PABLO PINCHEIRA
MAURICIO CALANI

Communicational Bias in Monetary Policy: Can Words Forecast Deeds?

A successful communication strategy requires a central bank to be credible. And this, in turn, means matching words with deeds.

—Mario Draghi

The implementation of monetary policy based on setting the overnight interest rate is usually complemented with a strong set of communicational tools that, first, inform markets about the reasons underlying current decisions and, second, indicate the most likely future path of the monetary policy rate (MPR), given the appraisal of the current economic environment. Inflation-targeting countries have been leaders in incorporating this set of tools, which usually comprises periodical inflation reports, financial stability reports, formal speeches, and minutes released immediately after monetary policy meetings.[1] In the case of Chile, the second-oldest inflation-targeting country, these minutes usually include a paragraph signaling the most likely future direction of the MPR. This signal is called the communicational bias, or c-bias.

An unambiguous c-bias can be extracted from several of the minutes, whereas in other cases, the reading of the signal could be subject to readers' and researchers' prejudice or misinterpretation. To avoid a mistaken perception

Pincheira and Calani are with the Central Bank of Chile.

We are very thankful to participants at the monetary policy meetings of the Central Bank of Chile who helped us analyze historical press releases. We also thank Oscar Landerretche for a wonderful discussion of our paper and Roberto Rigobón for his thought-provoking comments. Rodrigo Alfaro, Eduardo Engel, Pablo García, Klaus Schmidt-Hebbel, Barbara Rossi, and Claudio Soto provided helpful and insightful comments. Finally, we would like to thank Felipe Alarcón and Luis Ceballos for their help with the data and Carlos Medel for outstanding assistance.

1. Equivalent to the Federal Open Market Committee (FOMC) meeting of the Board of Governors of the U.S. Federal Reserve.

of the Central Bank of Chile's intended message, we consulted with people who participated in monetary policy meetings and asked them to classify the signals into simple categories, so as to construct a consensus c-bias, which is the final object of our analysis. Market agents might interpret the c-bias differently than the Board of the Central Bank of Chile. In this paper, however, we are concerned with the Central Bank's own consistency in matching words and deeds, which is why we focus on the Central Bank's own interpretation of the c-bias.[2] An evaluation of the Bank's consistency can be easily accommodated in a predictive ability framework because the consensus c-bias is nothing but a forecast of the future direction of the MPR.

How accurate should the c-bias be to support the hypothesis that the Central Bank of Chile is indeed matching words and deeds? In principle, there is no natural accuracy threshold above which we could confidently draw a conclusion in favor of the credibility of the c-bias. Furthermore, in an uncertain economic environment, a hundred percent accuracy is not even expected for the best possible forecasting device. Without any absolute threshold against which to compare the c-bias, we rely on a number of relative thresholds coming from several alternative forecasting strategies. The lack of a more accurate alternative forecasting strategy is consistent with a credible c-bias. In contrast, if the c-bias were systematically outperformed by alternative forecasting methods, one could hardly argue in favor of consistency between words and deeds, as relying on the c-bias would lead to systematic mistakes that could be avoided by using other forecasting approaches.

In this context, we focus on evaluating the informal null hypothesis that the c-bias is a sufficient indicator of the Board's decisions regarding the future direction of the MPR. We test this informal null hypothesis by making use of formal out-of-sample tests of predictive ability. Our basic framework considers the c-bias a natural forecast for the future direction of the MPR, but with no specific forecasting horizon. When the Board releases a c-bias, there is no indication about the timing of the decisions that are supposed to be made

2. Consistency between words and deeds is key to establishing a credible policy. Furthermore, the consensus new Keynesian model predicts that current economic developments and variables such as the exchange rate and long-term interest rates are dependent on the expectation of the future evolution of monetary policy (Galí, 2008). It is through these forward-looking variables that part of the transmission mechanism for taming inflation operates (Svensson, 2003). If the c-bias is a credible announcement, then it should be relevant for expectations and have an impact on economic outcomes, although this impact may be subtle and hard to identify using econometrics in small samples. A strong relationship between the future MPR and the c-bias is a necessary condition for this impact to exist. Should this relationship fail, the link between the c-bias and economic outcomes would lack logical support.

in the future, if conditions do not deviate far from the baseline scenario. Therefore, we explore up to twelve months to see if the Board has fulfilled its intentions as expressed in the c-bias. This point is important for the interpretation of our results. When analyzing twelve forecasting horizons, we could easily be confronted with mixed evidence: the alternative benchmark could outperform the c-bias at some horizons but not others. Fortunately, we face this situation only once: the Survey of Professional Forecasters (SPF) outperforms the c-bias when predicting one month ahead, whereas the c-bias is the clear winner at longer horizons. We interpret this as evidence supporting our informal null hypothesis.[3]

We compare the c-bias's predictive ability with several benchmarks. First, we take a random walk (in levels and first differences) and a uniformly distributed random variable, considering three equally likely scenarios: a tightening, easing, or neutral c-bias. We also consider the case of a Taylor rule model including predictors such as the output gap, inflation deviation from the target, and the persistence of the MPR. Survey-based forecasts, as mentioned, are also included in our analysis. Finally, we make use of market expectations derived from the forward rate curve.

According to our terminology, our results indicate that we cannot reject the informal null hypothesis of the c-bias being a sufficient indicator of the future direction of the MPR. In other words, the evidence is consistent with the hypothesis that the Central Bank of Chile matched words and deeds in the sample period. This is so because no other benchmark outperforms the c-bias's predictive ability at forecasting horizons longer than one month.

Beyond our empirical results, our work contributes to the literature in several different directions. First, we assemble a database for a qualitative variable that is, to our knowledge, novel among emerging economies. Second, we contribute to the literature that evaluates central banks' performance under inflation targeting. This literature focuses on evaluating several dimensions of central banks' performance through macroeconomic final outcomes only—not through time-consistency in matching words and deeds, which would be key to establishing a credible policy framework.[4] This last feature, although

3. By sufficient indicator, we mean a variable that cannot be outperformed by an alternative method in providing a likely estimate of the future direction of the MPR. We borrow this concept from the statistics literature, which defines the more precise notion of a sufficient statistic.

4. Macroeconomic outcomes usually comprise inflation and output volatility, shock resilience, inflation level convergence, and sacrifice ratios. Excellent reference papers include Ball and Sheridan (2005), Corbo, Landerretche, and Schmidt-Hebbel (2002), Cecchetti, Flores-Lagunes, and Krause (2006), and Mishkin and Schmidt-Hebbel (2007).

widely recognized, has been confined to the theoretical arena.[5] Additional distinctive features of this paper are the use of out-of-sample tests of predictive ability and the use of an ordered-response model to characterize the evolution of MPRs. This model is used to properly take into consideration the fact that the MPR is a discrete rather than continuous time series.

The rest of the paper is organized as follows. The next section reviews the literature and discusses the importance of communicational tools for monetary policy. We then describe the c-bias used at the Central Bank of Chile and the way in which we deal with the qualitative aspects of the data. Subsequent sections outline the chosen methodologies and present our empirical results. A final section concludes.

Monetary Policy Implementation and Its Communicational Toolkit

Expectations about the future MPR may play a major role in the conduct of monetary policy, since current overnight interest rates may not be as important as the expectation of their persistence and future changes.[6] In fact, the new Keynesian standard model gives expectations a prominent role in the determination of macroeconomic outcomes. But under what circumstances does a central bank's communication provide extra information for shaping expectations beyond that already contained in observed macroeconomic variables? Under rational expectations and perfect (symmetric) information: none. If one assumes rational expectations, any systemic pattern in the way the policy is being conducted should be correctly inferred from the central bank's behavior. Nevertheless, private agents cannot do so perfectly. Models in which information is not perfect, but in which agents must make inferences about how the central bank operates, give a significant role to transparency and communication as tools for overcoming the information gap.[7]

The extent to which any economy departs from rational expectations and perfect information is an empirical issue, and many articles have been written

5. Agénor (2002), for instance, argues that transparency is an unresolved analytical issue in the design of inflation-targeting regimes. Walsh (2007), using a simple new Keynesian framework, concludes that a policy's impact is significantly affected by the way policy announcements alter expectations.

6. See Woodford (2005).

7. See Woodford (2005) and more recently Blinder and others (2008) for a general discussion; see Cukierman and Meltzer (1986) and Orphanides and Williams (2005) for precise theoretical models in which providing information to private agents is not only nontrivial, but also welfare-improving.

in this arena. The growing attention to inflation-targeting countries, which have taken huge steps toward increasing transparency and accountability, has only made the topic even more appealing, because inflation-targeting countries usually complement the adoption of this regime with several publications and press releases.[8]

The empirical analysis of communicational tools on macroeconomic outcomes is mainly focused on its impact on interest rates and the yield curve.[9] For the United States, Gürkaynak, Sack, and Swanson attempt precisely to separate the effects of current MPR changes from the effects of announcements and statements by the Federal Open Market Committee (FOMC) of the Board of Governors of the U.S. Federal Reserve.[10] They label these effects current policy and the future path policy. Using a principal-components approach, they conclude that previous studies focusing only on the current MPR change missed most of the story, as the second factor (future policy) accounts for more than three-quarters of the total effect on longer interest rates. Following similar insights, Andersson, Dillén, and Sellin examine a wide set of monetary policy signals, including the publication of inflation reports and executive speeches from the Riksbank.[11] They conclude that current monetary policy actions have their greatest effect on the short end of the yield curve and that signaling appears to have some effects on longer interest rates. Siklos and Bohl examine whether communication is important for explaining interest rate movements by the Bundesbank, using a Taylor rule equation.[12] They find that the communication variable they construct is robust and significant. This communication variable is based on the number of speeches on a particular matter for which an auxiliary equation is estimated. This last equation, however, has on its left-hand side the number of speeches and on the right-hand side current and past values of interest rate changes. Thus, their empirical work might not be quantifying the impact of communication on future interest rates, but rather could be capturing the impact of past policy changes on current policy.

Two papers are more closely related to the present article in the sense that they focus on the predictive ability of communicational tools on future monetary policy changes. Lapp and Pearce study the (in-sample) predictive ability of the bias in the FOMC.[13] They conclude that the bias has some power

8. Batini and Laxton (2007).
9. Blinder and others (2008).
10. Gürkaynak, Sack, and Swanson (2005).
11. Andersson, Dillén, and Sellin (2006).
12. Siklos and Bohl (2007).
13. Lapp and Pearce (2000).

to predict future changes in the Federal funds rate. They show that a bias toward tightening implies, on average, a positive change in the Federal funds rate of 11 basis points, in contrast to a negative change of 37 basis points after an easing bias. Rosa and Verga analyze the recent experience of the European Central Bank using the introductory statements of its president in his monthly press conference.[14] They map wording into an index using the frequency of words associated with the tightness of monetary policy. They show that this index is positively and significantly correlated to subsequent repo rate changes. They also find that the European Central Bank's rhetoric is a complement to, rather than a substitute for, measures of activity and exchange rate movements within an empirical reaction function. They fail to show, however, that the European Central Bank's rhetoric can be a better predictor than Euribor rates. Finally, they regress the change in Euribor rates on the change in one-month forward rates and the first difference of the communication index, both of which are positively related with the dependent variable.

Communicational Bias in Chile

The Board of the Central Bank of Chile makes monetary policy decisions at its monthly monetary policy meetings. In these meetings, which are announced six months in advance, the Board sets the level of the MPR, which is the target rate based on which liquidity is provided to the financial industry.[15] This operational implementation is supported by extensive communication by the Central Bank with the public. In particular, policy decisions are communicated immediately after monetary policy meetings in an official news release or minutes.

