Forecasting metal prices: Do forecasters herd?

Christian Pierdzioch

Jan-Christoph Rülke

Georg Stadtmann

European University Viadrina Frankfurt (Oder)

Department of Business Administration and Economics

Discussion Paper No. 325

September 2012

ISSN 1860 0921

Forecasting metal prices: Do forecasters herd?

Christian Pierdzioch[a,*], Jan-Christoph Rülke[b], and Georg Stadtmann[c]

July 2012

Abstract

We analyze more than 20,000 forecasts of nine metal prices at four different forecast horizons. We document that forecasts are heterogeneous and report that anti-herding appears to be a source of this heterogeneity. Forecaster anti-herding reflects strategic interactions among forecasters that foster incentives to scatter forecasts around a consensus forecast.

JEL classification: G17; C33; L61
Keywords: Metal prices; Forecasting; Forecaster (anti-)herding

[a] *Helmut Schmidt University, Department of Economics, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany.*
[b] *Department of Economics, WHU – Otto Beisheim School of Management, Burgplatz 2, 56179 Vallendar, Germany.*
[c] *European-University Viadrina, P.O.B. 1786, 15207 Frankfurt (Oder), Germany, and Deutsches Institut für Wirtschaftsforschung (DIW Berlin), Mohrenstraße 58, 10117 Berlin, Germany.*

* Corresponding author. Tel.: +49 40 6541 2879. fax: +49 40 6541 3808.
E-mail addresses: c.pierdzioch@hsu-hh.de (C.Pierdzioch), Jan-C.Ruelke@whu.edu (J.C. Rülke), Stadtmann@europa-uni.de (G. Stadtmann).

We are grateful for the financial support received through the foundation "Geld and Währung" from the Deutsche Bundesbank (S126/10081/11).

1. Introduction

Metals are crucial imported input factors for many industrialized countries, and they are a major source of export revenues for some developing countries. Large swings in metal prices can have a large impact on the terms of trade. Corporate managers and policymakers, therefore, closely track changes in metal prices.[1] Moreover, researchers spend much effort to forecast future price trends. Forecasting future trends in metal prices, however, has turned out difficult as metal prices have experienced recently substantial swings and sharp price reversals. The media are full of reports that blame speculative trading activities and herding of market participants as major sources of significant price swings and market rallies.[2] A natural question is whether such herding – to the extent that it occurred – was driven by herding of metal-price forecasters. Forecaster herding arises if forecasters ignore their private information and instead follow the forecasts of others (Scharfstein and Stein 1990, Froot et al. 1992).

We implement a robust empirical test developed by Bernhardt et al. (2006) to study whether metal-price forecasters do, in fact, herd. This test is easy to implement, it is robust to various forms of misspecification, and it delivers results that can easily be interpreted in economic terms. In order to implement the test, we study more than 20,000 forecasts of nine metal prices, including forecasts of the prices of Gold and Silver. Forecasts are available at four different forecast horizons for a sample period that covers more than 15 years of data (1995 – 2011). Across all nine metal prices and all four forecasting horizons, we do not find signs of forecaster herding. On the contrary, we find strong evidence of forecaster *anti*-herding. Our findings are in line with the mounting evidence of forecaster anti-herding that has been documented in recent literature for the forecasts of stock analysts (Naujoks et al. 2009), fiscal forecasts (Stadtmann et al. 2011),

---

[1]See, for example, the United Nations (2011) report: *"G20 Study Group on Commodities"*.

[2]See, for example, Arends (2010), Schindler (2011), and Monk (2012) for media reports studying potential fundamental and non-fundamental (bubble and herding) determinants of the gold price.

and oil-price forecasts (Pierdzioch et al. 2010). To our knowledge, evidence of forecaster anti-herding has not been reported in earlier literature for forecasts of metal prices.

