\tau,h+2,i}$	0.050*** (0.018)	0.157*** (0.026)
Observations	6043	5971
Number of forecasters	41	41

*, **, and *** denotes statistical significance at the 10%, 5%, and 1% level, respectively.

GMM Standard Errors using the variance-covariance matrix given in the text are in parentheses.

Table 3: Evaluation of Strong Efficiency, Inflation and GDP Forecast Revisions, 1995-2010.

Variables	Inflation	GDP
γ	0.026 (0.023)	-0.056*** (0.019)
$x_{t,h+1} - x_{t,h+2}$	0.258*** (0.058)	0.001 (0.004)
Observations	7630	7536
Number of forecasters	41	41

*, **, and *** denotes statistical significance at the 10%, 5%, and 1% level, respectively.
 GMM Standard Errors using the variance-covariance matrix given in the text are in parentheses.

Figure 1: Annual Inflation Forecasts and Cross-Section Dispersion, 1995-2010

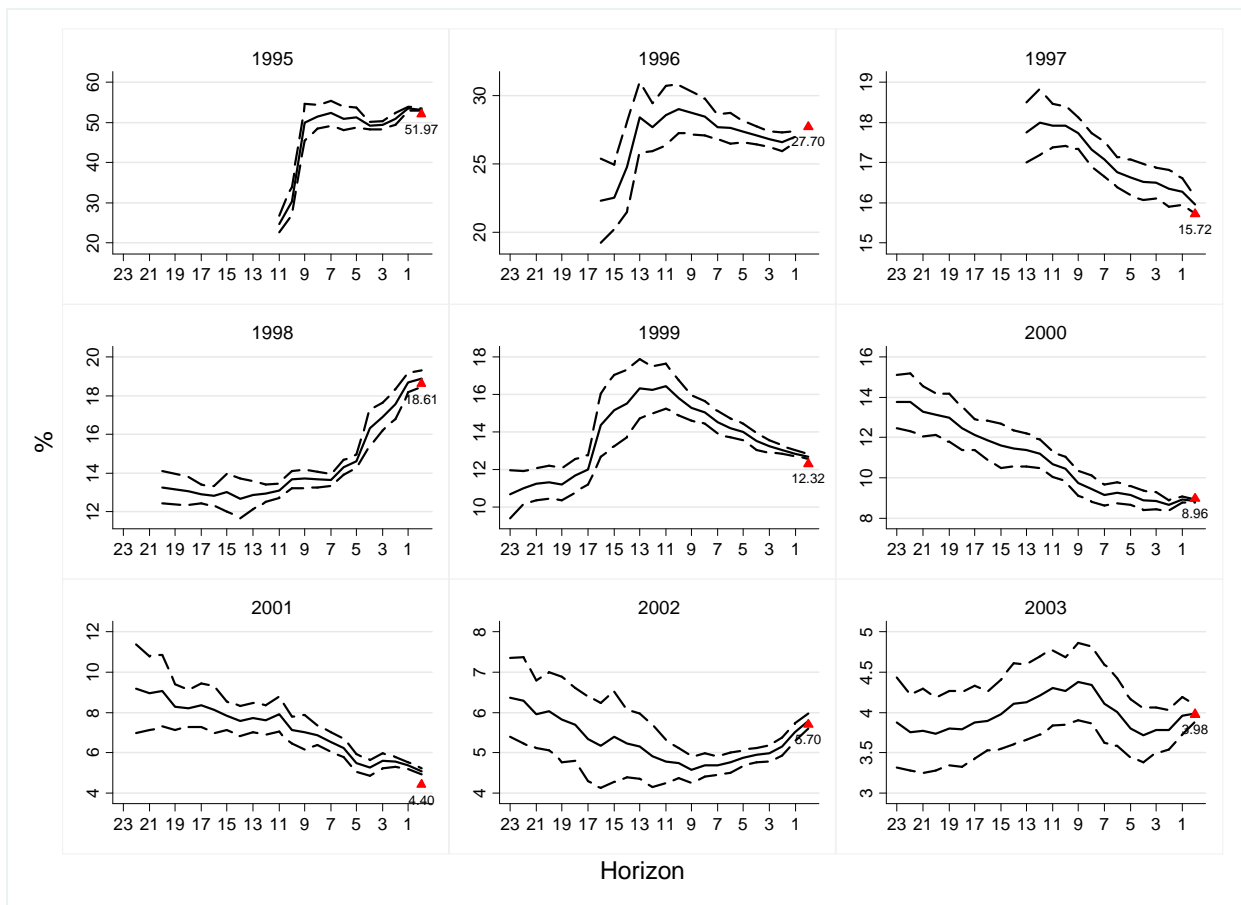



Figure 1 (cont.): Annual Inflation Forecasts and Cross-Section Dispersion, 1995-2010



*The horizontal axis plots the forecast horizon when the prediction was produced.