These minutes can be broken down into three sections: first, the policy decision is announced; second, the arguments behind the decision are sketched, in terms of domestic and international economic events; and, finally, the last paragraph of the minutes is devoted to providing hints about the "most likely course of future monetary policy" if conditions do not deviate far from the baseline scenario. It is this last signal that we call the communicational bias. The Central Bank of Chile has published these statements since September 1997, but it is only since 2000 that the publications have been issued every

14. Rosa and Verga (2007).
15. Central Bank of Chile (2007).

FIGURE 1. The Monetary Policy Rate in Chile

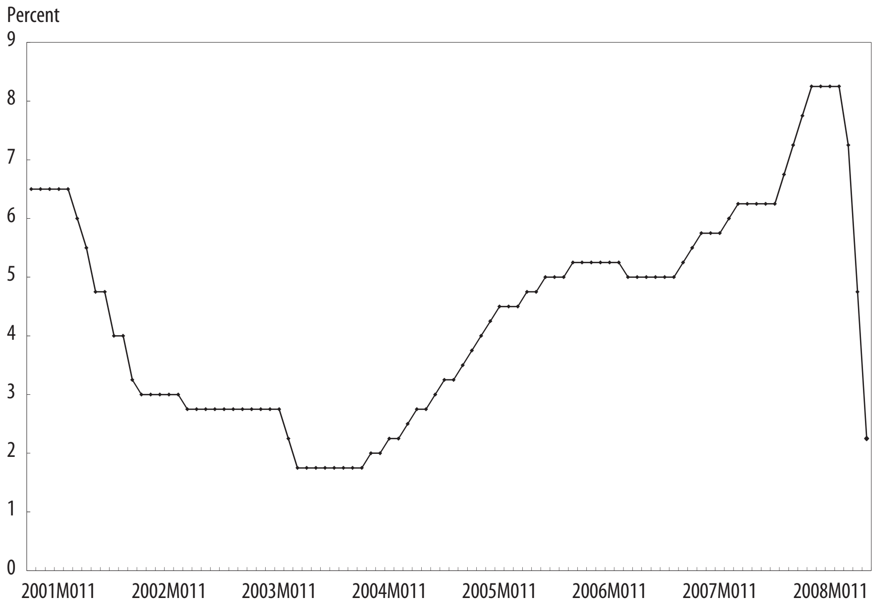

month without interruption. Moreover, in August 2001, the Central Bank changed its target instrument from an inflation-indexed MPR to a nominal interest rate.[16] We therefore decided to work with monthly data from August 2001 to March 2009, which is long enough to capture at least one whole cycle of monetary policy decisions. Figure 1 plots the evolution of the monetary policy rate in Chile during this period. As the figure shows, this variable has been quite persistent.

As mentioned earlier, an unambiguous c-bias can be extracted from several minutes, but the reading of the signal could be subject to researchers' misinterpretation. To avoid a mistaken perception of the Central Bank of Chile's intended message, we asked the people who participated in the monetary policy meetings (with or without voting rights) to classify the signals into the following categories: strong upward bias, moderate upward bias, no change, moderate downward bias, strong downward bias, and no bias.

16. This change implied a large decline in interest rate volatilities; see Fuentes and others (2003).

FIGURE 2. Communicational Bias[a]

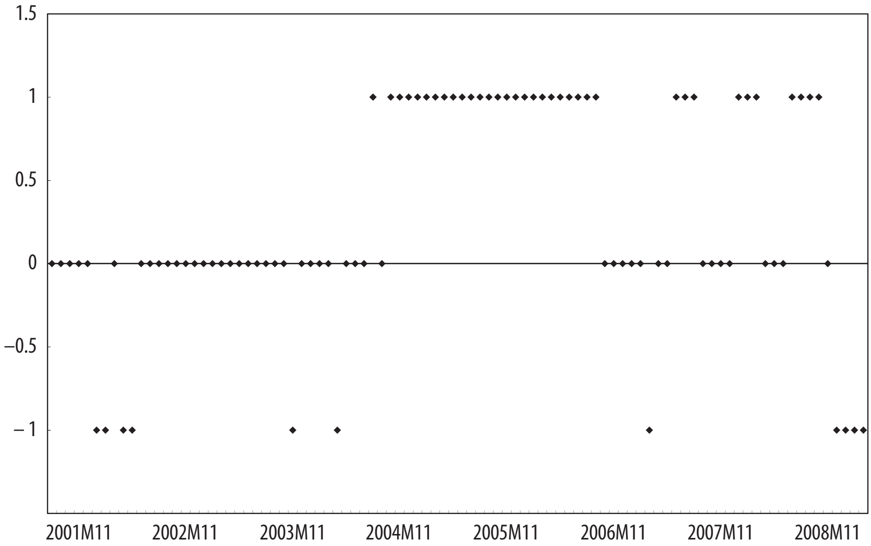

a. A value of one indicates tightening, zero indicates no change (neutral bias), and negative one indicates easing.

In 63 percent of the cases, the people with whom we consulted reported the same opinions on the signal. Only one statement is classified into three different categories, and the rest of the cases have two categories. For these cases, we asked the opinion of other staff economists at the Central Bank and thus reached a consensus on the message every statement provides. Because some of our categories have very few observations, we collapse them into the following three categories: upward bias (including the previous strong and moderate upward bias categories), neutrality (including the previous no change and no bias categories), and downward bias (including the previous strong and moderate downward bias categories). Figure 2 displays this variable, which we call the consensus c-bias. When we collapsed the observations into three categories, the percentage of coincident opinions rose sharply to 89 percent, indicating that most of the initial disagreement stems from different appreciations of the intensity of the c-bias. This strategy also allows us to work with a less ambiguous, and thus higher quality, forecast.

Several stylized facts arise naturally. First, a neutral bias is the most frequent state (50 percent of the time), followed closely by a tightening bias (38 percent). An easing bias, in contrast, only occurs 12 percent of the time.

Second, the c-bias is highly persistent, with only twenty-one changes over the course of ninety-two months. Third, the last third of the sample differs from the first two in terms of persistence. In the first two-thirds, the average maintenance time is five months, and the longest period of an unchanged c-bias is twenty-four months. In the last third of the sample, the average and longest maintenance periods are only three and five months, respectively. Fourth, the c-bias is not very hard to predict. In fact, it is straightforward to build a forecast of the future direction of the MPR from the Survey of Professional Forecasters in Chile. Because this survey is released several days before the monetary policy meetings take place, this forecast could also be used to predict the c-bias. How good of a predictor is it? The success rate in predicting the c-bias is 85 percent, both when the prediction is evaluated with the c-bias as a whole and when only a nonneutral c-bias is used. This high degree of predictability in the c-bias is not surprising and, furthermore, is consistent with authors who favor a predictable monetary policy.[17] Because this paper evaluates consistency between words and deeds, the key exercise should aim at detecting how well the c-bias predicts the future MPR. The results reported later in the paper show that the c-bias is a really tough benchmark to beat in this task.

When not neutral, the c-bias is a natural predictor of future changes in the direction of the MPR. When neutral, however, the c-bias cannot be interpreted as a forecast of some future policy decision. In the forty-six months in which the c-bias is neutral, only once was the original category no change. The rest of the time, the neutral label means that the c-bias was absent and no forecast was released. Not considering this point would probably lead to unfair conclusions, as we would be wrongly interpreting the absence of c-bias as a signal of future MPR movements. Consequently, we test the predictive ability of the c-bias only when this variable is not neutral. We do this following Giacomini and White's conditional predictive ability framework.[18]

Finally, the consensus c-bias is not a real-time variable. As mentioned earlier, we constructed this variable using the retrospective interpretation of people who participated in monetary policy meetings. Therefore, our forecasting evaluation is not an out-of-sample exercise, but rather a pseudo-out-of-sample exercise.[19]

17. See, for instance, Blattner and others (2008).

18. Giacomini and White (2006). A neutral c-bias may mean something and may have an impact on macroeconomic outcomes. In our analysis we are only saying that most of the time a neutral c-bias is not an intended forecast of future MPR movements.

19. We thank Barbara Rossi and Claudio Soto for highlighting this point.

Methodology

As mentioned earlier, if the c-bias is informative about future developments of the MPR, then it should have an impact on economic outcomes. This impact may be subtle and hard to identify using econometrics in small samples. Nevertheless, a strong relationship between the future MPR and the c-bias is a necessary condition for this impact to exist.[20]

Our empirical exercise entails two particular econometric challenges. First, most of the time series we deal with in this paper show high degrees of persistence. One may think that a similar pattern of persistence between the c-bias and the MPR may be driving the econometric results reported later in the paper, but we show that the c-bias's rate of success in predicting the MPR is similar in periods of high and low persistence. We also consider persistent benchmarks when comparing the predictive ability of the c-bias, including martingale models and a Taylor curve with an explicit term introducing persistence. If persistence were the only reason explaining that the c-bias is a good predictor of the MPR, then it would not outperform another benchmark displaying similar patterns of persistence. Our results show that the c-bias is, in general, better than any other benchmark.

Our econometric framework for analyzing predictive ability is designed to work with autocorrelated data. In particular, we construct our test statistics using heteroskedasticity- and autocorrelation-consistent (HAC) variance estimators.

The second econometric challenge concerns the discrete nature of both the c-bias and MPR changes. Jansen and de Haan are, to our knowledge, the first to take into consideration the discrete characteristic of the data, but they do not perform formal predictive ability tests and rely solely on the goodness of fit (pseudo-*R*-squared) of their estimations.[21] In this paper, we explicitly consider the discrete nature of the data and use the formal out-of-sample predictive ability test proposed by Giacomini and White to test for equal predictive ability between the c-bias and several benchmarks.[22] Before presenting our results, we briefly summarize the intuition behind this test.

20. Appendix A analyzes the impact of the c-bias on the forward curve using an event-study approach. Our results show that the c-bias has a statistically significant impact on the short end of the curve. See the appendix for further details.

21. Jansen and de Haan (2006).

22. Giacomini and White (2006).

The Giacomini-White Conditional Approach

We consider two competing parametric forecasting models for the conditional expectation of a scalar time series, Y_{t+1}.[23] We denote the forecasts from these two models as $y^1_{t+1}(\beta_1)$ and $y^2_{t+1}(\beta_2)$, where β_1 and β_2 are population parameters of the two competing models. For a given loss function, $\mathcal{L} = \mathcal{L}[Y_{t+1}, y^i_{t+1}(\beta_i)]$, $i = 1, 2$, the traditional unconditional approach suggests the following test of equal forecast accuracy:[24]

$$(1) \qquad H_0 : \mathbb{E}\left[\mathcal{L}\left(Y_{t+1}, y^1_{t+1}\left(\beta_1\right)\right) - \mathcal{L}\left(Y_{t+1}, y^2_{t+1}\left(\beta_2\right)\right)\right] = 0,$$

whereas the conditional approach suggests the following testing strategy:

$$(2) \qquad H_0 : \mathbb{E}\left[\mathcal{L}\left(Y_{t+1}, y^1_{t+1}\left(\hat{\beta}_{t1}\right)\right) - \mathcal{L}\left(Y_{t+1}, y^2_{t+1}\left(\hat{\beta}_{t2}\right)\right) \middle| \mathcal{F}_t\right] = 0,$$

$$\text{almost surely for all } t \geq 0$$

where $\hat{\beta}_{t1}$ and $\hat{\beta}_{t2}$ denote parameter estimates of β_1 and β_2 with information up until time t. The implementation of the conditional approach relies on the fact that equation (2) is equivalent to

$$\mathbb{E}\left\{h_t\left[\mathcal{L}\left(Y_{t+1}, y^1_{t+1}\left(\hat{\beta}_{t1}\right)\right) - \mathcal{L}\left(Y_{t+1}, y^2_{t+1}\left(\hat{\beta}_{t2}\right)\right)\right]\right\} = 0,$$

for all $\mathcal{F}_t$-measurable functions h_t.