Laster et al. (1999) have developed a widely studied model that illustrates why forecasters anti-herd. In their model, two groups of customers buy forecasts. The first group of customers buys forecasts regularly. This group is interested in accurate forecasts and, thus, buys forecasts from a forecaster who has delivered the most accurate forecasts over a longer time period. The second group of customers, in contrast, buys forecasts occasionally. This group of customers buys from a forecaster who provided the best forecast in the last period. The decision to buy forecasts only occasionally may be a simple heuristic, or it may be the result of a rational benefit-cost analysis. For example, movements of metal prices may have only a moderate impact on the business of the second group of customers and the costs of monitoring the accuracy of forecasts may be higher for this group than for the first group of customers. Because forecasters' profit function consists of revenues from both groups of customers, forecasters do not deliver the most accurate forecast. If the second group of customers dominates, forecasters have a strong incentive to differentiate their forecasts from the forecasts of others. The strong incentive to differentiate forecasts arises because, in case a forecaster delivers an "extreme" forecast, the number of other forecasters who deliver the very same "extreme" forecast is small. Thus, even though an "extreme" forecast may have a small probability of being accurate, the expected payoff of such a forecast can be high because, in the case of such a stroke of luck, a forecaster does not have to share with others revenues from the second group of customers. If a forecaster would publish a less extreme forecast that is close to the consensus forecast, in contrast, the probability is high that other forecasters make similar forecasts, implying that many forecasts come close to the "best" forecast. If this is the case, even an excellent forecast is likely to have only a rather moderate effect on a forecaster's income and reputation.

In earlier literature, researchers have focused on aspects of metal markets that significantly differ from the aspect of forecaster (anti-)herding, which is the focus of our empirical study. For example, much research has been undertaken to shed light on the speculative efficiency of metal markets (see the survey by Watkins and McAleer 2004). In an early study of the London Metal Exchange, Canarella and Pollard (1986) analyze whether futures prices are unbiased predictors of future spot prices. Sephton and Cochrane (1990) further study the efficiency of the London Metal Exchange by means of single-market and multiple-market models that employ the dynamics of forward and spot ("prompt") prices. Other researchers have focused on the cointegration of spot and metal futures prices (Brenner and Kroner 1995, Chow 1998, among others). Hsieh and Kulatilaka (1982) analyze whether forward metal prices equal market participants' expectations of future spot prices. Instead of using survey data on forecasts of metal prices, they use econometric models to proxy the dynamics of expectations. Dooley and Lenihan (2005) and Ahti (2009) show that time-series-based econometric models may be useful to forecast metal prices.

We organize the remainder of our study as follows: In Section 2, we describe our data set. In Section 3, we illustrate the test for forecaster (anti-)herding that we used in our empirical analysis. In Section 4, we report our empirical results. In Section 5, we offer some concluding remarks.

2. Theoretical background and data

We study monthly survey data of price forecasts for nine metals compiled by Consensus Economic Forecast (CEF) for the time period 1995−2011. The survey is conducted during the first week of a month and released at the beginning of the second week. We study forecasts of the prices of the following metals: Aluminium, Cobalt, Copper, Gold, Lead, Nickel, Platinum, Silver, and Uranium. Forecasts are available at four different forecasting horizons: one month, one quarter, one year, and two years. We thus can study short-term, medium-term, and long-term forecasts. We supplement the forecasts with the realized values of the metal prices as well as the forward rates (sourced from Datastream), where the latter are based on data for the first week of each month and the time horizon matches those of the forecast. Table 1 summarizes information on the sample means of forecasts and realizations of metal prices, the correlation between the consensus forecast and the forward rate, the number of forecasts, the number of forecasters, and the sample period for which forecasts are available. In total, we can analyze 20,464 forecasts.

Please insert Table 1 about here.

The CEF survey data contain information not only on individual forecasts, but also information on the company or institutions at which forecasters work.[3] Because this information allows the performance of the forecasting company to be evaluated, the accuracy of forecasts may affect the reputation of forecasters. Reputation may strengthen if forecasts are accurate, and this may give rise to less "extreme" forecasts and herding of forecasters. Alternatively, it may happen that concerns regarding forecaster reputation give rise to a scattering of forecasts. Such a scattering of forecasts arises, for example, if a "superstar" effect is at work that strengthens incentives to make extreme forecasts in an attempt to differentiate forecasts from the forecasts of others. If such forecast differentiation is prevalent in the forecasting industry, the result is anti-herding of forecasters.