Horizon 23 corresponds to the forecast made in January of a given year for the outcome of the following year.

Horizon 0 corresponds to the forecast released in December of a given year for the outcome of the current year.

**Red \triangle denotes actual outcome.

***Dashed lines correspond to the mean forecast +/- the sample standard deviation.

Figure 2: Inflation Forecast Revisions and Cross-Section Dispersion, 1995-2010

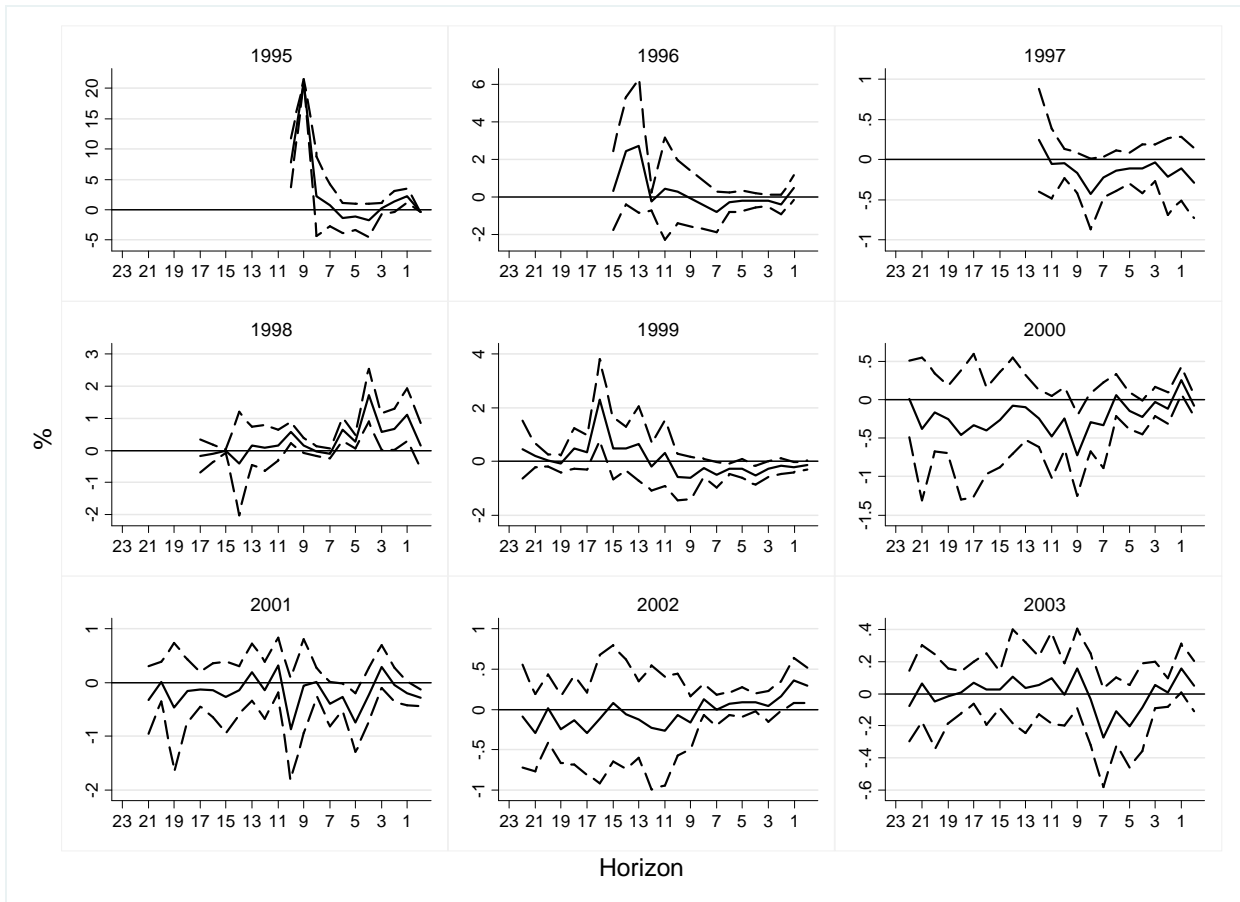



Figure 2 (cont.): Inflation Forecast Revisions and Cross-Section Dispersion, 1995-2010



*The horizontal axis plots the forecast horizon when the prediction was produced.

Horizon 23 corresponds to the forecast made in January of a given year for the outcome of the following year.