One-Step-Ahead Conditional Test

When the forecasting horizon is $\tau = 1$, $h_t \Delta\mathcal{L}_{R,t+\tau}$ is a martingale difference sequence if the null hypothesis is true. Giacomini and White propose the following statistic for the test of equal conditional predictive ability:[25]

$$(3) \qquad T^h_{P_n,R} = P_n\left(\bar{Z}'_{P_n,R}\Omega^{-1}_{P_n}\bar{Z}_{P_n,R}\right),$$

23. In this section, we follow closely Giacomini and White (2006).
24. This approach is attributed to Diebold and Mariano (1995) and West (1996).
25. Giacomini and White (2006).

where P_n denotes the total number of forecasts, $T + 1$ is the total number of available observations, R denotes the maximum size of the rolling estimation window, and

$$\bar{Z}_{p_n,R} = \frac{1}{P_n}\sum_{t=R}^{T} Z_{R,t+1};$$

$$Z_{R,t+1} = h_t \Delta\mathcal{L}_{R,t+1};$$

$$\hat{\Omega}_{P_n} = \frac{1}{P_n}\sum_{t=R}^{T} Z_{R,t+1} Z'_{R,t+1}.$$

Giacomini and White provide conditions under which the asymptotic distribution of $T^h_{P_n,R}|H_0$ is chi-square:[26]

$$T^h_{P_n,R}\Big|H_0 \xrightarrow{D} \chi^2_q \quad as\ P_n \to \infty.$$

When the dimension of the testing function h_t is one, the test is asymptotically normal.

Multi-Step Conditional Test

When the forecasting horizon is $\tau > 1$, Giacomini and White propose the following statistic for the test of equal conditional predictive ability:[27]

$$T^h_{P_n,R,\tau} = P_n\left(\bar{Z}'_{P_n,R}\tilde{\Omega}^{-1}_{P_n}\bar{Z}_{P_n,R}\right) \qquad (4)$$

where

$$\bar{Z}_{p_n,R} = \frac{1}{P_n}\sum_{t=R}^{T-\tau} Z_{R,t+\tau+1},$$

$$Z_{R,t+\tau+1} = h_t \Delta\mathcal{L}_{R,t+\tau},$$

and $\tilde{\Omega}_{P_n}$ is a HAC estimate of the variance of $\sqrt{P_n}\ \bar{Z}_{P_n,R}$.

26. Giacomini and White (2006).
27. Giacomini and White (2006).

Giacomini and White provide conditions under which the asymptotic distribution of $T^{h}_{P_n,R,\tau}|H_0$ is chi-square:[28]

$$T^{h}_{P_n,R,\tau}\Big|H_0 \xrightarrow{D} \chi^2_q \quad as\ P_n \to \infty.$$

Again, when the dimension of the testing function h_t is one, the test is asymptotically normal.

We test conditional predictive ability using a very simple testing function, h_t,

$$h_t\left(\text{c-bias}_t\right) = \begin{cases} 1 & \text{if c-bias}_t \text{ is not neutral} \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

and thus evaluate the predictive ability of the c-bias only when this signal represents a forecast.

Interpretation of the Test in Our Environment

The c-bias alone is a predictor of the direction of change in the MPR, but it does not provide a specific horizon. We assume that behind this c-bias there is a latent predictor of the future MPR, which we call $b_t(k)$ and define as follows:

$$b_t\left(k\right) = \begin{cases} \infty & \text{if c-bias}_t \text{ is upward biased for all } k = 1,2,\dots \\ 0 & \text{if c-bias}_t \text{ is neutral for all } k = 1,2,\dots \\ -\infty & \text{if c-bias}_t \text{ is downward biased for all } k = 1,2,\dots \end{cases}$$

We consider a generic loss function,

$$\mathcal{L} : \mathbb{R}^2 \to \mathbb{R}$$

and

$$\mathcal{L} = \mathcal{L}\left(Y_{t+k},\, y^p_t\left(k\right)\right),$$

28. Giacomini and White (2006).

where $y_t^p(k)$ is a predictor of Y_{t+k}, which uses information available up to time t. This loss function can often be expressed in terms of an increasing function of the difference between the predictor and the variable it attempts to predict:

$$\mathcal{L}\left(Y_{t+k}, y_t^p\left(k\right)\right) = l\left(Y_{t+k} - y_t^p\left(k\right)\right).$$

Even though the most commonly used loss function is quadratic, it is also common to use a loss function based on the direction of change, such as

$$\mathcal{L}\left(Y_{t+k}, y_t^p\left(k\right)\right) = \begin{cases} 1 & \text{if } \operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(y_t^p\left(k\right) - Y_t\right) \\ 0 & \text{otherwise,} \end{cases}$$

where

$$\operatorname{sign}\left(X\right) = \begin{cases} 1 & \text{if } X > 0 \\ -1 & \text{if } X < 0 \\ 0 & \text{if } X = 0 \end{cases}$$

In particular,

$$\mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right) = \begin{cases} 1 & \text{if } \operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(b_t\left(k\right) - Y_t\right) \\ 0 & \text{otherwise.} \end{cases}$$

The expected value of this loss function is then

$$\mathbb{E}\mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right) = \Pr\left\{\operatorname{sign}(Y_{t+k} - Y_t) \neq \operatorname{sign}\left(b_t\left(k\right) - Y_t\right)\right\},$$

which is nothing but the probability of the predictor $b_t(k)$ missing the direction of change of the variable Y_{t+k}, which is the same as the probability that the c-bias will provide a wrong prediction of the future change in the MPR. Let C denote the event in which {c-bias$_t \neq$ neutral} and let

$$\Delta\mathcal{L}_{t+k,t} = \mathcal{L}\left(Y_{t+k}, y_t^p\left(k\right)\right) - \mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right).$$

Then,

$$\mathbb{E}\left[h_t\mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right)\right] = \Pr\left\{\left[\operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(b_t\left(k\right) - Y_t\right)\right] \cap C\right\}$$
$$= \Pr\left[\operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(b_t\left(k\right) - Y_t\right) \middle| C\right]\Pr\left(C\right).$$

Therefore,

$$\mathbb{E}\left[\left(\Delta\mathcal{L}_{t+k,t}\right)h_t\right] = \mathbb{E}\left[h_t\mathcal{L}\left(Y_{t+k}, y_t^p\left(k\right)\right)\right] - \mathbb{E}\left[h_t\mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right)\right]$$
$$\Pr\left[\operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(y_t^p\left(k\right) - Y_t\right) \middle| C\right]\Pr\left(C\right)$$
$$- \Pr\left[\operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(b_t\left(k\right) - Y_t\right) \middle| C\right]\Pr\left(C\right).$$

This expression shows that the expected value of the loss function difference times the testing function, h_t, is proportional to the difference in the failure rate in predicting the direction of change of the future MPR, conditioned on the Board actually communicating a forecast. Most of our analysis uses this econometric framework, comparing the failure rates of two competing predictors.

The next two subsections explicitly demonstrate that the null hypothesis of the Giacomini-White approach translates into very simple conditions for two leading cases among our benchmarks.[29] We analyze the special case of the uniform distribution and the case of a martingale difference model for monthly changes in the MPR.

The Special Case of the Uniform Distribution

One of the benchmarks we use to compare the predictive ability of the c-bias is a pure luck model, which measures the Board's assessment against a lucky guess. To this end, we consider a model in which statements about the future stance of monetary policy are generated independently by a random number generator. This random device associates equal probabilities (of one-third) with the possible future outcomes: tightening, easing, and

29. Giacomini and White (2006).

no change. An obvious problem is that no sample path for these forecasts is available. Nevertheless, a little algebra allows us to properly express the Giacomini-White null hypothesis for this pure luck model in a very simple manner.

—Proposition 1: Let us consider the following random forecasting device,

$$r(Y_t) = \begin{cases} Y_t + \varepsilon \text{ with } \Pr\left(r(Y_t) = Y_t + \varepsilon\right) = \dfrac{1}{3} \\ Y_t \text{ with } \Pr\left(r(Y_t) = Y_t\right) = \dfrac{1}{3} \\ Y_t + \varepsilon \text{ with } \Pr\left(r(Y_t) = Y_t - \varepsilon\right) = \dfrac{1}{3} \end{cases}$$

where Y_t represents the actual MPR at time t, $r(Y_t)$ is such that sign$(r(Y_t) - Y_t)$ is independent of current and future monetary policy decisions. In particular, we assume that sign$(r(Y_t) - Y_t)$ is independent of the direction (sign) of future changes in the MPR and of any function of the c-bias. If this random device is used to generate MPR forecasts, then the Giacomini-White null hypothesis could be expressed as follows:

$$H_0 : \mathbb{E}\left[\mathcal{L}\left(Y_{t+k}, b_t(k)\right)\right] = \frac{2}{3}$$

for the unconditional case and

$$H_0 : \mathbb{E}\left[\mathrm{N}\left(Y_{t+k}, b_t(k)\right)\right] = 0$$

with

$$\mathrm{N}\left(Y_{t+k}, b_t(k)\right) = \frac{2}{3}h_t - h_t\mathcal{L}\left(Y_{t+k}, b_t(k)\right)$$

when the testing function h_t is given by

$$h_t\left(\text{c-bias}_t\right) = \begin{cases} 1 & \text{if c-bias}_t \text{ is not neutral} \\ 0 & \text{otherwise.} \end{cases}$$

The proof of this proposition is presented in appendix B.

The Special Case of the Martingale Difference Model for Changes in the MPR

We also explore the predictive ability of the c-bias with respect to another simple benchmark: namely, a martingale difference model for monthly changes in the MPR. We consider the following model:

$$\text{MPR}_{t+1} - \text{MPR}_t = \text{MPR}_t - \text{MPR}_{t-1} + \xi_{t+1};$$

$$\mathbb{E}\left(\xi_{t+1} \middle| I_t\right) = 0;$$

$$I_t = \left\{\text{information available at time } t\right\}.$$

With this model, we have

$$\mathbb{E}\left(\text{MPR}_{t+k} - \text{MPR}_t \middle| I_t\right) = k\left(\text{MPR}_t - \text{MPR}_{t-1}\right).$$

We therefore have the following predictor for MPRs at time $t + k$:

$$y_t^p\left(k\right) \equiv \mathbb{E}\left(\text{MPR}_{t+k} \middle| I_t\right) = \text{MPR}_t + k\left(\text{MPR}_t - \text{MPR}_{t-1}\right).$$

Thus,

$$\mathcal{L}\left(\text{MPR}_{t+k}, y_t^p\left(k\right)\right) = \begin{cases} 1 & \text{if sign}\left(\text{MPR}_{t+k} - \text{MPR}_t\right) \\ & \neq \text{sign}\left(k\left(\text{MPR}_t - \text{MPR}_{t-1}\right)\right) \\ 0 & \text{otherwise.} \end{cases}$$

The next section reports the results of our *horse race* between the c-bias and all the benchmarks we are considering.