Scharfstein and Stein (1990, p. 476) argue that a "superstar" effect arises if, for example, top-ranking forecasters receive a disproportionately high reputation and income. Similarly, Rosen (1981, p. 845) argues that two constituent features of the "superstar" effect are *"first, a close connection between personal reward and the size of one's market; and second, a strong tendency for both market size and reward to be skewed to the most talented people in the activity"*. Laster et al. (1999) develop a formal model of forecaster

[3] The forecasters work for investment banks, commercial banks, consultancies, and in the automotive industry. A complete list of participants is available upon request from the authors. For more information, see www.consensuseconomics.com.

anti-herding in which these two constituent features of a "superstar" effect are at work. In their model forecasters are rewarded not only for forecast accuracy, but also for giving the best forecast at a single point in time. The latter component of forecaster income gives rise to a scattering of forecasts and, thus, forecaster anti-herding. In their model, forecasters' profit function can be represented as follows:

$$\Pi = -\alpha(s_{t+k} - E_{i,t}[s_{t+k}])^2 + (1-\alpha)\left[\frac{\Sigma}{n} \text{ if } E_{i,t}[s_{t+k}] = s_{t+k}, \ 0 \text{ else}\right], \tag{1}$$

where Π = profit from forecasting, $E_{i,t}[s_{t+k}]$ = forecast of forecaster i made in period t, s_{t+k} = realization of the metal price being forecasted, and $0 \leq \alpha \leq 1$ is a weighting parameter. The quadratic term on the right-hand side represents the profits from making an accurate forecast. Accordingly, any deviation of the metal price from the forecast lowers profits. The term in brackets on the right-hand side captures that a forecaster can win an amount of Σ in the case of an exact forecast, where this amount is divided among all those forecasters, n, who made such an exact forecast. If the forecast turns out to be incorrect, the term in brackets is zero. The second term of the profit function, thus, implies a close connection between a forecasters' income and the size of the market, where income is skewed to the most talented forecaster ($1-\alpha$ and Σ are large, and n is small).

The two elements of the profit function represent the profits from two groups of customers. The first group of customers consists of intensive forecast users who are interested in accurate forecasts. The profit from selling forecasts to this group of customers increases in the accuracy of forecasts. The second group of customers consists of occasional forecast users. In Equation (1), a forecaster receives profits from this group of customers only in the case of an exact forecast. Laster et al. (1999, p. 297) motivate this modeling choice as follows: *"The motivation for modeling the competition for occasional users as winner-takes-all is the media attention given to the forecaster who, in a given period, proves to be the most accurate among those participating in a survey. This publicity enhances a forecaster's reputation, credibility, and name recognition among* *occasional users, who are otherwise unfamiliar with most forecasters."*

The higher the relative importance of the second group of customers is $(1-\alpha)$, and the higher the revenues from these customers (Σ) are, the stronger is the incentive to make an "extreme" forecast. If a forecaster makes an "extreme" forecast, the probability of winning part of the revenues, Σ, from the second group of customers is low. At the same time, however, the number, n, of other forecasters who make the very same "extreme" forecast is small. As a consequence, forecast differentiation (that is, anti-herding) can lead to an increase in a forecaster's expected profit.

Forecaster anti-herding and the concomitant scattering of forecasts should result in cross-sectional heterogeneity of forecasts. Figure 1 shows that such a cross-sectional heterogeneity of forecasts, in fact, is a characteristic feature of the CEF survey data. The figure shows the actual metal prices (solid lines) and the range of the one-month-ahead forecasts (shaded areas). The range of forecasts measures the cross-sectional heterogeneity of forecasts and is defined as the maximum minus the minimum of forecasts at a given point in time.[4] The empirical results we shall document in Section 4 indicate that forecaster anti-herding contributes to the cross-sectional heterogeneity of forecasts. It is important to mention, however, that our empirical results do not quantify the extent to which forecaster anti-herding contributes to the cross-sectional heterogeneity of forecasts.

[4]A similar cross-sectional heterogeneity of forecasts has also been reported for commodity prices. See Pierdzioch et al. (2010) for an empirical analysis of the cross-sectional heterogeneity of oil-price forecasts.

