



University of Pretoria
Department of Economics Working Paper Series

Inflation Forecasts and Forecaster Herding: Evidence from South African Survey Data

Christian Pierdzioch
Helmut-Schmidt-University
Monique B. Reid
Stellenbosch University
Rangan Gupta
University of Pretoria
Working Paper: 2014-55
October 2014

Department of Economics
University of Pretoria
0002, Pretoria
South Africa
Tel: +27 12 420 2413

On the Directional Accuracy of Inflation Forecasts: Evidence from South African Survey Data

October 2014

Abstract

We study the directional accuracy of South African survey data of short-term and longer-term inflation forecasts. Upon applying techniques developed for the study of relative operating characteristic (*ROC*) curves, we find evidence that forecasts contain information with respect to the subsequent direction of change of the inflation rate.

JEL classification: C53, D82, E37

Keywords: Inflation rate; Forecasting; Directional Accuracy

Addresses:

* Christian Pierdzioch, Helmut-Schmidt-University, Department of Economics, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany. Email: c.pierdzioch@hsu-hh.de.

Monique B. Reid, Stellenbosch University, Department of Economics, Private Bag X1, Matieland, South Africa, 7602. Email: mreid@sun.ac.za.

Rangan Gupta, University of Pretoria, Department of Economics, Pretoria, 0002, South Africa. Email: rangan.gupta@up.ac.za.

* Corresponding author.

1 Introduction

Inflation expectations play a key role in explaining inflation dynamics and, thus, are of key interest for monetary policy and anyone who monitors inflation. Because the Phillips-curve equation implies that inflation expectations are likely to feed into actual inflation at some stage (Nunes 2010, Adams and Pantula 2011, Fuhrer 2012, among others), it is not surprising that economists have developed sophisticated econometric techniques to study the formation and the dynamics of inflation expectations. Knowing how inflation expectations are formed and whether inflation expectations contain information as to where the inflation rate is heading is particularly important for South Africa, given that it has been targeting inflation between 3-6% from the beginning of 2000.

Because survey data of inflation forecasts are a rich data source for studying inflation expectations it is important to study the accuracy of survey data of inflation forecasts. The accuracy of forecasts can be measured along various dimensions. One dimension is the forecast horizon. We study inflation forecasts at two different horizons: short-term forecasts and longer-term forecasts. Another dimension that economists have been particularly interested in concerns the unbiasedness and rationality of inflation forecasts. In the case of South African survey data of inflation forecasts, Ehlers and Steinbach (2007) find for inflation expectations from the Bureau for Economic Research (BER) survey and data from the Reuters Inflation Expectations (RIE) survey that the expectations of financial analysts and short-term expectations by the trade unions appear unbiased. In contrast, their findings suggest that the expectations of the business representatives and the longer-term expectations of trade unions are biased. Moreover, Ehlers and Steinbach (2007) find that all three groups use information inefficiently, implying that weakly rationality of forecasts can be rejected. The forecasts for the current-quarter horizon from the RIE Survey are the only exception. In a similar vein, Kabundi and Schaling (2013) argue that inflation expectations violate the rational expectation hypothesis because expectations are closely tied to the lagged inflation rate. In a follow-up study, Kabundi et al. (2014) confirm that the forecasts of business representatives and trade unions are linked to the lagged inflation rate. Only the forecasts of financial analysts appear to rely relatively less on past inflation, indicating heterogeneity across different groups of forecasters. Pierdzioch et al. (2014) confirm heterogeneity of forecasts for short-term inflation forecasts taken from Bloomberg and, in addition, report that forecasts are subject to a behavioral bias due to forecaster herding. Forecaster herding arises if forecasters do not publish unbiased forecasts but rather orient their forecasts towards the forecasts of others (consensus forecast). Moreover, forecaster herding appears to dominate in times of high inflation volatility.

Yet another often studied dimension of the accuracy of forecasts concerns their directional accuracy (Schnader and Stekler 1990, Stekler 1994, Kolb and Stekler 1996, Ash et al. 1998, Pons 2001,

Sinclair et .al. 2010, to name just a few). Unlike other measures of forecast accuracy, the directional accuracy of forecasts is likely to be closely correlated with the economic value of forecasts (Leitch and Tanner 1991; on the economic evaluation of directional forecasts, see also Blaskowitz and Herwartz 2011). Moreover, the directional accuracy of inflation forecasts is a concept highly relevant for any central bank, but especially for one that targets inflation, like the South African Reserve Bank.¹ “A central bank wants to know if inflation will accelerate or decelerate to decide if the interest rate should be raised or lowered.” (Öller and Barot 2000, p. 306). We, thus, use survey data on short-term inflation forecasts taken from Bloomberg and longer-term inflation forecasts taken from the BER surveys to re-examine the informational content of forecasts for subsequent developments of the inflation rate. While market-timing tests have been widely studied in earlier research to study the directional accuracy of forecasts, we apply techniques that have been developed for the analysis of relative operating characteristic (*ROC*) curves to study the directional accuracy of inflation forecasts. *ROC* curves have been extensively studied in disciplines like biostatistics and medicine, but have gained popularity in economics only in recent years (Berge and Jordà 2011, Lahiri and Wang 2013, Liu and Moench 2014, among others). A *ROC* curve visualizes the directional accuracy of forecasts by confronting the rate of true signals (sensitivity) with the rate of false signals (one minus specificity) for alternative values of a decision criterion. The decision criterion discriminates between signal and nonsignals. For our data, we define a signal as a forecast of a rising inflation rate and a nonsignal as a forecast of an unchanged or falling inflation rate. Hence, the decision criterion may specify that a signal occurs if the forecast of the inflation rate exceeds the actual inflation rate at the time when a forecast is being made. Such a choice of the decision criterion would correspond to the kind of decision criterion typically used to set up widely-studied market-timing tests. Market-timing tests, however, are typically based only on a single decision criterion, and setting the decision criterion to zero obviously is an arbitrary choice. A *ROC* curve, therefore, summarizes information for a wide range of values of the decision criterion, each leading to different signals and nonsignals and, therefore, to a different assessment of the directional accuracy of inflation forecasts. Based on our analysis of *ROC* curves, we find that inflation forecasts contain significant information with regard to the direction of change of the inflation rate. This results holds both for the CPI inflation rate and the CPIX (i.e., CPI excluding mortgage costs) inflation rate, and for short-term and longer-term forecasts (constant one-year forecast horizon).