Empirical Results

Figure 3 shows the c-bias's success rate in predicting future changes in the MPR. Gray bars indicate the unconditional success rate, including episodes in which the c-bias is neutral; black bars show the success rate conditional on the Central Bank of Chile issuing a signal (nonneutral c-bias). Because the c-bias is a forecast with no specific forecasting horizon, we explore predictability up to twelve months ahead, a horizon that should be long enough to capture the policy-relevant predictability of the c-bias. The black bars in panel A of the figure indicate that the conditional success rate peaks at more than 80 percent in the fourth month.[30] This success rate is slightly lower at longer horizons. Panel B shows even more interesting results, as it depicts the success rate of the c-bias calculated only in periods in which the c-bias changed. In the figure, the conditional success rate is, on average, just a little lower than the average in panel A, indicating that the behavior of the c-bias as a predictor of the future direction of the MPR is similar in periods of inertia and innovation. In fact, the average success rate during the first six months is 75 percent in panel A and 74 percent in panel B. When averaged over the twelve horizons, the success rate is 77 percent in panel A and 73 percent in panel B.

These high rates suggest that the c-bias is a strong signal of the Central Bank's future deeds. Nevertheless, this simple analysis does not indicate whether this predictability is easy or hard to achieve. To clarify this point, we compare the c-bias as a predictor of the future direction of the MPR against different models. We use the Giacomini-White framework outlined in the previous section and focus on the testing function defined in expression (5).[31]

The C-Bias and Very Simple Benchmarks

We start by considering a simple model assuming that the MPR follows a martingale difference process. The core statistic, presented in the second column of table 1, is proportional to the difference in the failure rate in predicting the direction of future MPR changes, conditional on the Board actually

30. The gray bars roughly show a decreasing success rate as the forecasting horizon lengthens. We do not pay much attention to these results because they are obtained assuming that a neutral c-bias is predicting no change in the MPR, which is not correct because most of the time (98 percent) a neutral c-bias corresponds to no signal whatsoever.

31. Giacomini and White (2006).

FIGURE 3. The Conditional Success Rate of the Communicational Bias[a]

A. Conditional success rate computed using the whole sample

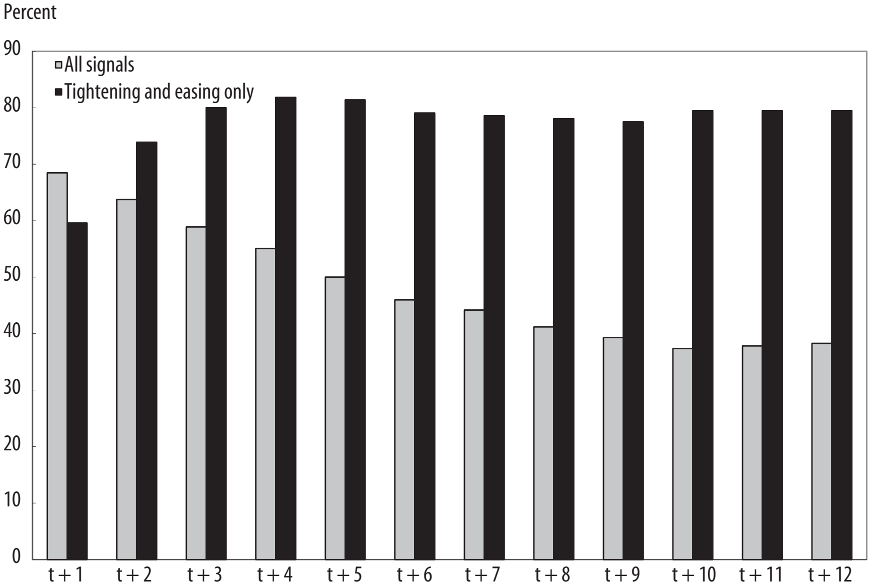

B. Conditional success rate computed when there is a change in the c-bias only.

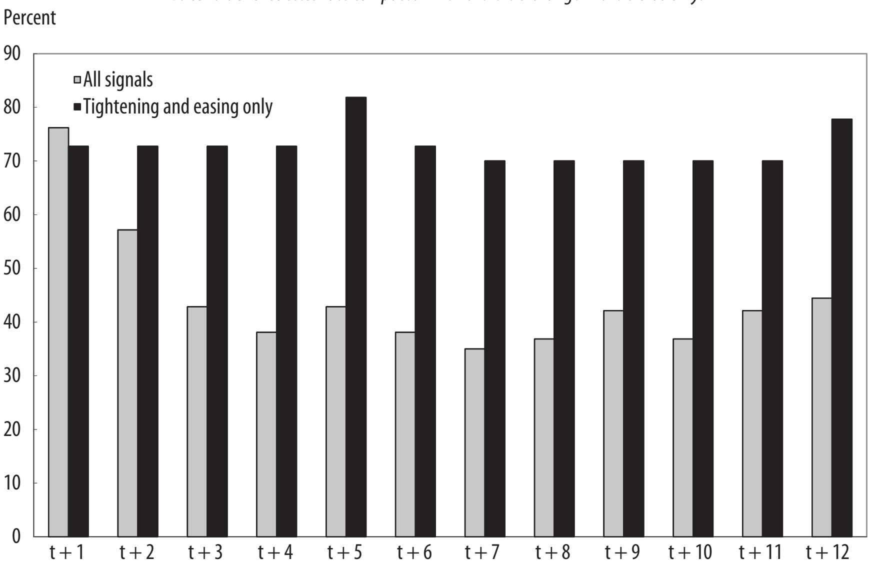

a. Gray bars indicate the unconditional success rate, including episodes in which the c-bias is neutral. Black bars show the success rate conditional on the Central Bank of Chile issuing a signal (nonneutral c-bias).

TABLE 1. Predictive Ability Test for the C-Bias against a Martingale Model[a]

Forecasting horizon	*Core statistic*[b]	*Standard error*	t *statistic*	p *value*
1 month	0.087	0.074	1.182	0.119
2 months	0.231	0.097	2.388	0.008
3 months	0.300	0.098	3.050	0.001
4 months	0.348	0.089	3.907	0.000
5 months	0.375	0.084	4.453	0.000
6 months	0.368	0.089	4.151	0.000
7 months	0.360	0.086	4.188	0.000
8 months	0.341	0.089	3.841	0.000
9 months	0.321	0.093	3.457	0.000
10 months	0.313	0.097	3.232	0.001
11 months	0.317	0.095	3.335	0.000
12 months	0.321	0.095	3.375	0.000

a. The benchmak model is a random walk for the MPR. The exercise uses monthly data from August 2001 to March 2009. Standard errors are based on heteroskedasticity- and autocorrelation-consistent (HAC) variance estimators (Newey and West, 1987, 1994). The *t* statistics are from the Giacomini-White (GW) test and the *p* value is from a one-tailed test.

b. Positive values imply that the martingale is less accurate.

communicating a forecast. The third and fourth columns provide information about the standard errors and the corresponding *t* statistics. The fifth column shows *p* values for a one-sided test of equal predictive ability against the alternative of the c-bias having a better forecasting performance. A positive core statistic means that the loss function associated with the martingale is greater than that associated with the c-bias. In other words, a positive core statistic indicates that the c-bias is a better predictor than the martingale model. As the table shows, this is the case for all forecasting horizons. Furthermore, results are statistically significant in favor of the c-bias for every horizon except the first one, which is only marginally significant at the usual significance levels.

A martingale difference model for the MPR essentially predicts that the MPR will not change in the future. An alternative basic benchmark is a random experiment that imputes equal probabilities (of one-third) to the three possible future outcomes. The results of this comparison indicate that the c-bias does contain statistically significant information to predict the future MPR at every single horizon (see table 2).

Finally, we explore the predictive ability of the c-bias with respect to a martingale difference model for monthly changes in MPR. Table 3 shows our results when using the testing function (5). The c-bias outperforms this benchmark at all horizons except the first one. At the first forecasting horizon, however, the two methods are statistically indistinguishable.

TABLE 2. Predictive Ability Test for the C-Bias against a Random Generator[a]

Forecasting horizon	*Core statistic*[b]	*Standard error*	t *statistic*	p *value*
1 month	0.127	0.039	3.265	0.001
2 months	0.198	0.057	3.503	0.000
3 months	0.226	0.059	3.818	0.000
4 months	0.232	0.059	3.970	0.000
5 months	0.227	0.061	3.716	0.000
6 months	0.218	0.064	3.410	0.000
7 months	0.213	0.064	3.333	0.000
8 months	0.208	0.064	3.252	0.001
9 months	0.202	0.064	3.159	0.001
10 months	0.209	0.064	3.258	0.001
11 months	0.211	0.062	3.440	0.000
12 months	0.214	0.061	3.495	0.000

a. This exercise uses monthly data from August 2001 to March 2009. Standard errors are based on HAC variance estimators (Newey and West, 1987, 1994). The *t* statistics are from the Giacomini-White (GW) test and the *p* value is from a one-tailed test.

b. Positive values imply that the random generator is less accurate.

Ordered Response Taylor Rule Model

The c-bias predicts the direction of future changes of the MPR, but it could be proxying macroeconomic variables that are commonly followed by central banks. We therefore conducted a much more acid test for the c-bias's predictive ability, in which the benchmark is a discrete linear model inspired in a

TABLE 3. Predictive Ability Test for the C-Bias against a Martingale Model for the Difference in the MPR[a]

Forecasting horizon	*Core statistic*[b]	*Standard error*	t *statistic*	p *value*
1 month	0.033	0.050	0.653	0.257
2 months	0.077	0.058	1.331	0.092
3 months	0.111	0.053	2.112	0.017
4 months	0.157	0.049	3.233	0.001
5 months	0.170	0.044	3.835	0.000
6 months	0.172	0.045	3.850	0.000
7 months	0.151	0.043	3.514	0.000
8 months	0.129	0.048	2.674	0.004
9 months	0.131	0.049	2.680	0.004
10 months	0.120	0.053	2.274	0.011
11 months	0.122	0.050	2.443	0.007
12 months	0.136	0.047	2.878	0.002

a. The benchmak model is a random walk for the difference in MPR. This exercise uses monthly data from August 2001 to March 2009. Standard errors are based on HAC variance estimators (Newey and West, 1987, 1994). The *t* statistics are from the Giacomini-White (GW) test and the *p* value is from a one-tailed test.

b. Positive values imply that benchmark model is less accurate.

standard Taylor rule. We impose this structure based on the assumption that future policy rates will change in discrete multiples of 25 basis points, as has been usual in the past. Let $\Delta r_{t+k,t}$ represent possible MPR changes in the period from t to $t + k$, and let k be the forecasting horizon. During k periods, the MPR can change in any direction and in several magnitudes. Let $J(k, t)$ be the number of possibilities of change in the MPR, which depend on both k and t. For example, in the four-year period from July 2003 to June 2007, with $k = 2$, $\Delta r_{t+k,t}$ took six values (namely, –1.00 percent, –0.50 percent, –0.25 percent, 0.00, 0.25 percent, 0.50 percent), and thus $J = 6$. As the forecast horizon lengthens, the number of possibilities of change rises, as do its extreme values. We use an ordered probit model to generate our forecasts using information on inflation and output.