Figure 1: Actual Metal Prices and Forecast Ranges

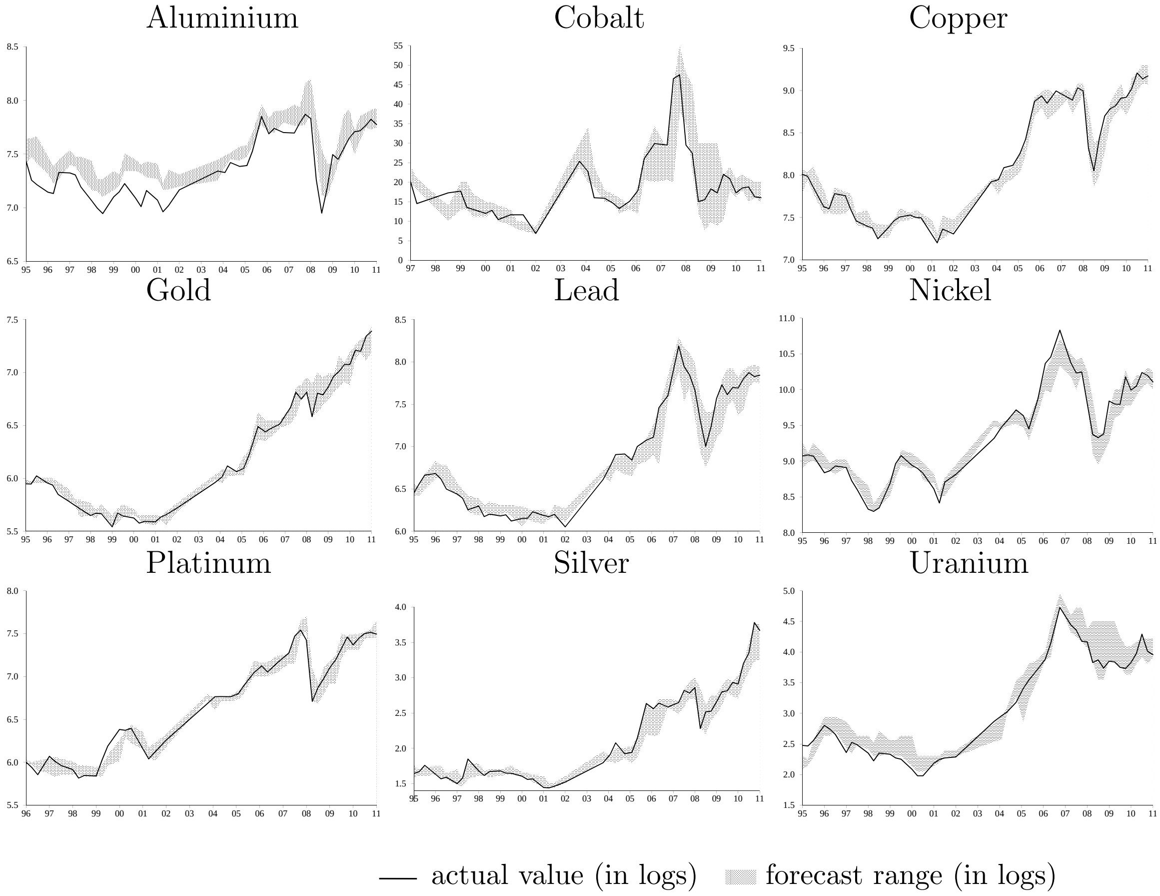

Note: Figure 1 shows the range (shaded area) of the forecasts and the actual value of the metal prices (solid lines).

Forecaster anti-herding can have a negative effect on the informational quality of metal-price forecasts. While forecaster anti-herding need not distort the average of forecasts (Laster et al. 1999, p. 306), forecaster anti-herding inflates the cross-sectional heterogeneity of forecasts. As a result, forecasts of metal prices give, for an outside observer, a more dispersed and, thus, less precise account of expected future movements of metal prices than it would be the case if private sector forecasters delivered unbiased forecasts. In terms of empirical research, forecaster anti-herding implies that results of standard panel tests of unbiased forecasts are difficult to interpret. Rejection of the hypothesis of unbiased forecasts may reflect deviations from forecaster rationality, or they may reflect rational biases due to forecaster anti-herding. Forecaster anti-herding, thus, limits the informational content of tests of the informational efficiency of the process of forecast formation.

Finally, Figure 1 shows that several metal prices rose more or less steadily until around 2006, while substantial dynamics occurred during the recent past. One should also note that Consensus Economics discontinued to survey metal prices between September 2002 and April 2004. Hence, there are no forecasts available for this period of time.