We organize the remainder of our research as follows. We first describe our empirical method in Section 2. We then summarize our empirical results in Section 3, where we also describe our data on short-term and long-term inflation rates. We conclude in Section 4.

¹The South African Reserve Bank had moved to an official inflation targeting regime, within a band of 3% to 6%, since the February of 2000.

2 Empirical Method

Our empirical analysis aims at testing the directional accuracy of period t -forecasts of the inflation rate. The direction of change of the inflation rate is defined as the difference between the future inflation rate, π_{t+1} , and the inflation rate, π_t , in the period in which a forecast is being formed. Accordingly, we define an “event” as a rise in the inflation rate between period of time t and period of time $t + 1$. A “nonevent” occurs if the inflation rate does not change or falls between period of time t and period of time $t + 1$. Hence, we define

$$\text{Event: } \pi_{t+1} > \pi_t, \tag{1}$$

$$\text{Nonevent: } \pi_{t+1} \leq \pi_t. \tag{2}$$

We denote by $\pi_{t,t+1}^e$ a period- t forecast of the inflation rate in $t + 1$. In order to trace out whether forecasts signal the direction of change of the inflation rate, we assume that a “signal” occurs whenever the difference between a period- t forecast and the latest information on the (real-time) period- t inflation rate is not less than a decision criterion, $c \in (-\infty, +\infty)$. The complementary “nonsignal” occurs whenever the difference between a period- t forecast and the period- t inflation rate is less than a decision criterion. We have

$$\text{Signal: } \pi_{t,t+1}^e - \pi_t \geq c, \tag{3}$$

$$\text{Nonsignal: } \pi_{t,t+1}^e - \pi_t < c. \tag{4}$$

Sensitivity, SE , measures the ratio of true signals relative to all events, and *specificity*, SP , measures the ratio of true nonsignals relative to all nonevents. It follows from the definition of signals and nonsignals, that the sensitivity and specificity of forecasts of the inflation rate are functions of the decision criterion, c . Accordingly, we define ($\#$ = number of counts)

$$SE(c) = \frac{\#(\pi_{t,t+1}^e - \pi_t \geq c \cap \pi_{t+1} - \pi_t > 0)}{\#(\pi_{t+1} - \pi_t > 0)}, \tag{5}$$

$$SP(c) = \frac{\#(\pi_{t,t+1}^e - \pi_t < c \cap \pi_{t+1} - \pi_t \leq 0)}{\#(\pi_{t+1} - \pi_t \leq 0)}. \tag{6}$$

If $c \rightarrow \infty$, then no signals occur and $SE(c) = 0$ and $SP(c) = 1$, such that the rate of false signals is $1 - SP(c) = 0$. Similarly, if $c \rightarrow -\infty$, then $SE(c) = 1$ and $1 - SP(c) = 1$. Hence, upon letting c vary from a very small to a very large value, we get a *ROC* curve that starts in a $[0, 1] \times [0, 1]$ unit quadrant with the rate of false signals on the horizontal axis and the rate of true signals on the vertical axis at the point $[0, 0]$ and then increases until it approaches the point $[1, 1]$. The following four cases summarize how a *ROC* curve can connect the points $[0, 0]$ and $[1, 1]$:

1. Forecasts of the inflation rate do not perform better than a pure noise signal irrespective of the choice of the decision criterion: $SE(c) = 1 - SP(c) \forall c$. \Rightarrow A *ROC* curve equals the bisecting line.
2. Forecasts of the inflation rate perform better than a pure noise signal irrespective of the choice of the decision criterion: $SE(c) > 1 - SP(c) \forall c$ (except at the endpoints, where always $SE(c) = 1 - SP(c)$). \Rightarrow A *ROC* curve settles above the bisecting line.
3. Forecasts of the inflation rate perform worse than a pure noise signal irrespective of the choice of the decision criterion: $SE(c) < 1 - SP(c) \forall c$. \Rightarrow A *ROC* curve settles below the bisecting line.
4. Forecasts of the inflation rate perform better than a pure noise signal for some choices of the decision criterion and worse for other choices: $SE(c) < 1 - SP(c)$ for some c and $SE(c) > 1 - SP(c)$ for other c . \Rightarrow A *ROC* curve crosses the bisecting line.

The area, *AUROC*, under a *ROC* curve summarizes the directional accuracy of forecasts of the inflation rate. If forecasts of the inflation rate perfectly predict the direction of changes of the inflation rate, then $AUROC = 1$. In contrast, if the directional accuracy of forecasts of the inflation rate is the same as the directional accuracy of a pure noise signal (Case 1), then $AUROC = 0.5$. Finally, if forecasts of the inflation rate systematically predict the wrong sign of subsequent changes of the inflation rate (Case 3), then $AUROC < 0.5$. In this case, however, we could reverse the definition of a “signal” such that again $AUROC > 0.5$

The *AUROC* statistic can be estimated using a non-parametric approach that uses the fact that the *AUROC* statistic is linked to the Wilcoxon-Mann-Whitney U statistic (Bamber 1975, Hanley and McNeil 1982). Following Greiner et al. (2000, page 38–39), we compute

$$AUROC = \frac{n_0 n_1 - U}{n_0 n_1}, \quad (7)$$

where $n_0 = \#$ nonevents, $n_1 = \#$ events, and $U = R - 0.5n_0(n_0 + n_1)$ = two-sample Mann-Whitney rank-sum test with $R =$ rank sum of the nonevents. The standard error (SE) of the *AUROC* statistic is given by (see Hanley and McNeil 1982 and Greiner et al. 2000)

$$SE = \sqrt{\frac{A + B + C}{n_0 n_1}}, \quad (8)$$

where $A = AUROC(1 - AUROC)$, $B = (n_1 - 1)(Q_1 - AUROC^2)$, $C = (n_0 - 1)(Q_2 - AUROC^2)$, with $Q_1 = AUROC/(2 - AUROC)$ and $Q_2 = 2AUROC^2/(1 + AUROC)$.