Ordered response models for $\Delta r_{t+k,t}$ can be derived from a latent variable model. Let $\Delta r^*_{t+k,t}$ be a latent variable,

$$\Delta r^*_{t+k,t} = \mathbf{X}_t'\beta + e_{t+k},\ e_{t+k} \sim \mathcal{N}(0.1), \tag{6}$$

where β is $m \times 1$ and $\mathbf{X}_t$ does not contain a constant. Let $\mu_1 < \mu_2 < \mu_J$ be threshold parameters and define (for our example)

$$\begin{array}{lll} \Delta r = -1.00\% & \text{if} & \Delta r^* < \mu_1 \\ \Delta r = -0.50\% & \text{if} & \mu_1 < \Delta r^* < \mu_2; \\ \vdots & \vdots & \vdots \\ \Delta r = +0.50\% & \text{if} & \Delta r^* > \mu_{J-1} \end{array} \tag{7}$$

We can thus easily define the (conditional) probability distribution function for Δr_h, given that Δr_h can take a limited set of values:

$$P(\Delta r = -1.00\% | \mathbf{X}) = P(\Delta r^* \le \mu_1 | \mathbf{X}) = P(\mathbf{X}\beta + e \le \mu_1 | \mathbf{X}) = \Phi(\mu_1 - \mathbf{X}\beta);$$

$$P(\Delta r = -0.50\% | \mathbf{X}) = P(\mu_1 < \Delta r^* \le \mu_2 | \mathbf{X}) = \Phi(\mu_2 - \mathbf{X}\beta) - \Phi(\mu_1 - \mathbf{X}\beta);$$

$$\vdots$$

$$P(\Delta r = +0.25\% | \mathbf{X}) = P(\mu_4 < \Delta r^* \le \mu_5 | \mathbf{X}) = \Phi(\mu_5 - \mathbf{X}\beta) - \Phi(\mu_4 - \mathbf{X}\beta);$$

$$P(\Delta r = +0.50\% | \mathbf{X}) = P(\Delta r^* > \mu_5 | \mathbf{X}) = 1 - \Phi(\mu_5 - \mathbf{X}\beta).$$

We can estimate the parameters μ and β through maximum likelihood (ML) and use these ML estimates for β to compute fitted values for $\Delta r^*_{t+k,t}$. Similarly, we can use the ML estimates for μ to infer a discrete response of $\Delta r_{t+k,t}$.

FUNCTIONAL FORMS. We assume the standard Taylor rule:[32]

$$i_t = c + \alpha\left(\pi_t - \bar{\pi}\right) + \beta\left(y_t - y_t^p\right) + \varepsilon_t, \tag{8}$$

where $\bar{\pi}$ is the target inflation rate, π_t is current inflation (that is, the change in the log consumer price index over the previous twelve months), i_t is the annualized policy rate (MPR), and $(y_t - y_t^p)$ is the output gap, which we abbreviate as y_t^G. Adding persistence to the process creates a better description of the data.[33] If we take a persistence-augmented version of equation (8) for period $t + h$ and then subtract equation (8), we obtain

$$i_{t+h} - i_t = \rho\left(i_{t+h-1} - i_{t-1}\right) + \alpha\left(\pi_{t+h} - \pi_t\right) + \beta\left(y^G_{t+h} - y^G_t\right) + \varepsilon_{t+h} - \varepsilon_t. \tag{9}$$

This expression clearly depends on unrealized data $(t + h > t + h - 1 > t)$. Because we need an expression that links $i_{t+h} - i_t$ to available data at time t, we iterate the first term in equation (9), assuming that inflation and the output gap can be approximated by autoregressive processes, to find an expression in which $i_{t+h} - i_t$ depends only on data available at time t:

$$\begin{aligned} i_{t+h} - i_t = {} & \rho^h\left(i_t - i_{t-h}\right) + \rho^{h-1}\alpha\left(\pi_{t+1} - \pi_{t-h+1}\right) + \rho^{h-1}\beta\left(y^G_{t+1} - y^G_{t-h+1}\right) + \cdots \\ & + \rho\alpha\left(\pi_{t+h-1} - \pi_{t-1}\right) + \rho\beta\left(y^G_{t+h-1} - y^G_{t-1}\right) + \alpha\left(\pi_{t+h} - \pi_t\right) + \beta\left(y^G_{t+h} - y^G_t\right) \\ & + \rho^{h-1}\left(\varepsilon_{t+1} - \varepsilon_{t-h+1}\right) + \cdots + \rho\left(\varepsilon_{t+h-1} - \varepsilon_{t-1}\right) + \varepsilon_{t+h} - \varepsilon_t. \end{aligned}$$

Next, we assume that we can, for instance, approximate $(\pi_{t+1} - \pi_{t-h+1})$ by a first-order autoregressive, or AR(1), process. Then,

$$\left(\pi_{t+1} - \pi_{t-h+1}\right) = \phi_\pi\left(\pi_t - \pi_{t-h}\right) + v_{t+1}. \tag{10}$$

32. Taylor (1993); Woodford (2003).
33. Judd and Rudebusch (1998).

We can do the same with the change in the output gap in period h:

$$(11) \qquad \left(y^G_{t+1} - y^G_{t-h+1}\right) = \phi_y\left(y^G_t - y^G_{t-h}\right) + \omega_{t+1}.$$

Iterating on these results, we obtain

$$(12) \qquad i_{t+h} - i_t = \tilde{\rho}\left(i_t - i_{t-h}\right) + \tilde{\alpha}\left(\pi_t - \pi_{t-h}\right) + \tilde{\beta}\left(y^G_t - y^G_{t-h}\right) + \xi_{t+1,t+h},$$

where $\xi_{t+1,t+h}$ is a function of the shocks $\varepsilon_{t+1}, \ldots, \varepsilon_{t+h}$; $\nu_{t+1}, \ldots, \nu_{t+h}$; and $\omega_{t+1}, \ldots, \omega_{t+h}$. We use this final expression in equation (6) as the model governing the latent variable in the determination of the discrete response Taylor rule.

PREDICTIVE ABILITY TESTS. We use the model in equations (6) and (12) to generate threshold parameters, μ_i. We then fit the model with actual data and save the corresponding discrete forecasts as in equation (7). The estimation procedure uses the first observations in the sample. We take a rolling estimation window of forty observations and compute one- to twelve-month-ahead forecasts to build pseudo-out-of-sample forecast errors. In this experiment, we are not using the vintages of the output gap, but rather are working with revised data, which is an additional source of noise that distinguishes our exercise from a real-time experiment.

There is a clear trade-off between estimation accuracy and the number of observations we use for prediction. We consider forty observations to be appropriate for estimation purposes. Figure 4 shows how the forecasts of the discrete Taylor rule look. Unlike the simpler martingale benchmarks, this model is able to predict both positive and negative future values of ΔMPR.

Finally, we compare the predictive ability of the c-bias and our Taylor rule. Results are displayed in table 4. Positive values of the core statistic indicate that the Taylor rule is, on average, less accurate in predicting the direction of the MPR change than the c-bias. With the exception of a few horizons in which no statistically significant evidence is found, our statistic is indeed positive, and we can confidently reject the null hypothesis in favor of the c-bias for horizons of two, ten, and eleven months ahead. The results are also marginally significant in favor of the c-bias in the last horizon. As in the other cases analyzed thus far, the c-bias is not outperformed at any single horizon.

FIGURE 4. Forecasts of the Change in the MPR Based on the Discrete Taylor Rule Model

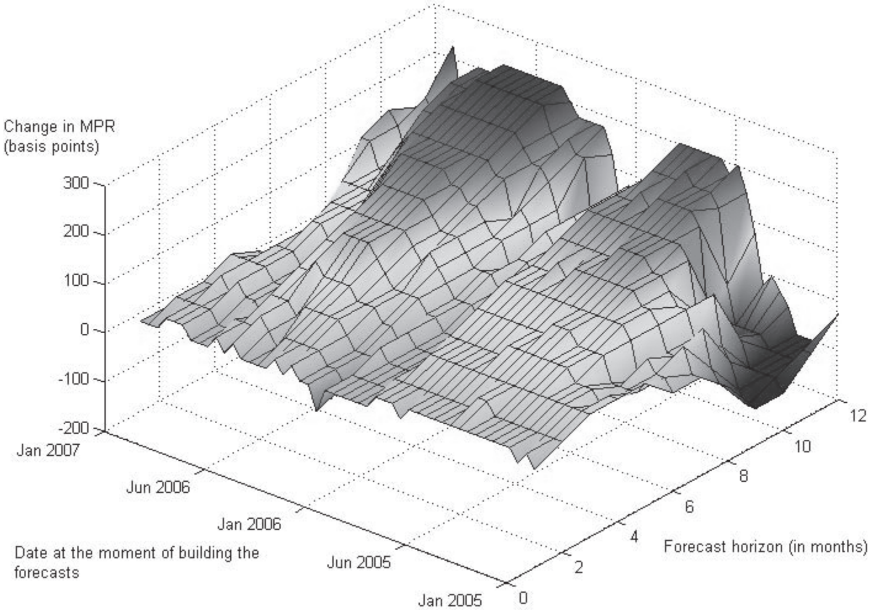

TABLE 4. Ordered Response Taylor Rule Model[a]

Forecasting horizon	*Core statistic*[b]	*Standard error*	t *statistic*	p *value*
1 month	0.038	0.077	0.496	0.310
2 months	0.160	0.113	1.420	0.078
3 months	0.042	0.112	0.371	0.355
4 months	–0.043	0.089	–0.487	0.687
5 months	–0.023	0.109	–0.209	0.583
6 months	–0.048	0.122	–0.390	0.652
7 months	0.050	0.160	0.312	0.378
8 months	0.079	0.174	0.454	0.325
9 months	0.167	0.166	1.004	0.158
10 months	0.265	0.126	2.093	0.018
11 months	0.250	0.132	1.892	0.029
12 months	0.200	0.158	1.269	0.102

a. The table reports a comparison of an ordered response Taylor rule model and the c-bias, using monthly data from December 2004 to March 2009. Standard errors are based on HAC variance estimators (Newey and West, 1987, 1994). The *t* statistics are from the Giacomini-White (GW) test.

b. Positive values imply the c-bias is more accurate than the Taylor rule model.

TABLE 5. Predictive Ability Test for the C-Bias against the Survey of Professional Forecasters[a]

Forecasting horizon	*Core statistic*[b]	*Standard error*	t *statistic*	p *value*
1 month	–0.066	0.030	–2.164	0.985
2 months	–0.044	0.038	–1.184	0.882
3 months	–0.022	0.030	–0.762	0.777
4 months	0.023	0.025	0.906	0.182
5 months	0.057	0.032	1.776	0.038
6 months	0.058	0.033	1.778	0.038
7 months	0.059	0.033	1.779	0.038
8 months	0.060	0.033	1.781	0.037
9 months	0.060	0.034	1.782	0.037
10 months	0.049	0.037	1.310	0.095
11 months	0.062	0.028	2.169	0.015
12 months	0.050	0.022	2.226	0.013

a. The alternative benchmark is the forecast from the Survey of Professional Forecasters, using monthly data from September 2001 to March 2009. Standard errors are based on HAC variance estimators (Newey and West, 1987, 1994). The *t* statistics are from the Giacomini-White (GW) test and the *p* value is a from one-tailed test.

b. Positive values imply the c-bias is more accurate than the alternative benchmark.

The C-Bias and the Survey of Professional Forecasters

The Survey of Professional Forecasters has been carried out periodically by the Central Bank of Chile since February 2000. Around forty individual analysts are asked to provide a number of forecasts for different economic variables at different forecasting horizons. While the individual information is confidential, the Central Bank releases the median of the individual answers on a monthly basis. In particular, expectations for the MPR two months ahead are released since September 2001. We use this information to generate a survey-based forecast of the future direction of the MPR in Chile. The results are presented in table 5. For episodes in which the c-bias is not neutral, survey-based forecasts outperform the c-bias at the first predictive horizon, are statistically equal to the c-bias at the second, third, and fourth forecasting horizons, and are outperformed by the c-bias for the rest of the forecasting horizons (five to twelve months ahead).

The C-Bias and the Forward Rate

In this subsection, we compare the c-bias's predictive ability to that of the forward rate.[34] As shown in table 6, our results indicate that for episodes in

34. Data for the forward rates are based on the estimations of the yield curve performed by RiskAmerica, which are available from October 2002 to March 2009.