Horizon 0 corresponds to the forecast released in December of a given year for the outcome of the current year.

**Dashed lines correspond to the mean forecast +/- the sample standard deviation.

Figure 3: Annual GDP Growth Forecasts and Cross-Section Dispersion, 1995-2010

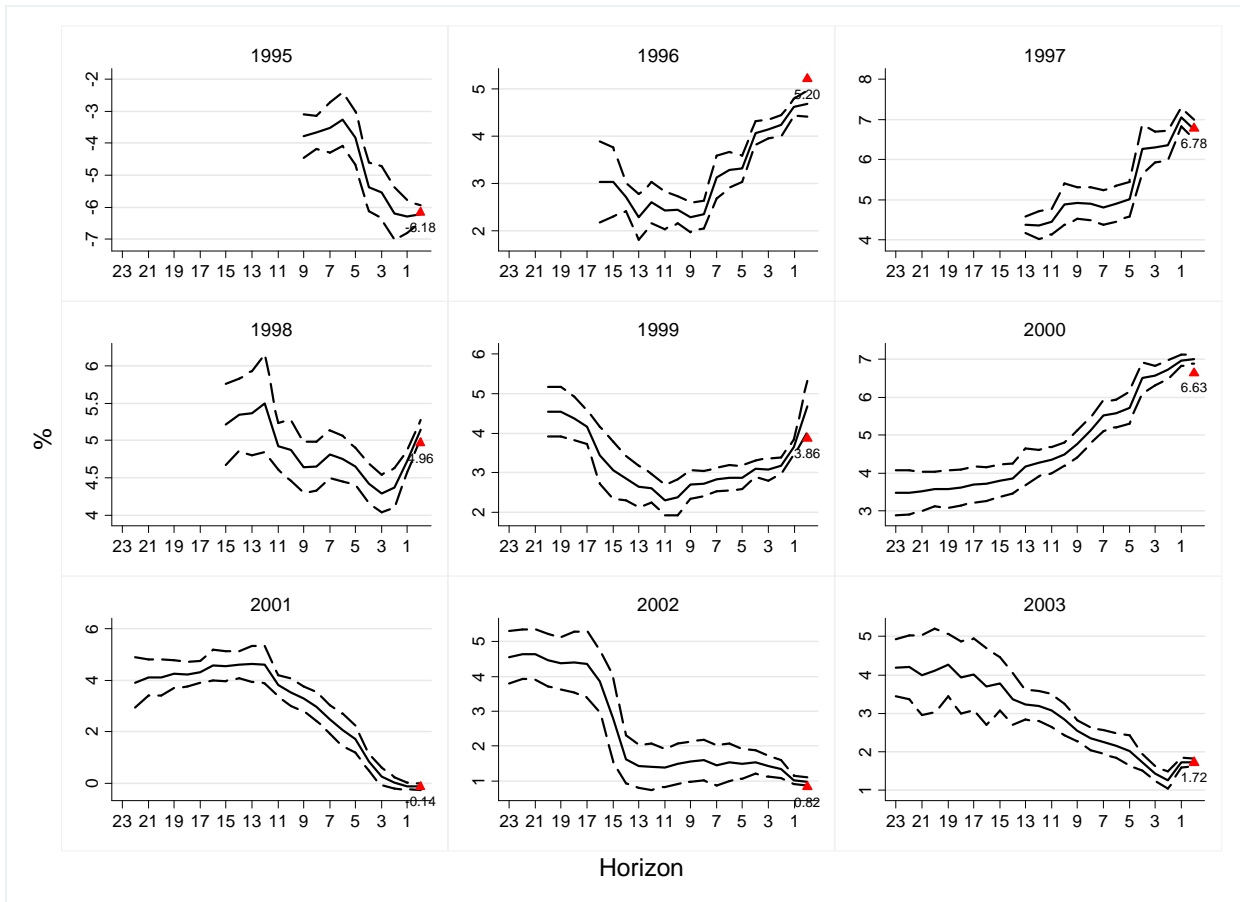



Figure 3 (cont.): Annual GDP Growth Forecasts and Cross-Section Dispersion, 1995-2010



*The horizontal axis plots the forecast horizon when the prediction was produced.

Horizon 23 corresponds to the forecast made in January of a given year for the outcome of the following year.

Horizon 0 corresponds to the forecast released in December of a given year for the outcome of the current year.