3 Empirical Analysis

3.1 Short-term Forecasts

Monthly survey data on forecasts of the South African inflation rate were taken from Bloomberg. As for the timing of events, Bloomberg sends questionnaires to the forecasters, the forecasters then submit their forecasts, Bloomberg collects the incoming forecasts and publishes them as it receives them. Finally, Bloomberg simultaneously publishes all forecasts on the so called observation date, which typically proceeds the official release of real-time data on inflation. Forecasts are for the year-on-year CPI inflation rate and the year-on-year CPIX inflation rate (excluding mortgage costs). The forecasts are short-term forecasts because, for example, forecasts form forecasts in May of the inflation rate released in June. The CPI data cover the sample period 2000/05–2014/06, and the CPIX data are available for the sample period 2000/05–2008/12. The data is available as an unbalanced panel because not all forecasters participated in all surveys. In case of the CPI inflation rate, 82 forecasters published a total of 2,691 forecasts, where mean of the number of forecasts per forecaster is 33 (rounded) and the median is 17. For the CPIX inflation rate, we have available 1,418 forecasts, where the mean of forecasts per forecaster is 23 (rounded) and the median is 16. We compute events and signals using real-time data on inflation to proxy forecasters' information set. Moreover, central banks and other users of forecasts make decisions in real time, and so it is likely that they are interested in the directional accuracy of real-time signals and events.

– Please include Figure 1 about here. –

Figure 1 plots the CPI inflation data (Panel A) and the CPIX inflation data (Panel B). The figure shows the fluctuations of the inflation rate along with the 3%–6% inflation-targeting band, the number of forecasts available in every month, and the cross-sectional forecast errors (mean, maximum, minimum). The figure shows that the inflation rate climbed above the upper boundary of the inflation-targeting band by several percentage points in 2002/2003 and 2007/2008. As for the rise in the inflation rate in 2002/2003, it is important to note that South Africa experienced a currency crisis in autumn/winter 2001. The exact causes of the currency crisis were controversial at the time and actually lead the President of South Africa to appoint the Myburgh Commission of Inquiry to investigate the causes of the rand depreciation during the fourth quarter of 2001. Bhundia and Ricci (2005) conclude that there is little evidence that contagion from Argentina was the cause of the depreciation. Instead, a range of political and economic factors were identified (an acceleration in money growth and exchange rate overshooting, the South African Reserve Bank's

net open forward book, the delay in the privatisation of *Telkom*², an announcement by the the South African Reserve Bank that it would tighten exchange controls, a slowdown of global economic activity). The rise in the inflation rate in 2007/2008 reflects a spike in food and energy prices. At that time, a national power crisis (Eskom, the electricity provider in South Africa experienced major capacity constraints at that time) led to a sharp increase in the price of electricity. Finally, the impact of the financial crisis added to that.

Despite the two marked peaks in the inflation rate, Figure 1 shows that the inflation rates returned relatively quickly to lower levels and also stayed within the inflation-targeting band for extended periods of time (for example, during 2003–2008). This quick return of the inflation rates to lower levels leads to the question of whether the South African Reserve Bank improved its inflation performance during the inflation targeting period despite volatile economic and global conditions. There indeed is evidence that the South African Reserve Bank managed to establish a substantial degree of credibility and transparency using the repo rate to control consumption and investment spending and inflation expectations. For example, Reid (2009a, 2009b) finds that inflation expectations in South Africa are relatively well anchored, and South Africa compares favourably in comparative studies of central bank transparency (Aron and Muellbauer 2007).

Figure 1 also shows that the number of forecasts increased over time, but this increase was not strictly monotonic. The population of forecasters, thus, is heterogeneous with respect to the frequency of participation in the survey. Some forecasters participated in the survey more often than others. Forecasters are also heterogeneous with regard to the sign and magnitude of forecast errors. While the forecast errors (defined as the actual the inflation rate minus the cross-sectional average of forecasts) hovered around zero, the minimum / maximum of the forecasting error defined as the difference between the actual inflation rate and the cross-sectional maximum / minimum of forecasts often depart to a substantial extent from the average forecast error. Figure 2 illustrates the cross-sectional heterogeneity of forecast errors for the CPI inflation data (Panel A) and the CPIX inflation data (Panel B). The figure shows histograms of the cross-sectional average forecast error and the maximum / minimum cross-sectional forecast error.

– Please include Figure 2 and 3 about here. –

Figure 3 plots the *ROC* curves for the CPI inflation data and the CPIX inflation data. The *ROC* curves have been estimated on the full sample of data, pooling inflation forecasts across time

²Telkom Group Limited is a wireline and wireless telecommunications provider in South Africa, and operates in more than 38 countries across Africa. Telkom is a semi-privatised, (39%) state-owned company.

and across forecasters. The shape of both *ROC* curves clearly illustrates that inflation forecasts contain substantial information with regard to subsequent changes of the inflation rate. The *ROC* curves settle above the bisecting line, showing that forecasts of the inflation rate outperform a pure noise signal. Moreover, the better performance of inflation forecasts relative to a pure noise signal does not depend on the specific choice of the decision criterion, *c*.