TABLE 6. Predictive Ability Test for the C-Bias against the Forward Rate[a]

Forecasting horizon	*Core statistic*[b]	*Standard error*	t *statistic*	p *value*
1 month	0.000	0.032	0.000	0.500
2 months	–0.013	0.028	–0.471	0.681
3 months	–0.026	0.030	–0.864	0.806
4 months	0.000	0.027	0.000	0.500
5 months	–0.014	0.030	–0.446	0.672
6 months	–0.014	0.031	–0.445	0.672
7 months	0.000	0.028	0.000	0.500
8 months	0.000	0.028	0.000	0.500
9 months	0.000	0.029	0.000	0.500
10 months	0.000	0.029	0.000	0.500
11 months	0.000	0.030	0.000	0.500
12 months	–0.015	0.026	–0.579	0.719

a. The alternative benchmark is the forward rate curve, using monthly data from October 2002 to March 2009. Standard errors are based on HAC variance estimators (Newey and West, 1987, 1994). The *t* statistics are from the Giacomini-White (GW) test and the *p* value is from a one-tailed test.

b. Positive values imply the c-bias is more accurate than the alternative benchmark.

which the c-bias is not neutral, the forward rate and the c-bias have statistically equal predictive ability.[35] This result means that, on average, the c-bias and the forward rate are equally accurate in predicting future changes in the MPR. Nevertheless, the information in the c-bias could still be useful for improving the predictive ability of the forward rate. We assess this possibility next.

If the forward rate curve is the best predictor of the MPR under quadratic loss, based on available information at time *t,* then its forecast errors should be orthogonal to information available at the moment of prediction. If orthogonality does not hold, then we could improve the predictive ability of the forward rate by using these nonorthogonal variables. In particular, if the c-bias (when not neutral) contains valuable information that can minimize the prediction error of the forward rate, then the following conditional expectation should be different from zero:

$$\mathbb{E}\left[e_t^f\left(k\right)\middle|\text{c-bias}_t \wedge \text{c-bias}_t \text{ is not neutral}\right] \neq 0, \tag{13}$$

where

$$e_t^f\left(k\right) = \text{MPR}_{t+k} - f_t\left(k\right)$$

35. This is similar to Rosa and Verga’s (2007) findings for the European Union.

represents the forward-curve-based forecasting error at time $t + k$ and $f_t(k)$ corresponds to the monetary policy rate forecast at time $t + k$ coming from the forward curve. Equation (13) is equivalent to

(14) $$\mathbb{E}\left[e_t^f(k)\middle|Tb_t(k) \wedge \{h_t = 1\}\right] \neq 0,$$

where

$$Tb_t(k) = \begin{cases} 1 & \text{if c-bias}_t \text{ is upward biased for all } k = 1, 2, \ldots \\ 0 & \text{if c-bias}_t \text{ is neutral for all } k = 1, 2, \ldots \\ -1 & \text{if c-bias}_t \text{ is downward biased for all } k = 1, 2, \ldots \end{cases}$$

Under the assumption that the conditional expectation of $e^f_t(k)$ with respect to $Tb_t(k)$ is piecewise linear, we have

(15) $$\mathbb{E}\left[e_t^f(k)\middle|Tb_t(k)\right] = \beta_1(k)d_{1t} + \beta_0(k)d_{0t} + \beta_{-1}(k)d_{-1t},$$

where

(16) $$d_{1t} = \begin{cases} 1 & \text{if c-bias}_t \text{ is upward biased} \\ 0 & \text{otherwise} \end{cases}$$

(17) $$d_{0t} = \begin{cases} 1 & \text{if c-bias}_t \text{ is neutral} \\ 0 & \text{otherwise} \end{cases}$$

and

(18) $$d_{-1t} = \begin{cases} 1 & \text{if c-bias}_t \text{ is downward biased} \\ 0 & \text{otherwise} \end{cases}$$

Therefore,

(19) $$\mathbb{E}\left[e_t^f(k)\middle|Tb_t(k); \{h_t = 1\}\right] = \begin{cases} \beta_1(k) & \text{if c-bias}_t \text{ is upward biased} \\ \beta_{-1}(k) & \text{if c-bias}_t \text{ is downward biased} \end{cases}$$

FIGURE 5. Coefficient Estimates of the Forward Rate Forecast Error Regressed on the C-Bias

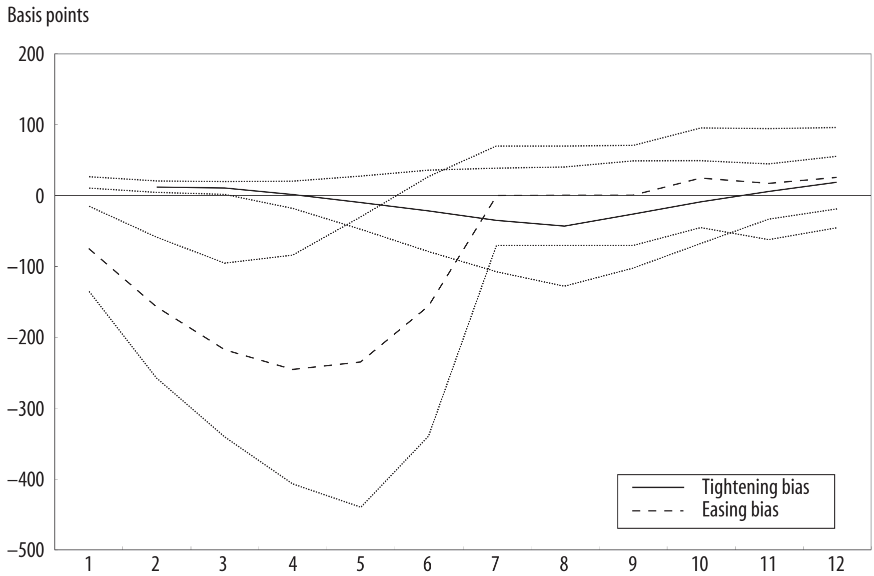

a. The dotted lines represent a 90 percent confidence interval.

Evidence of statistically significant coefficients $\beta_1(k)$ and $\beta_{-1}(k)$ would thus indicate rejection of the null hypothesis,

$$H_0: \mathbb{E}\left[e_t^f\left(k\right) \middle| Tb_t\left(k\right) \wedge \left\{h_t = 1\right\} \right] = 0, \tag{20}$$

and consequently that the c-bias provides useful information for financial agents to predict the MPR.

Figure 5 shows our estimates of $\beta_1(k)$ and $\beta_{-1}(k)$, together with their respective 90 percent HAC-confidence interval for one-sided tests. Under the null hypothesis of forecast errors being random, we expect $\beta_1(k) = \beta_{-1}(k) = 0$. Nevertheless, we find that a tightening c-bias is associated with underprediction of the forward rate for the first four consecutive months. That is, a positive c-bias indicates that the forward rate should adjust upward to recenter the mean of $e(t+h)$ on zero. In terms of the beta coefficients, we find that $\beta_1(1)$, $\beta_1(2)$, and $\beta_1(3)$ are statistically different from zero, indicating that the c-bias contains information that could be useful for improving forecasts from the forward curve. The results are similar when the c-bias signals an easing in monetary policy, in that a downward c-bias is associated with overprediction of the forward rate for the first six consecutive months. However, the size of the revision suggested by our analysis is much bigger than in the case of tightening. In terms of the beta coefficients, we find that $\beta_{-1}(1)$, $\beta_{-1}(2)$, $\beta_{-1}(3)$,$\beta_{-1}(4)$, and $\beta_{-1}(5)$ are statistically different from zero, suggesting that the c-bias contains information that could be useful for improving forecasts from the forward curve.

We check for robustness of these results by augmenting expression (15), first with the actual change in the MPR and then with the actual change in the MPR and one lag.[36] We do this because evidence of statistically significant coefficients might be the result of the omission of the actual change in the MPR, and this variable could be the real driver of our previous results. We run two additional augmented regressions. Figure 6 reports robust estimates of the coefficients using a Bayesian model averaging (BMA) strategy following Brock and Durlauf.[37] We still find some statistically significant coefficients, but the evidence is weaker than before. Now a tightening c-bias is associated with statistically significant underprediction of the forward rate only for the first month. Similarly, an easing c-bias is associated with statistically significant overprediction of the forward rate only for the second and third months. In spite of this reduction in the number of statistically significant coefficients, our robust strategy indicates that the c-bias seems to contain valuable information for improving the predictive ability of the short end of the forward curve.[38]

Conclusions

Monetary policy under inflation targeting relies heavily on the credibility a central bank can build over time. Presumably, this credibility enhances the efficiency of monetary policy and ultimately results in welfare gains. Policymakers are increasingly practicing transparent communication with the

36. Results from the first augmented regression show statistically significant coefficients only when the c-bias signals an easing in monetary policy. In this case, $\beta_{-1}(2)$ and $\beta_{-1}(3)$ are statistically different from zero. In the second augmented regression, we also find that $\beta_{1}(1)$ is statistically significant at the 10 percent significance level.

37. Brock and Durlauf (2001).

38. Appendix C provides a detailed description of the BMA strategy.

FIGURE 6. Robust Coefficient Estimates of the Forward Rate Forecast Error Regressed on the C-Bias[a]

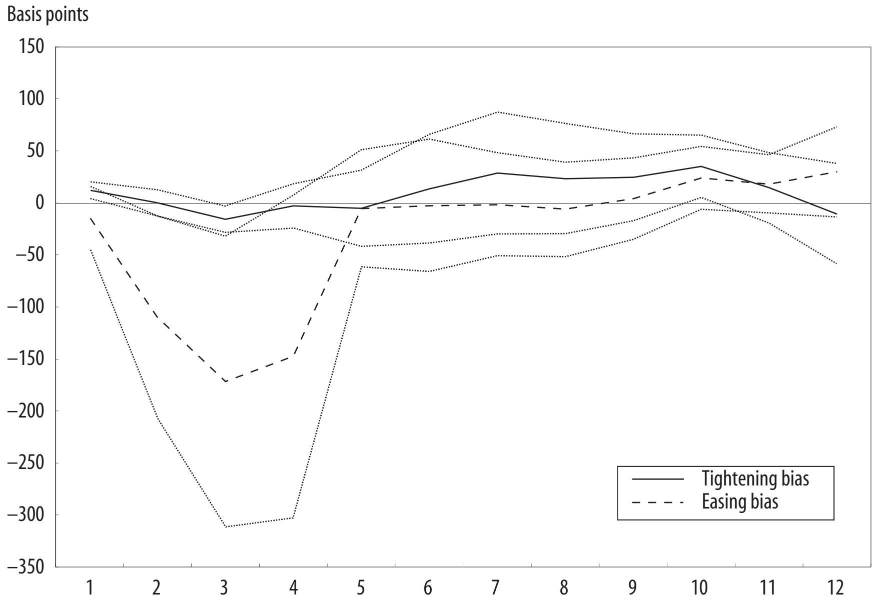

a. Robust estimates using a Bayesian model averaging approach. The dotted lines represent a 90 percent confidence interval.

public to complement their decisions on interest rate setting. We examine one particular feature of the communicational practice of the Central Bank of Chile contained in the press releases published immediately after monetary policy meetings: namely, the Board's assessment of the most likely future of monetary policy, which we call communicational-bias or simply c-bias.

To evaluate the Central Bank of Chile's own consistency in matching words and deeds, we examine whether the communicational bias translates into future monetary policy rate decisions. This analysis can be easily accommodated in a predictive ability framework, because the c-bias is nothing but a forecast of the future direction of the MPR.