3. Testing for forecaster (anti-)herding

We use a test that has recently been proposed by Bernhardt et al. (2006) to analyze whether forecasters (anti-)herd. Their test is easy to implement, and the economic interpretation of the test results is straightforward. In order to lay out the economic intuition that motivates their test, it is useful to consider a forecaster i who forms an efficient private forecast, $E^P_{i,t}[s_{t+k}]$, of a future metal price in period $t+k$, derived from an optimal forecasting model and all information available in period t when the forecast is being made. The private forecast, thus, is (median) unbiased and the probability that the unbiased private forecast overshoots or undershoots the future metal price should be 0.5.

The published forecast, however, need not to be identical to the unbiased private forecast. For example, the published forecast, $E_{i,t}[s_{t+k}]$, made by forecaster i differs from the private forecast when a forecaster at least in part ignores the private forecast and instead follows the forecasts of others. The forecasts of others can be represented by the so called "consensus" forecast, $\tilde{E}_t[s_{t+k}]$, that is, the average forecast made by all forecasters, at a given point in time. In the case of forecaster herding, the published forecast is biased towards the consensus forecast, $\tilde{E}_t[s_{t+k}]$. In the case the private forecast, $E^P_{i,t}[s_{t+k}]$, exceeds the consensus forecast, $\tilde{E}_t[s_{t+k}]$, the published forecast, thus, is smaller than the private forecast, implying $E^P_{i,t}[s_{t+k}] > E_{i,t}[s_{t+k}] > \tilde{E}_t[s_{t+k}]$. As a result, the probability of undershooting is smaller than 0.5. Similarly, if the private forecast is smaller than the consensus forecast, we have $E^P_{i,t}[s_{t+k}] < E_{i,t}[s_{t+k}] < \tilde{E}_t[s_{t+k}]$, implying that the probability that the future metal price overshoots the published forecast is also smaller than 0.5.

In contrast, in the case of forecaster anti-herding, the published forecasts is farther away from the consensus forecast than the private forecast. For example, if the private forecast falls short of the consensus forecast, we have $E_{i,t}[s_{t+k}] < E^P_{i,t}[s_{t+k}] < \tilde{E}_t[s_{t+k}]$, implying that the probability that the future metal price overshoots the published forecast is larger in the case of anti-herding than in the case in which a forecaster publishes an unbiased forecast. The probability of undershooting, thus, is larger than 0.5. Similarly, if the private forecast exceeds the consensus forecast, we have $E_{i,t}[s_{t+k}] > E^P_{i,t}[s_{t+k}] > \tilde{E}_t[s_{t+k}]$, implying that the probability of overshooting is larger than in the case in which a forecaster publishes an unbiased forecast. It follows that, if the private forecast is smaller than the consensus forecast, the probability that the future metal price overshoots the published forecast is also larger than 0.5.

The probabilities of undershooting and overshooting can be used to set up a simple test of forecaster (anti-)herding. The null hypothesis is that published forecasts of metal prices are unbiased (no herding or anti-herding). The probability, P, that an unbiased forecast of a future metal price, $E_{i,t}[s_{t+k}]$, made by forecaster i overshoots (undershoots) the future realization of the metal price, s_{t+k}, should then be 0.5, regardless of the consensus forecast, $\tilde{E}_t[s_{t+k}]$. As a result, the conditional probability of undershooting in case a forecast exceeds the consensus forecast should be

$$P(s_{t+k} < E_{i,t}[s_{t+k}] \,|\, E_{i,t}[s_{t+k}] > \tilde{E}_t[s_{t+k}], s_{t+k} \neq E_{i,t}[s_{t+k}]) = 0.5. \tag{2}$$

The conditional probability of overshooting in the case that an unbiased forecast is smaller than the consensus forecast should be

$$P(s_{t+k} > E_{i,t}[s_{t+k}] \,|\, E_{i,t}[s_{t+k}] < \tilde{E}_t[s_{t+k}], s_{t+k} \neq E_{i,t}[s_{t+k}]) = 0.5. \tag{3}$$