– Please include Table 1 about here. –

Table 1 summarize the estimates of the *AUROC* statistic. The table also depicts the standard error of the *AUROC* statistic, and the number of nonevents and events (increases in the inflation rate). The key message conveyed by the table is that a test of the null hypothesis that inflation forecasts are indistinguishable from a pure noise signal yields highly significant results. Hence, inflation forecasts contain useful information with regard to the sign of subsequent changes of the inflation rate.

In addition to the full sample results, we present results for several subsamples. As a first subsample period, we consider the period up to and including 2008/12 because the South African Reserve Bank at that time switched from targeting CPIX inflation to targeting CPI inflation. Next, we consider subsample periods during which the inflation rate fluctuated within (outside) of the official 3%–6% inflation-targeting band. Finally, we consider subsamples that only use forecasts from those forecasters who published not less than a minimum number of forecasts. The minimum number of forecasts is defined in terms of the full-sample mean (median) number of forecasts per forecaster. These two subsamples account for the fact for the cross-sectional heterogeneity of forecasters with regard to the number of forecasts that they published. The results for the various subsamples corroborate the results for the full sample of data. For all subsamples, we find that inflation forecasts predict the sign of the change of the inflation rate from the current to the next month.

– Please include Figure 4 about here. –

Figure 3 plots rolling-window estimates of the *AUROC* statistic. For both the CPI inflation rate and the CPIX inflation rate, the *AUROC* statistic shows a slight tendency to decline. The *AUROC* statistic, however, always stays above the benchmark line of 0.5 (and significantly so). We conclude that the directional accuracy of inflation forecasts is not concentrated in a few subsample periods, but that directional accuracy holds throughout the sample period being analyzed.

A *ROC* curve shows all possible combinations of sensitivity and one minus specificity, where the combinations depend on the decision criterion, c . It is interesting to ask how a user of inflation forecasts should set the decision criterion, c . Alternatively, one might ask which point on the *ROC* a user of inflation forecasts should choose. Several approaches are available to answer this question (for a survey, see Greiner et al. 2000). A simple approach is to maximize the Youden index (Youden 1950). The Youden index is defined as, $Y(c) = Se(c) + Sp(c) - 1$, and the c is chosen such that $c_Y = \max_c Y(c)$. For our data, the maximum of the Youden index obtains for a decision criterion of $c = 0.1$ (estimated with high precision). Hence, users may find it optimal not to interpret a forecast exceeding the actual inflation rate as a signal of an increase in the inflation rate, but rather a forecast that exceeds the actual inflation rate by 0.1 percentage points.

3.2 Longer-term Forecasts

A natural question is whether the good directional accuracy of short-term inflation forecasts also holds for longer-term inflation forecasts. In order to study this question, we use the BER inflation data. The BER forecasts are available at a quarterly frequency. Forecasters belong to one of the following three groups: financial sector, business sector, and trade unions. Unlike the case of the Bloomberg data, individual forecasts are not available, that is, the data are available only at the sector level. The sample period is 2000Q3–2014Q2. Forecasters predict the CPIX inflation rate from 2000Q3 to 2008Q4 and the CPI inflation rate from 2009Q1–2014Q2, reflecting the change in 2009 of the proxy for inflation that was officially targeted by the South African Reserve Bank. Accordingly, we use as predictand the CPIX inflation rate up to the end of 2008 and the CPI inflation rate thereafter. Like Reid (2012), we treat both the actual inflation rate and the inflation-forecast series as continuous time series because, as witnessed by Figure 5, there is no apparent structural break. As in the case of short-term forecasts, we compute events and signals using real-time data on inflation.

– Please include Figure 5 about here. –

The original BER survey data consists of calendar year forecasts, implying that the forecast horizon changes from survey to survey. In Q1 respondents are asked to forecast a larger proportion of the current year than in Q3 – the horizons of the forecasts change over the year. We construct from the data a series with a constant one-year forecast horizon. To this end, we add in each quarter a fraction of the current year to a fraction of the following year, depending on the quarter in which the survey is being conducted. For example, the 2001Q4 survey (for which the field work is

conducted in October) asked respondents to forecast what the CPIX inflation rate would be at the end of 2001 and at the end of 2002. So the one-year forecast constructed for this study consists of 25% of the current year’s forecast and 75% of the following year’s forecast. Figure 5 shows the constructed one-year inflation forecasts for the three groups of forecasters.

– Please include Figure 6 about here. –

Figure 6 plots the resulting *ROC* curves. For constructing the *ROC* curves, we define an event as an increase in the inflation rate between, for example, 2000Q4 and 2001Q4. Moreover, we assume that forecasters in period of time t have access to the inflation rate in period of time $t - 1$, and a “signal” of an increasing inflation rate occurs if the period- t one-year horizon forecast exceeds the period- $(t - 1)$ inflation rate. For example, we compare the forecast of the inflation rate made in 2001Q1 with the actual inflation rate in 2000Q4. Matching events and signals implies that the effective full sample period (after removing missing data) is 2001Q4–2014Q2. As compared to the *ROC* curves shown in Figure 3, the *ROC* curves estimated on the BER data are less smooth due to the smaller number of observations. The *ROC* curves settle above the bisecting line, suggesting that inflation forecasts have predictive power with regard to the direction of change of the inflation rate over a one-year forecast horizon.

– Please include Table 2 about here. –

The results summarized in Table 2 show that the *AUROC* statistics reflect the shapes of the shape of the *ROC* curves estimated for the three groups of forecasters. The point estimate of the *AUROC* statistic for the financial analysts is somewhat larger than the point estimate for the other two groups of forecasters. The *AUROC* statistics are significant for all three groups of forecasters.

Because of the small sample size and the inherent overlapping data structure we also use a bootstrap simulation to study the significance of the *AUROC* statistic. To this end, we resample blocks of the data using the block bootstrap proposed by Politis and Romano (1994) (for an application of the block bootstrap to the study of directional forecasts in the presence of serial correlation, see Blaskowitz and Herwartz 2014). The block bootstrap accounts for the temporal dependence of the data, where we set the mean length of a sampled block of data to one year. We use 1,000 bootstrap simulation runs and compute for every simulation run the *AUROC* statistic. The simulated standard errors of the *AUROC* statistic (Table 2) are somewhat larger for the business sector and trade unions than those computed by means of Equation (8), but for all three groups of forecasters the simulated confidence intervals show that forecasts contain information with respect to the direction of change of the inflation rate.