In this context, we focus on evaluating the informal null hypothesis of the c-bias being a sufficient indicator of the Board's decisions regarding the future direction of the MPR. We test this informal null hypothesis by making use of formal out-of-sample tests of predictive ability. Our basic framework considers the c-bias to be a natural forecast for the future direction of the MPR, but with no specific forecasting horizon. When the Board releases a c-bias, it provides no indication of the timing of the decisions that are supposed to be made in the future, if conditions do not deviate far from the baseline scenario. We therefore give the Board up to twelve months to match its intentions as expressed in the c-bias.

Based on our results and terminology, we cannot reject the informal null hypothesis that the c-bias is a sufficient indicator of the future direction of the MPR. In other words, the evidence is consistent with a central bank that has matched words and deeds in the sample period. This is so because no other benchmark outperforms the c-bias's predictive ability at forecasting horizons longer than one month.

In particular, we find that the c-bias's conditional success rate in predicting the future direction of the MPR peaks at higher than 80 percent, irrespective of whether the calculation is made over the whole sample or using a subsample displaying lower persistence. We also show that the c-bias more accurately predicts future MPR changes than a martingale model in levels and differences. Similarly, the c-bias strongly outperforms random forecasts generated by a uniformly distributed random variable. Moreover, the pseudo-out-of-sample predictive ability of a more sophisticated model that considers inflation and output can be outperformed by the c-bias at some forecasting horizons. The Survey of Professional Forecasters provides a more competitive benchmark. This survey outperforms the c-bias at the first forecasting horizon, but it never outperforms the c-bias at longer horizons. On the contrary, the c-bias takes the lead when forecasting five to twelve months ahead. Finally, the c-bias is equally accurate as the forward rate curve. Nevertheless, we show that the predictive accuracy of the forward curve could be improved through the use of information from the c-bias.

This evidence is consistent with the c-bias being a strong predictor of the future direction of the MPR. This strong predictive ability is a necessary condition for the c-bias to have an impact on macroeconomic variables.

One additional comment deserves mentioning. The c-bias does not seem to be an outstanding forecast of future developments in the MPR in the very short run. This is so because the c-bias is only able to outperform one of the six models used in the forecasting exercise when prediction is evaluated one month ahead. Moreover, the only case in which the c-bias is outperformed occurs at this specific horizon. The c-bias displays a much better predictive performance at longer horizons. In particular, when predicting five or more months ahead, the c-bias outperforms at least four of the six benchmarks.

A couple of extensions are worth considering for additional investigation. First, the construction of a real-time consensus c-bias would support a reexamination of the results presented in this article from a totally out-of-sample perspective, which would be an improvement over our pseudo-out-of-sample approach. Second, a thorough evaluation of the c-bias's impact on macroeconomic outcomes would also be an interesting object of future research.

Appendix A: The Impact of the C-Bias on the Forward Curve

In this appendix, we provide a first empirical examination of the impact of the c-bias on the forward curve. Larraín examines a related issue, namely, the impact of the MPR on the yield curve.[1] He uses an event-study approach by estimating the following regression:

$$\Delta i_{t,n} = \alpha_n + \beta_n \Delta \mathrm{MPR}_t^u + \varepsilon_{t,n},$$

where the dependent variable is the change in the yield of a bond with maturity n before and after monetary policy meetings, and the independent variable is the unexpected component of the change in the MPR ($\Delta \mathrm{MPR}^u_t$). This component is defined as the difference between the actual MPR and its expected value, which is approximated by the instantaneous forward rate (f_{0t}):

$$\Delta \mathrm{MPR}_t^u \equiv \mathrm{MPR}_t - f_{0t}.$$

We use a similar event-study approach with a few relevant differences. First, our dependent variable is the change in the forward curve, rather than the change in the yield curve. Second, we include the additional regressors of equations (16), (17), and (18), which correspond to the c-bias when signaling a tightening, no change, and an easing in the MPR, respectively. We do this considering that the c-bias is released on the same day that policy rates are set, so it might have an influence on the forward curve. Third, we focus on the short end of the forward curve by analyzing the impact on the first twelve monthly forward rates only. Fourth, our regression includes the expected and

1. Larraín (2007).

unexpected MPR changes separately, where the expected MPR (ΔMPR_t^e) is defined as

$$\Delta MPR_t^e = \Delta MPR_t - \Delta MPR_t^u = f_{0t} - MPR_{t-1}.$$

This appendix thus assesses whether market operators change their perception of the future based on the information in the c-bias. We use the following specification:

$$\Delta f_{i,t} = \beta_{-1i} d_{-1t} + \beta_{0i} d_{0t} + \beta_{1i} d_{1t} + \beta_i^u \Delta MPR_t^u + \beta_i^e \Delta MPR_t^e + w_{t,i},$$

where we regress changes in the forward rates with the different components of the communicational bias and with the unexpected and expected components of MPR changes. We consider twelve regressions: one for each forward rate, from $i = 1$ to $i = 12$. The results are displayed in table A1.

Both the expected and unexpected changes in the MPR have a statistically significant impact on the very short end of the forward curve. This impact is positive and not negligible. The impact of the unexpected component is stronger than that of the expected component.

Table A1 further shows that the c-bias also affects forward rates in the very short term, with a statistically significant impact on the first three forward rates. Two particular results deserve mentioning. First, our regressions suggest that a neutral bias has an impact on the short end of the forward curve, which is a little surprising. Second, the sign of the coefficients associated with the c-bias variables are generally not as expected. For example, an upward bias seems to have a negative impact on the forward curve, which is somewhat counterintuitive.

A number of factors could explain these puzzling results regarding the impact of the c-bias on the forward curve. Rigobón and Sack describe the traditional limitations of event-study approaches.[2] Additional constraints include the small sample size (only seventy-seven observations for regressions in table A1) and the fact that financial markets are not as liquid as an econometrician would wish. Finally, the definition of a c-bias surprise is not clear-cut, so the correct identification of the unexpected component of the c-bias may help clarify our findings. Since this is not the main focus of this article, we leave the interesting task of thoroughly evaluating the c-bias's impact on macroeconomic outcomes to future research.

2. Rigobón and Sack (2004).

TABLE A1. Impact of the C-Bias on the Forward Curve[a]

	Monthly forward rates											
Direction of bias	*i = 1*	*i = 2*	*i = 3*	*i = 4*	*i = 5*	*i = 6*	*i = 7*	*i = 8*	*i = 9*	*i = 10*	*i = 11*	*i = 12*
Negative bias	0.011	0.006	0.002	−0.009	−0.030	−0.014	−0.009	−0.025	−0.006	−0.059	−0.013	0.024
	(0.050)	(0.052)	(0.055)	(0.065)	(0.056)	(0.058)	(0.058)	(0.060)	(0.074)	(0.064)	(0.057)	(0.049)
Neutral bias	−0.137***	−0.082***	−0.049	−0.032	−0.025	−0.020	−0.019	−0.018	−0.036	−0.010	−0.021	−0.021
	(0.031)	(0.025)	(0.030)	(0.034)	(0.037)	(0.040)	(0.039)	(0.040)	(0.040)	(0.044)	(0.044)	(0.038)
Positive bias	−0.166***	−0.091***	−0.058**	−0.036	−0.026	−0.210	−0.029	−0.022	−0.033	−0.016	−0.013	−0.024
	(0.035)	(0.022)	(0.027)	(0.032)	(0.038)	(0.038)	(0.038)	(0.039)	(0.040)	(0.042)	(0.041)	(0.037)
Unexpected change in MPR	0.639***	0.396***	0.150***	0.076	0.032	0.016	0.008	−0.010	−0.014	−0.027	−0.012	−0.035
	(0.031)	(0.040)	(0.520)	(0.060)	(0.055)	(0.066)	(0.061)	(0.063)	(0.071)	(0.058)	(0.070)	(0.065)
Expected change in MPR	0.558***	0.240***	0.102	0.035	0.006	−0.028	−0.015	−0.041	−0.047	−0.069	−0.052	−0.056
	(0.045)	(0.049)	(0.066)	(0.080)	(0.078)	(0.085)	(0.085)	(0.087)	(0.088)	(0.091)	(0.092)	(0.088)
Summary statistic												
R^2	0.78	0.64	0.29	0.09	0.03	0.03	0.01	0.01	0.02	0.03	0.02	0.03
Durbin-Watson statistic	2.08	1.71	1.50	1.53	1.55	1.54	1.54	1.56	1.71	1.70	1.69	1.65

**Statistically significant at the 5 percent level.

***Statistically significant at the 1 percent level.

a. Standard deviations are in parentheses.

Appendix B: Proof of Proposition 1

We consider the following random forecasting device:

$$r(Y_t) = \begin{cases} Y_t + \varepsilon \text{ with } \Pr\left(r(Y_t) = Y_t + \varepsilon\right) = \frac{1}{3} \\ Y_t \text{ with } \Pr\left(r(Y_t) = Y_t\right) = \frac{1}{3} \\ Y_t + \varepsilon \text{ with } \Pr\left(r(Y_t) = Y_t - \varepsilon\right) = \frac{1}{3} \end{cases}$$

where $\varepsilon > 0$ and where Y_t represents the actual MPR at time t. We consider a loss given by

$$\mathcal{L}\left(Y_{t+k}, r(Y_t)\right) = \begin{cases} 1 & \text{if } \operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(r(Y_t) - Y_t\right), \\ 0 & \text{otherwise} \end{cases}$$

but

$$\operatorname{sign}\left(r(Y_t) - Y_t\right) = \begin{cases} \operatorname{sign}(\varepsilon) = 1 \text{ with probability } \frac{1}{3} \\ \operatorname{sign}(0) = 0 \text{ with probability } \frac{1}{3} \\ \operatorname{sign}(-\varepsilon) = -1 \text{ with probability } \frac{1}{3} \end{cases}$$

Recall that

$$\mathbb{E}\left[\mathcal{L}\left(Y_{t+k}, r(Y_t)\right)\right] = \Pr\left[\operatorname{sign}\left(Y_{t+k} - Y_t\right) \neq \operatorname{sign}\left(r(Y_t) - Y_t\right)\right]$$

and consider the following notation:

$$S_{k,t} = \operatorname{sign}\left(Y_{t+k} - Y_t\right);$$

$$Sr = \text{sign}\left(r\left(Y_t\right) - Y_t\right).$$

Then,

$$\Pr\left[\text{sign}\left(Y_{t+k} - Y_t\right) \neq \text{sign}\left(r\left(Y_t\right) - Y_t\right)\right] = \Pr\left(S_{k,t} \neq Sr\right).$$

Let us calculate this probability:

$$\begin{aligned}\Pr\left(S_{k,t} \neq Sr\right) &= \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq Sr\right) \cap \left(Sr \in \left(i\right)\right)\right] \\ &= \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq i\right) \cap \left(Sr \in \left(i\right)\right)\right].\end{aligned}$$

Using the assumption of independence, we have

$$\Pr\left(S_{k,t} \neq Sr\right) = \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq i\right) \cap \left(Sr \in \left(i\right)\right)\right],$$

so

$$\begin{aligned}\Pr\left(S_{k,t} \neq Sr\right) &= \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq i\right) \cap \left(Sr \in \left(i\right)\right)\right] \\ &= \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq i\right)\right] \Pr\left[Sr \in \left(i\right)\right] \\ &= \frac{1}{3} \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq i\right)\right],\end{aligned}$$

because

$$\Pr\left[Sr \in \left(i\right)\right] = \frac{1}{3}, i = -1, 0, 1.$$

Therefore,

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{1}{3} \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \neq i\right)\right];$$

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{1}{3}\left\{\Pr\left[S_{k,t} \in (0,1)\right] + \Pr\left[S_{k,t} \in (-1,1)\right] + \Pr\left[S_{k,t} \in (-1,0)\right]\right\};$$