**Red \triangle denotes actual outcome.

***Dashed lines correspond to the mean forecast +/- the sample standard deviation.

Figure 4: Annual GDP Growth Forecast Revisions and Cross-Section Dispersion, 1995-2010

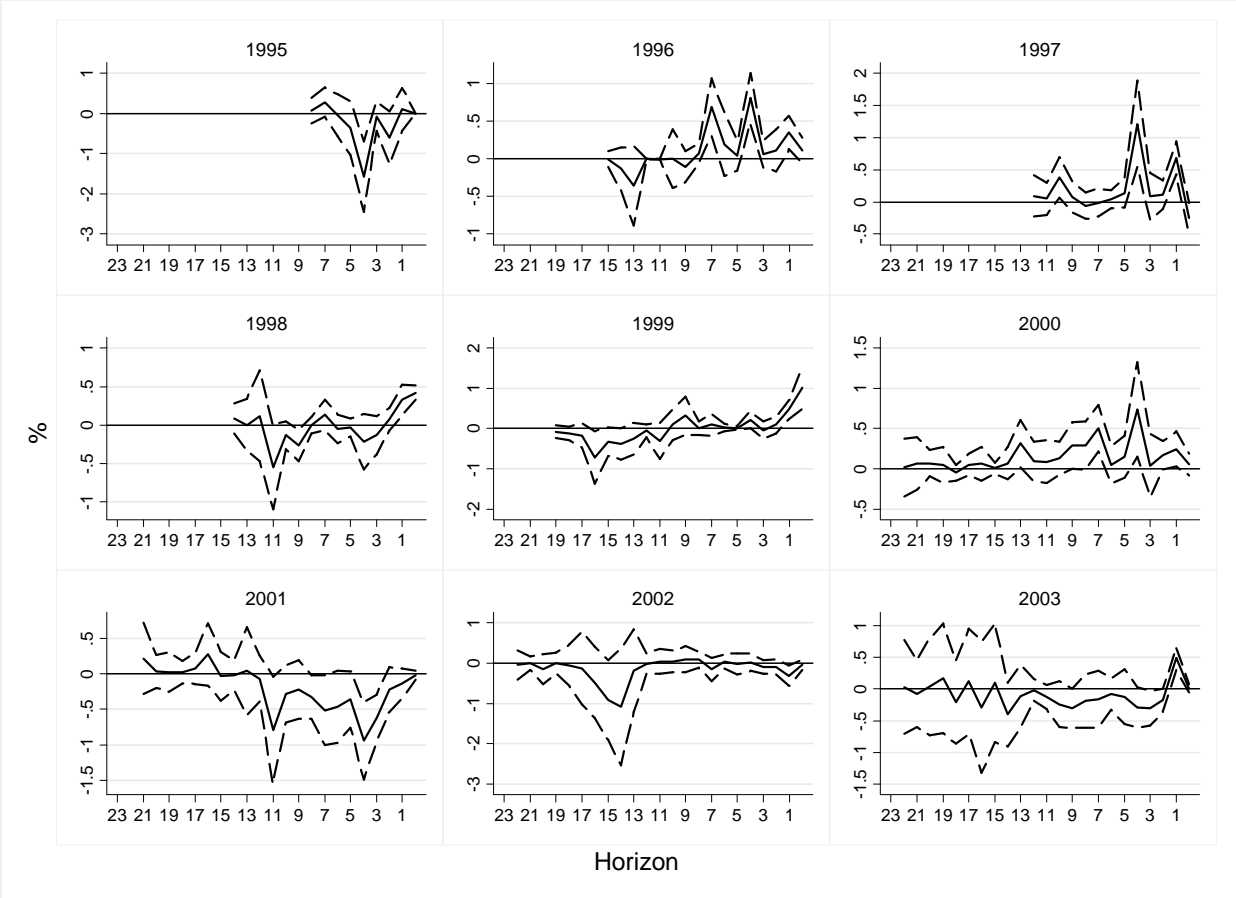
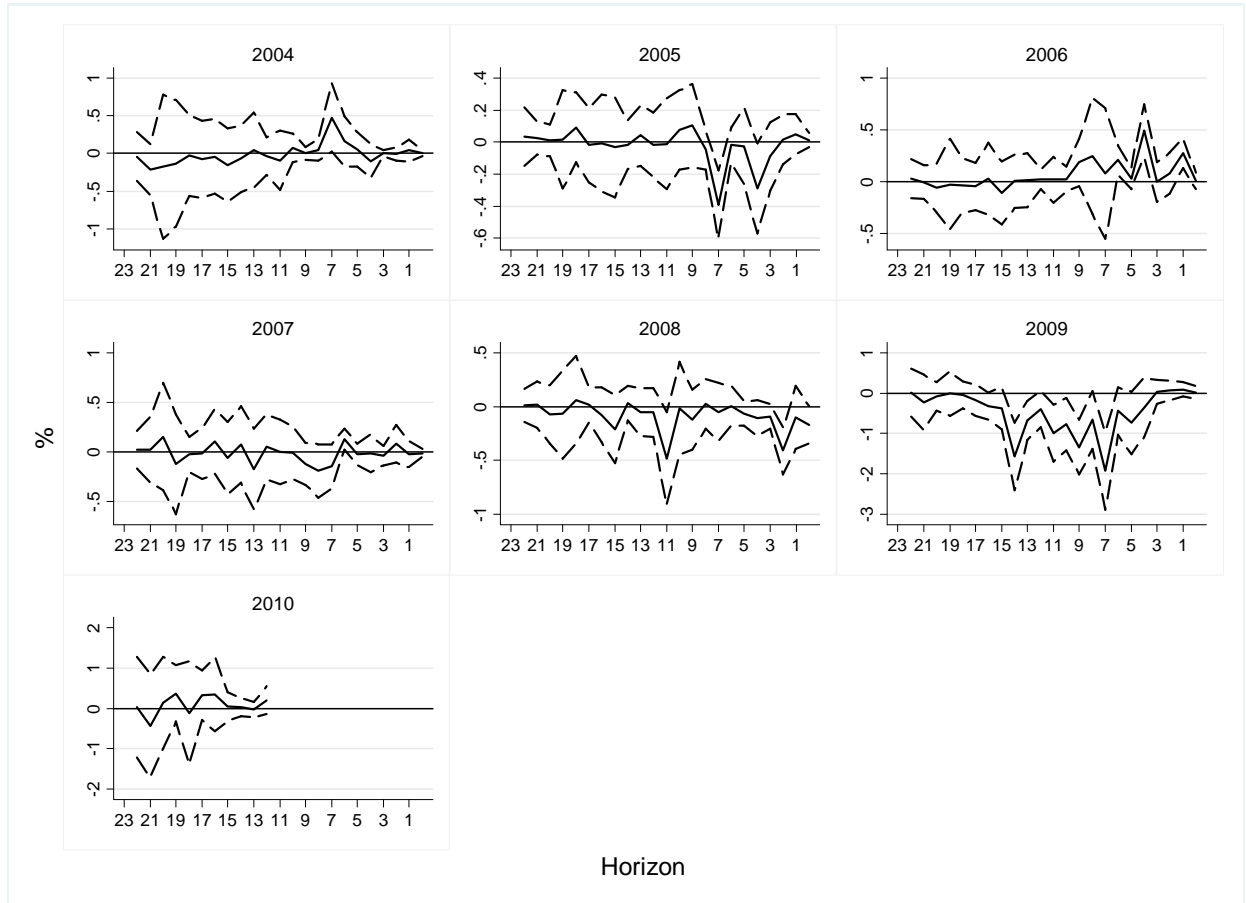