4 Concluding Remarks

In earlier studies of inflation forecasts for South Africa, researchers have focused on (deviations from) rationality, forecast biases, and strategic interactions among forecasters. In contrast, we have focused on the directional accuracy of forecasts. While (deviations from) rationality, forecast biases, and strategic interactions among forecasters certainly are important features of inflation forecasts for South Africa, our results suggest that inflation forecasts are still useful for predicting the direction of change of the inflation rate. In this respect, we note that our empirical results are about being able to forecast the direction of change of inflation at certain forecast horizons rather than about the causal effects of inflation expectations on actual inflation, et vice versa.

We have derived our results by applying techniques of *ROC*-analysis. These techniques have not been widely studied in economics. In future research, it will be interesting to apply techniques of *ROC*-analysis not only to study the directional accuracy of inflation forecasts, but also to forecasts of, for example, business-cycle expansions and recessions. While researchers have applied techniques of *ROC*-analysis to study U.S. business-cycle fluctuations (Berge and Jordà 2011, Liu and Moench 2014), it is interesting to explore whether the techniques also are useful for studying business-cycle fluctuations in other developed and developing countries.

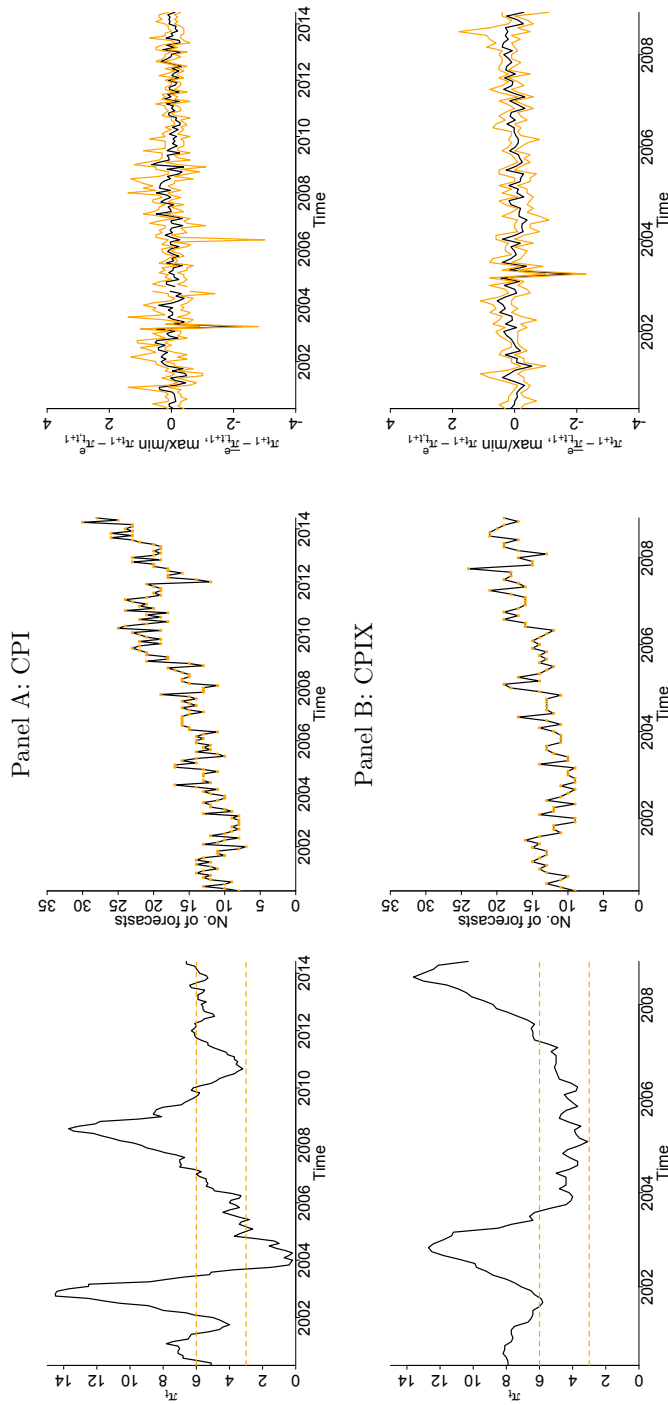
References

- Aron, J., J. Muellbauer (2007). Review of Monetary Policy in South Africa Since 1994. *Journal of African Economies* 16: 705–744.
- Adam K, M. Padula (2011). Inflation Dynamics and Subjective Expectations in the United States. *Economic Inquiry* 49: 13–25.
- Ash, J. C. K., Smyth, D. J., S. M. Heravi (1998). Are OECD Forecasts Rational and Useful? A Directional Analysis. *International Journal of Forecasting* 14: 381 – 391.
- Bamber, J. D. (1975). The Area Above the Ordinal Dominance Graph and the Area Below the Receiver Operating Characteristic Graph. *Journal of Mathematical Psychology* 12: 387–415.
- Berge, T. J., Ò. Jordà (2011). Evaluating the Classification of Economic Activity Into Recessions and Expansions. *American Economic Journal: Macroeconomics* 3: 246–277.
- Blaskowitz, O., H. Herwartz (2011). On Economic Evaluation of Directional Forecasts. *International Journal of Forecasting* 27: 1058–1065.
- Blaskowitz, O., H. Herwartz (2014). Testing the Value of Directional Forecasts in the Presence of Serial Correlation. *International Journal of Forecasting* 30: 30–42.