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{2}{3}\left[\Pr\left(S_{k,t} = -1\right) + \Pr\left(S_{k,t} = 0\right) + \Pr\left(S_{k,t} \in 1\right)\right];$$

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{2}{3}.$$

Therefore, if we take the testing function $h_t = 1$, then the null hypothesis,

$$H_0 : \mathbb{E}\left[\mathcal{L}\left(Y_{t+k}, r\left(Y_t\right)\right) - \mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right)\right]h_t = 0, \tag{21}$$

is equivalent to

$$H_0 : \mathbb{E}\left[\mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right)\right] = \frac{2}{3}.$$

Let

$$\mathrm{X}\left(Y_{t+k}, b_t\left(k\right)\right) = 1 - \mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right).$$

Then,

$$H_0 : \mathbb{E}\left[\mathrm{X}\left(Y_{t+k}, b_t\left(k\right)\right)\right] = \frac{1}{3}.$$

Notice that $\mathrm{X}(Y_{t+k}, b_t(k))$ is a Bernoulli random variable with an expected value equal to the probability of the c-bias succeeding in predicting the direction of future MPR changes. Under regularity assumptions (see last footnote at the end of this proof) and given k,

$$\frac{\sqrt{n}\,\frac{1}{n}\sum_{t=1}^{n}\left(X\left(Y_{t+k}, b_t\left(k\right)\right) = \frac{1}{3}\right)}{\bar{\sigma}} \xrightarrow{A} N\left(0,1\right);$$

$$\bar{\sigma}^2 = \lim_{n\to\infty} \mathrm{var}\left[\sqrt{n}\,\frac{1}{n}\sum_{t=1}^{n} \mathrm{X}\left(Y_{t+k}, b_t\left(k\right)\right)\right] < \infty.$$

We are not really interested in using the testing function $h_t = 1$. When using the relevant testing function,

$$h_t = \begin{cases} 1 & \text{if c-bias is not neutral} \\ 0 & \text{otherwise} \end{cases}$$

the null hypothesis in equation (21) is different. We notice that

$$\mathbb{E}\mathcal{L}\left(Y_{t+k}, r\left(Y_t\right)\right)h_t = \Pr\left\{\left[\mathrm{sign}\left(Y_{t+k} - Y_t\right) \neq \mathrm{sign}\left(r\left(Y_t\right) - Y_t\right)\right] \cap \left(h_t = 1\right)\right\},$$

which is simply the probability of making a mistaken forecast when the c-bias is not neutral. Let us consider the following set, C:

$$C = \left\{\text{c-bias}_t \neq \text{neutral}\right\} = \left\{h_t = 1\right\}.$$

Then,

$$\Pr\left\{\left[\mathrm{sign}\left(Y_{t+k} - Y_t\right) \neq \mathrm{sign}\left(r\left(Y_t\right) - Y_t\right)\right] \cap \left(h_t = 1\right)\right\} = \Pr\left[\left(S_{k,t} \neq Sr\right) \cap C\right],$$

and thus

$$\mathbb{E}\mathcal{L}\left(Y_{t+k}, r\left(Y_t\right)\right)h_t = \Pr\left[\left(S_{k,t} \neq Sr\right) \cap C\right].$$

We calculate the following probability:

$$\begin{aligned} \Pr\left[\left(S_{k,t} \neq Sr\right) \cap C\right] &= \sum_{i=-1}^{1} \Pr\left\{\left[\left(S_{k,t} \neq Sr\right) \cap C\right] \cap \left(Sr = i\right)\right\} \\ &= \sum_{i=-1}^{1} \Pr\left\{\left[\left(S_{k,t} \notin \left(i\right)\right) \cap C\right] \cap \left(Sr = i\right)\right\}. \end{aligned}$$

Using the assumption of independence, we have

$$\begin{aligned}\Pr\left[\left(S_{k,t} \neq Sr\right) \cap C\right] &= \sum_{i=-1}^{1} \Pr\left\{\left[\left(S_{k,t} \notin (i)\right) \cap C\right] \cap \left(Sr = i\right)\right\} \\ &= \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \notin (i)\right) \cap C\right] \Pr\left(Sr \in (i)\right) \\ &= \frac{1}{3} \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \notin (i)\right) \cap C\right].\end{aligned}$$

Thus,

$$\Pr\left[\left(S_{k,t} \neq Sr\right) \cap C\right] = \frac{1}{3} \sum_{i=-1}^{1} \Pr\left[\left(S_{k,t} \notin (i)\right) \middle| C\right] \Pr(C);$$

$$\begin{aligned}\Pr\left(S_{k,t} \neq Sr\right) = \frac{\Pr(C)}{3}\Big\{&\Pr\left[S_{k,t} \in (0,1) \middle| C\right] + \Pr\left[S_{k,t} \in (-1,1) \middle| C\right]; \\ &+ \Pr\left[S_{k,t} \in (-1,0) \middle| C\right]\Big\}\end{aligned}$$

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{2\Pr(C)}{3}\left[\Pr\left(S_{k,t} = -1 \middle| C\right) + \Pr\left(S_{k,t} = 0 \middle| C\right) + \Pr\left(S_{k,t} \in 1 \middle| C\right)\right];$$

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{2\Pr(C)}{3} \Pr\left[S_{k,t} \in (-1,0,1) \middle| C\right];$$

$$\Pr\left(S_{k,t} \neq Sr\right) = \frac{2\Pr(C)}{3} = \frac{2}{3}\mathbb{E}h_t$$

Therefore, the null hypothesis,

$$H_0 : \mathbb{E}\left[\mathcal{L}\left(Y_{t+k}, r\left(Y_t\right)\right) - \mathcal{L}\left(Y_{t+k}, b_t(k)\right)\right] h_t = 0,$$

is equivalent to

$$H_0 : \frac{2}{3}\mathbb{E}h_t = \mathbb{E}\left[h_t \mathcal{L}\left(Y_{t+k}, b_t(k)\right)\right].$$

Let

$$\mathrm{N}\left(Y_{t+k}, b_t\left(k\right)\right) = \frac{2}{3}h_t - h_t\mathcal{L}\left(Y_{t+k}, b_t\left(k\right)\right).$$

Then,

$$H_0\colon \mathbb{E}\left[\mathrm{N}\left(Y_{t+k}, b_t\left(k\right)\right)\right] = 0.$$

Under standard assumptions for the central limit theorem for dependent observations and given k,[3]

$$\frac{\sqrt{n}\,\frac{1}{n}\sum_{t=1}^{n}\left[\mathrm{N}\left(Y_{t+k}, b_t\left(k\right)\right)\right]}{\bar{\sigma}} \xrightarrow{A} N\left(0,1\right);$$

$$\bar{\sigma}^2 = \lim_{n\to\infty}\mathrm{var}\left\{\sqrt{n}\,\frac{1}{n}\sum_{t=1}^{n}\left[\mathrm{N}\left(Y_{t+k}, b_t\left(k\right)\right)\right]\right\} < \infty.$$

Appendix C: Bayesian Model Averaging

As Brock and Durlauf argue, the standard econometric approach in the literature relies on the choice of a particular model, M, which is considered a good approximation of the true model.[4] Given a data set, D, and the chosen model, M, estimates of the parameters of interest, β, and their variances can be obtained. The analogous Bayesian strategy involves the calculation of the posterior density of the parameter, $\mu(\beta|D, M)$.

Brock and Durlauf (and many others) analyze the problem of model uncertainty, which basically originates in the researcher's ignorance about the true model. Under this type of uncertainty, any estimate of the parameter

3. These standard assumptions require $\mathrm{N}(Y_{t+k}, b_t(k))$ to be a stationary ergodic mixingale with γ_m of size -1. A sequence $\{Z_t\}$, such that $\mathbb{E}Z_t^2 < \infty$, is a mixingale if we can find sequences of nonnegative numbers $\{a_t\}$ and $\{\gamma_m\}$ such that $\{\mathbb{E}[\mathbb{E}(Z_t|\mathcal{F}_{t-m})^2]\}^{1/2} \leq a_t\gamma_m$ and $\lim_{m\to\infty}\gamma_m = 0$, where $\{\mathcal{F}_t\}$ represents a filtration for which $\{Z_t\}$ is an adapted process. See White (2001) for further details.

4. Brock and Durlauf (2001).

of interest, β, will be conditioned on the particular choice of a model, M. Therefore, although the researcher is interested in the density $\mu(\beta|D)$, he or she will only be able to uncover $\mu(\beta|D, M)$.

To remove the problem of model uncertainty, the Bayesian approach proposes the definition of a space of possible models, $\mathcal{M}$. Integrating out the dependence of $\mu(\beta|D, M_m)$ on the particular model $M_m \in \mathcal{M}$ leads to the unconditional density $\mu(\beta|D)$. To do this, Bayes' theorem provides the following expression

$$\mu\left(\beta|D\right) = \sum_{M_m \in \mathcal{M}} \mu\left(\beta|D, M_m\right)\mu\left(M_m|D\right),$$

which reduces to

$$\mu\left(\beta|D\right) \propto \sum_{M_m \in \mathcal{M}} \mu\left(\beta|D, M_m\right)\mu\left(D|M_m\right)\mu\left(M_m\right),$$

where $\mu(D|M_m)$ is the likelihood of the data given the particular model $M_m \in \mathcal{M}$ and $\mu(M_m)$ represents the prior density defined over $\mathcal{M}$. Basically, these results show that the posterior density of the parameter β is a weighted average of the conditional densities of the parameter for different assumptions about the true model. This technique is called Bayesian model averaging (BMA).

Leamer provides expressions for the conditional expectation and variance of β given the data set D:[5]

$$\mathbb{E}\left(\beta|D\right) = \sum_{M_m \in \mathcal{M}} \mu\left(M_m|D\right)E\left(\beta|D, M_m\right),$$

$$\text{(22)} \qquad \begin{aligned} \text{var}\left(\beta|D\right) &= \mathbb{E}\left(\beta^2|D\right) - \left[\mathbb{E}\left(\beta|D\right)\right]^2, \\ &= \sum_{M_m \in \mathcal{M}} \mu\left(M_m|D\right)\text{var}\left(\beta|D, M_m\right) \\ &\quad + \sum_{M_m \in \mathcal{M}} \mu\left(M_m|D\right)\left[\mathbb{E}\left(\beta|D, M_m\right) - \mathbb{E}\left(\beta|D\right)\right]^2, \end{aligned}$$

5. Leamer (1978).

where

$$\mu\left(M_m \middle| D\right) = \frac{\mu\left(D \middle| M_m\right)\mu\left(M_m\right)}{\sum_{M_m \in \mathcal{M}} \mu\left(D \middle| M_m\right)\mu\left(M_m\right)}.$$

Therefore, the conditional variance of β given the data set D in equation (22) is broken down into two additive components: an intra-model variance and a cross-model variance.

The literature provides several approximations for the numerical implementation of the BMA technique. We use the Laplace approximation described by Volinsky and others:[6]

$$\log\left[\mu\left(D \middle| M_m\right)\right] \approx l - d_k \log\left(n\right), \tag{29}$$

where d_k represents the number of β parameters to estimate and l denotes the log-likelihood evaluated in the estimated parameters. Equation (23) is called the Bayesian information criterion (BIC) approximation, as shown by Hoeting.[7]

As Brock and Durlauf suggest, we compute estimates of $\mathbb{E}(\beta | D, M_m)$ and $\text{var}(\beta | D, M_m)$ using OLS and relying on a uniform prior.[8]

6. Volinsky and others (1997).
7. Hoeting and others (1999).
8. Brock and Durlauf (2001).