Figure 4 (cont.): Annual GDP Growth Forecast Revisions and Cross-Section Dispersion, 1995-2010




*The horizontal axis plots the forecast horizon when the prediction was produced.

Horizon 23 corresponds to the forecast made in January of a given year for the outcome of the following year.

Horizon 0 corresponds to the forecast released in December of a given year for the outcome of the current year.

**Dashed lines correspond to the mean forecast +/- the sample standard deviation.

Figure 5: Information Content of Forecasts Over Different Horizons



Inflation forecasts

GDP Growth Forecasts


*The measure of information content is $pred_{h,23} = (1 - \frac{MSE_h}{MSE_{23}})100$ due to Diebold and Kilian (2001).

It gives the improvement in the forecast (over the 24 months-ahead prediction) as the horizon decreases.

**For the inflation forecasts we use information over the period 1999-2009 when 23 horizons are available for each event.

***For the GDP growth forecasts we use information over the period 2000-2009 when 23 horizons are available for each event.


Figure 6: Aggregate Shocks to Expected Inflation



*Estimates of the aggregate shocks are estimated by averaging forecast revisions on each month across those individual respondents who passed the weak efficiency tests at the 10% level of statistical significance.

**Estimates of the volatility of aggregate shocks are estimated by taking the interquartile range on each month across those individual respondents who passed the weak efficiency tests at the 10% level of statistical significance.

Figure 7: Aggregate Shocks to Expected GDP Growth



*Estimates of the aggregate shocks are estimated by averaging forecast revisions on each month across those individual respondents who passed the weak efficiency tests at the 10% level of statistical significance.

**Estimates of the volatility of aggregate shocks are estimated by taking the interquartile range on each month across those individual respondents who passed the weak efficiency tests at the 10% level of statistical significance.