- Bhundia, A. J., L. A. Ricci (2005). The Rand Crises of 1998 and 2001: What Have We Learned? In Nowak, M., L A. Ricci (Eds.), *Post-Apartheid South Africa: The First Ten Years*. Washington, DC: International Monetary Fund. 156–173
- Ehlers, N., M. R. Steinbach (2007). The Formation of Inflation Expectations in South Africa. Working Paper WP/07/06, South African Reserve Bank.
- Fuhrer J. (2012). The Role of Expectations in Inflation Dynamics. *International Journal of Central Banking* 81: 137–165.
- Greiner, M., D. Pfeiffer, R. D. Smith (2000). Principles and Practical Application of the Receiver-Operating Characteristic Analysis for Diagnostic Tests. *Preventive Veterinary Medicine* 45: 23-41.
- Hanley, J. A., B. J. McNeil (1982). The Meaning and Use of the Area Under a Receiver Operating Characteristic Curve. *Radiology* 143: 29–36.
- Kabundi, A., E. Schaling (2013). Inflation and Inflation Expectations in South Africa: An Attempt at Explanation. *South African Journal of Economics* 81: 346–355.
- Kabundi, A., Schaling, E., M. Some (2014). Monetary Policy and Heterogeneous Inflation Expectations in South Africa. Working Paper No. 422, Economic Research Southern Africa.
- Kolb, R. A., H. O. Stekler (1996). How Well Do Analysts Forecast Interest Rates? *Journal of Forecasting* 15: 385–394.
- Lahiri, K., and J. G. Wang (2013). Evaluating Probability Forecasts for GDP Declines Using Alternative Methodologies. *International Journal of Forecasting* 29: 175–190.
- Leitch, G., T. E. Tanner (1991). Economic Forecast Evaluation: Profits versus the Conventional Error Measures. *American Economic Review* 81: 580–590.
- Nunes. R. (2010). Inflation Dynamics: The Role of Expectations. *Journal of Money, Credit, and Banking* 42: 1161–1172.
- Öller, L.E., B. Barot (2000). The Accuracy of European Growth and Inflation Forecasts. *International Journal of Forecasting* 16: 293–315.
- Pierdzioch, C., Reid, M. B., R. Gupta (2014). Inflation Forecasts and Forecaster Herding: Evidence from South African Survey Data. Department of Economics, University of Pretoria, Working Paper No. 2014-55.
- Politis, D. N., J. P. Romano (1994). The Stationary Bootstrap. *Journal of the American Statistical Association* 89: 1303–1313.
- Pons, J. (2001). The Rationality of Price Forecasts: A Directional Analysis. *Applied Financial Economics* 11: 287–290.
- Liu, W., E. Moench (2014). What Predicts U.S. Recessions? Staff Report No. 691, Federal Reserve Bank of New York Staff Reports.
- R Development Core Team (2014). R: A Language and Environment for Statistical Computing. Internet address: <http://www.R-project.org>.

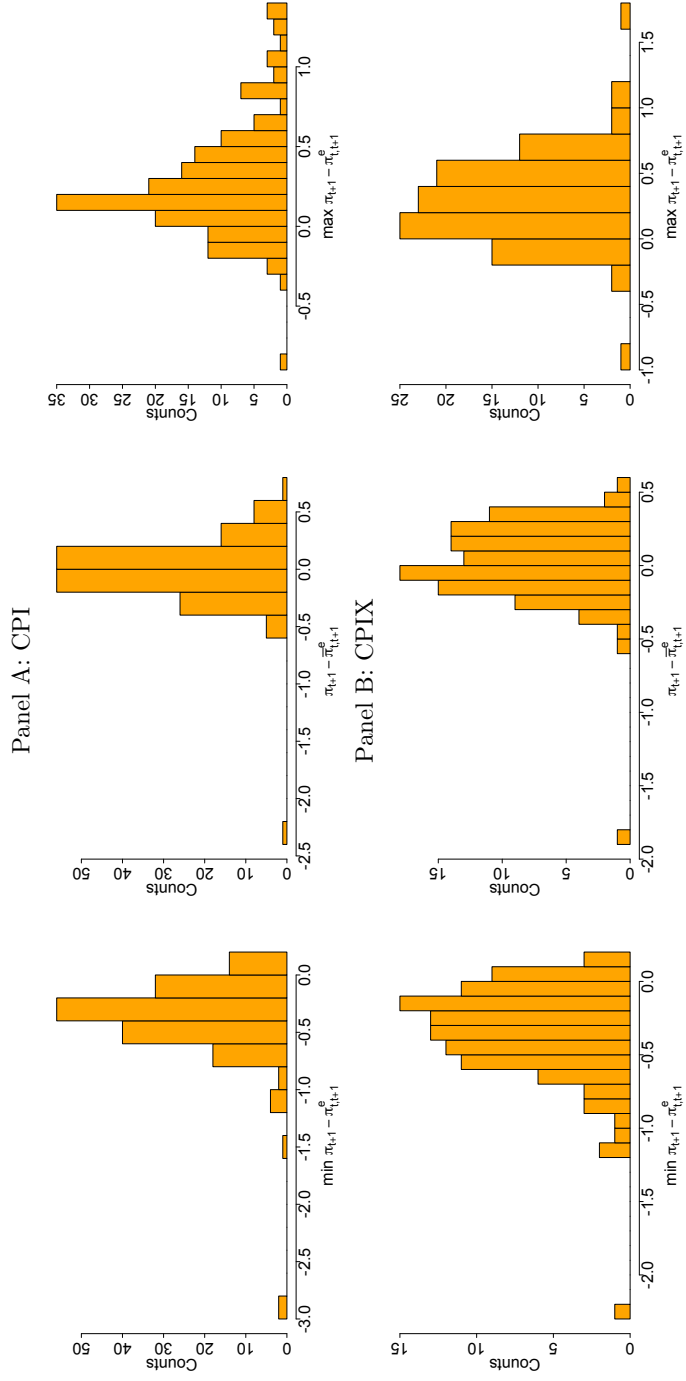
- Reid, M. B. (2009a). Isolating a Measure of Inflation Expectations for the South African Financial Markets Using Forward Interest Rates. *South African Journal of Economics* 77: 399–413.
- Reid, M. B. (2009b). The Sensitivity of South African Inflation Expectations to Surprises. *South African Journal of Economics* 77: 413–429.
- Reid, M. B. (2012). Inflation Expectations of the Inattentive General Public. Working Paper No. 278, Economic Research Southern Africa.
- Schnader, M. H., H. O. Stekler (1990). Evaluating Predictions of Change. *Journal of Business* 63:99–107.
- Sinclair, T. M., Stekler, H. O., L. Kitzinger (2010). Directional Forecasts of GDP and Inflation: A Joint Evaluation With an Application to Federal Reserve Predictions. *Applied Economics* 42: 2289–2297.
- Stekler, H. O. (1994). Are Economic Forecasts Valuable? *Journal of Forecasting* 13: 495–505.
- Youden, W. J. (1950). Index for Rating Diagnostic Tests. *Cancer* 3: 32–35.

Figure 1: Time Series of Short-Term Forecasts




Note: Dashed horizontal lines = boundaries of the inflation-targeting band of 3% - 6%. Until the end of 2008, the inflation target was defined in terms of the CPI. Since 2009, the inflation target has been defined in terms of the CPI. Since the introduction of inflation targeting, the inflation target has been defined in terms of a 3% - 6% band for longer-term CPIX / CPI inflation. $\pi_{t+1} - \pi_{t,t+1}^e$ = actual inflation rate minus the cross-sectional average of forecasts. $\pi_{t+1} - \max / \min \pi_{t,t+1}^e$ = actual inflation rate minus the cross-sectional maximum / minimum of forecasts. This figure and all other figures and results were computed using the free programming environment R (R Development Core Team 2014).

Figure 2: Histograms of Short-Term Forecast Errors



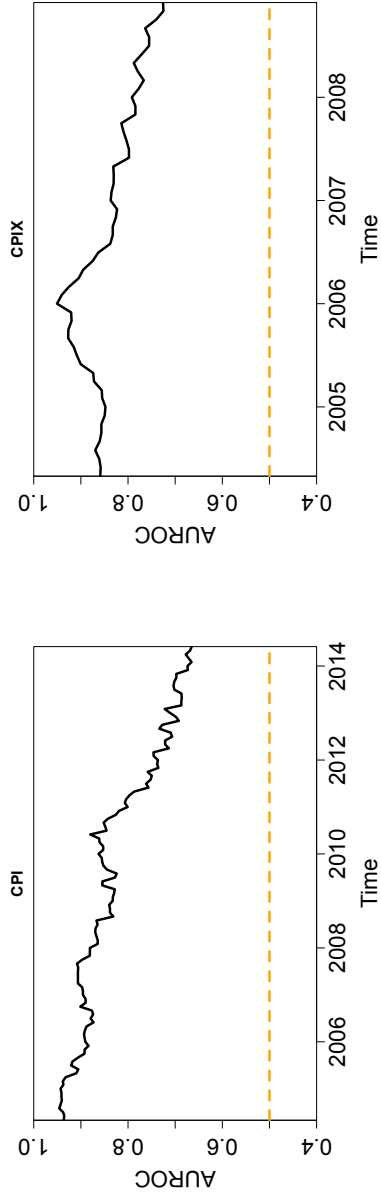
Note: $\pi_{t+1} - \pi_{t+1}^e$ = actual inflation rate minus the cross-sectional average of forecasts. $\max \pi_{t+1} - \pi_{t+1}^e$ = actual inflation rate minus the cross-sectional maximum of forecasts. $\min \pi_{t+1} - \pi_{t+1}^e$ = actual inflation rate minus the cross-sectional minimum of forecasts.

Figure 3: ROC Curves for Short-Term Forecasts



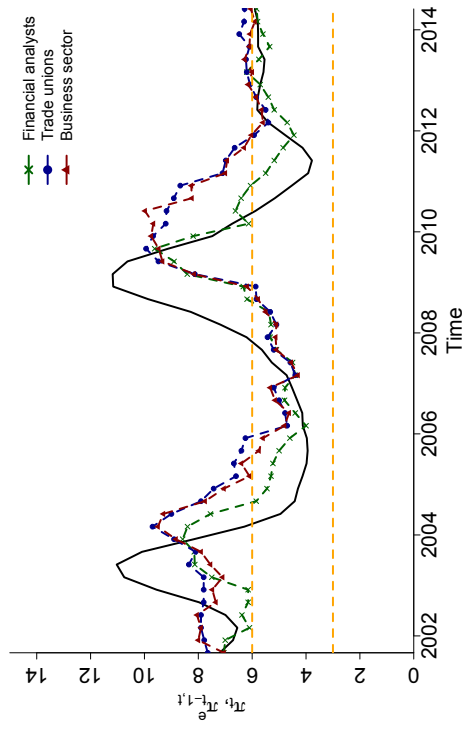
Notes: Bisecting line = ROC curve when forecasts do not contain any information with regard to the sign of subsequent changes of the inflation rate.
Dark line = actual ROC curve implied by forecasts.

Figure 4: Rolling-Window Estimates of the *AUROC* Statistic for Short-Term Forecasts



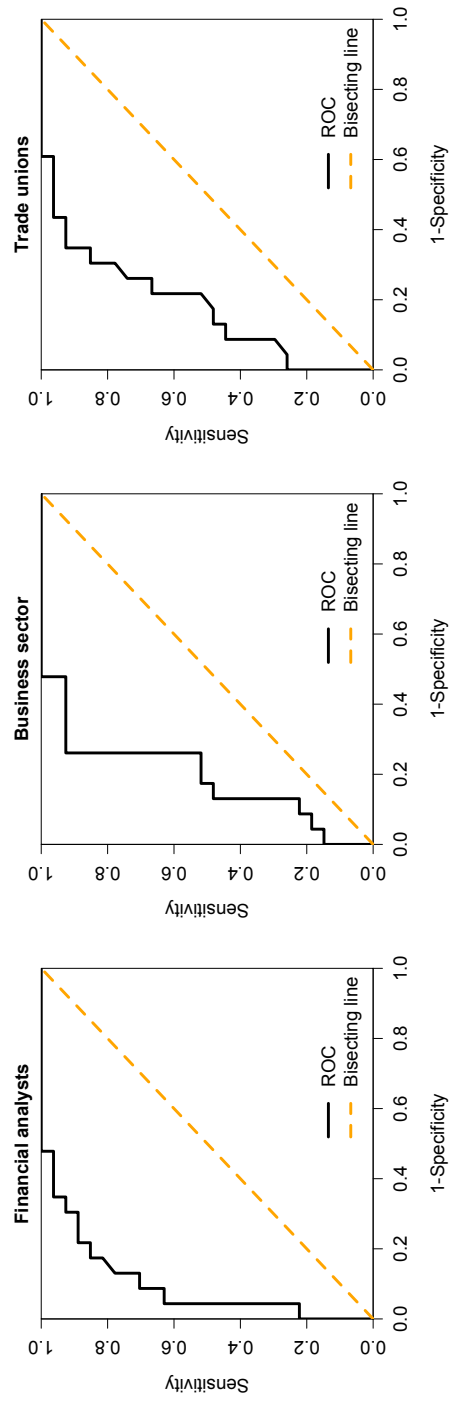
Notes: Dashed horizontal line = 0.5 (benchmark value if inflation forecasts do not predict the sign of the subsequent change of the inflation rate). Solid line = *AUROC* statistic computed using a rolling-estimation window. The length of the rolling-estimation window is 48 months.

Figure 5: Longer-Term Forecasts



Note: Solid black line = CPIX inflation rate from 2000Q3 to 2008Q4 and the CPI inflation rate from 2009Q1–2014Q2. Dashed horizontal lines = boundaries of the inflation-targeting band of 3% – 6%. The forecasts are for a constant one-year forecasting horizon.

Figure 6: ROC Curves for Longer-Term Forecasts



Notes: Bisecting line = ROC curve when forecasts do not contain any information with regard to the sign of subsequent changes of the inflation rate. Dark line = actual ROC curve implied by forecasts. The forecasts are for a constant one-year forecasting horizon. The inflation rate is defined as the CPIX inflation rate from 2000Q3 to 2008Q4 and the CPI inflation rate from 2009Q1-2014Q2.

Table 1: *AUROC* Statistic

Inflation	Subsample period	<i>AUROC</i>	Std. Error	No. of events	No. of nonevents
CPI	Full sample	0.8954***	0.0064	1,399	1,284
CPI	2000/05 – 2008/12	0.9126***	0.0081	592	692
CPI	2009/01 – 2014/06	0.8754***	0.0101	792	592
CPI	Inside the 3% – 6% band	0.8670***	0.0098	754	687
CPI	Outside the 3% – 6%band	0.9019***	0.0093	664	568
CPI	Min(forecasts) \geq 33	0.8935***	0.0074	1,047	990
CPI	Min(forecasts) \geq 17	0.8979***	0.0066	1,254	1,168
CPIX	Full sample	0.8588***	0.0099	667	743
CPIX	2000/05 – 2008/12	0.8547***	0.0081	592	692
CPIX	Inside the 3% – 6% band	0.8401***	0.0162	332	299
CPIX	Outside the 3% – 6%band	0.8820***	0.0122	354	412
CPIX	Min(forecasts) \geq 23	0.8577***	0.0111	526	594
CPIX	Min(forecasts) \geq 16	0.8577***	0.0106	581	662

Note: Subsample periods were defined according to the measure of inflation the South African Reserve Bank targeted. Until the end of 2008, the inflation target was defined in terms of the CPIX. Since 2009, the inflation target has been defined in terms of the CPI. Since the introduction of inflation targeting, the inflation target has been defined in terms of a 3% – 6% band for longer-term CPIX / CPI inflation. Min(forecasts) = For constructing the herding statistic, only forecasters were included in the sample who published not less than a minimum number of forecasts. The minimum number of forecasts was defined in terms of the full-sample mean (median) number of forecasts per forecaster. *** = significance at the 1% level of a test of the null hypothesis that inflation forecasts are indistinguishable from a pure noise signal.

Table 2: *AUROC* Statistic for Longer-Term Forecasts

Sector	<i>AUROC</i>	Std. Error	Boot Std. Error	Boot. 95% CI	Boot. 90% CI	No. of events	No. of nonevents
Financial analysts	0.9042***	0.0437	0.0385	(0.8188, 0.9730)	(0.8381, 0.9617)	23	27
Business sector	0.8132***	0.0605	0.0656	(0.6639, 0.9226)	(0.6944, 0.9060)	23	27
Trade unions	0.8237***	0.0589	0.0664	(0.6612, 0.9321)	(0.6975, 0.9151)	23	27

Note: Estimates are for the full sample period 2001Q4–2014Q2 (after removing missing data). The forecasts are for a constant one-year forecasting horizon. The inflation rate is defined as the CPIX inflation rate from the beginning of the sample to 2008Q4 and the CPI inflation rate from 2009Q1–2014Q2. Boot. Std. Error = standard error computed based on the bootstrapped sampling distribution of the *AUROC* statistic. Boot. 95% CI = Boundaries of a 95% confidence interval computed based on the percentiles of the bootstrapped sampling distribution of the *AUROC* statistic. The bootstrap was implemented by means of a stationary bootstrap (number of simulation runs = 1,000). Smoothing parameter of the stationary bootstrap set such that the mean block length is equal to one year. *** = significance at the 1% level of a test of the null hypothesis that inflation forecasts are indistinguishable from a pure noise signal.