In case a forecaster herds, the published forecasts are closer to the consensus forecast than in the case of unbiased forecasts. The published forecasts, thus, are biased towards the consensus forecast. For those forecasts that exceed the consensus forecast, the probability of undershooting thus is less than 0.5. Similarly, biased published forecasts that are less than the consensus forecast imply a probability of overshooting that is also less than 0.5. We have

$$P(s_{t+k} < E_{i,t}[s_{t+k}] \,|\, E_{i,t}[s_{t+k}] > \tilde{E}_t[s_{t+k}], s_{t+k} \neq E_{i,t}[s_{t+k}]) < 0.5. \tag{4}$$

$$P(s_{t+k} > E_{i,t}[s_{t+k}] \,|\, E_{i,t}[s_{t+k}] < \tilde{E}_t[s_{t+k}], s_{t+k} \neq E_{i,t}[s_{t+k}]) < 0.5. \tag{5}$$

In the opposite case of forecaster anti-herding, the published forecasts is farther away from the consensus forecast than in the case of unbiased forecasts. If forecasters anti-herd, the two conditional probabilities, thus, are larger than 0.5. In this case, we have

$$P(s_{t+k} < E_{i,t}[s_{t+k}] \,|\, E_{i,t}[s_{t+k}] > \tilde{E}_t[s_{t+k}], s_{t+k} \neq E_{i,t}[s_{t+k}]) > 0.5. \tag{6}$$

$$P(s_{t+k} > E_{i,t}[s_{t+k}] \,|\, E_{i,t}[s_{t+k}] < \tilde{E}_t[s_{t+k}], s_{t+k} \neq E_{i,t}[s_{t+k}]) > 0.5. \tag{7}$$

In order to test for forecaster (anti-)herding, Bernhardt et al. (2006) suggest to compute a test statistic, S, which is defined as the average of the sample estimates of the two conditional probabilities. Unbiased forecasts imply $S = 0.5$, herding implies $S < 0.5$, and anti-herding implies $S > 0.5$. A test of the null hypothesis $S = 0.5$ can be set up by using the result that the test statistic, S, asymptotically has a normal sampling distribution.

Bernhardt et al. (2006) show that the test statistic, S, has a number of interesting properties. First, it is robust to phenomena like correlated forecast errors and market-wide shocks. The robustness of the test statistic is due to the fact that it is defined as the average of the conditional probabilities of overshooting and undershooting. A market-wide shock, for example, that drives metal prices up increases the conditional probability of overshooting and decreases the conditional probability of undershooting, leaving the average of the two conditional probabilities unaffected. Second, the averaging of the two conditional probabilities also implies that the test statistic, S, yields reliable results in case forecasters do not target the median but the mean of an asymmetric distribution over future metal prices. More generally, the averaging of the two probabilities makes the test statistic robust to systematic biases in forecasts unrelated to forecaster (anti-)herding. Third, the test statistic is robust to outliers in the data, data entry errors, or sharp trend reversals in metal prices. The robustness of the test statistic results from the fact that the conditional probabilities are computed as the relative frequencies of events from a large number of forecasts. Finally, the test statistic is conservative insofar as its variance attains a maximum under the null hypothesis of unbiased forecasts, implying that it is more difficult to reject the null hypothesis of unbiased forecasts when we should do so.

4. Empirical results

Table 2 depicts the S-statistic and the upper and lower bounds of a confidence interval for the four different forecasting horizons that we analyze. The key finding conveyed by the table is that forecasters do not herd. Rather, our main finding is that forecasters anti-herd. In the cases of all nine metal prices, the S-statistic exceeds the value of 0.5 that it would assume if forecasters delivered unbiased forecasts. There are only two exceptions: for long-term Gold and Silver forecasts (at a forecasting horizon of two years) the null hypothesis of unbiased forecasts cannot be rejected. In all other cases, there is statistically significant evidence of forecaster anti-herding.

Please insert Table 2 about here

A potential problem could arise because forecasters who contribute to the CEF survey data publish their forecasts simultaneously. This simultaneity is in contrast to the study by Bernhardt et al. (2006), who analyze sequentially published forecasts of stock analysts. The simultaneous publication of forecasts could imply that forecasters do not know the consensus forecast when deciding on the forecast that they want to publish. It is, thus, important to assess whether a simultaneity bias beleaguers our results. A simultaneity bias cannot arise if we use forward metal prices to measure the consensus forecast. Forward metal prices inform forecasters about market-wide expectations of future metal prices. Importantly, forward metal prices are publicly known, implying that forecasters can take them into account when forming their forecasts. In addition, forward metal prices closely track the consensus forecast. Finally, forward metal prices are available for the four different forecasting horizons that we studied in our empirical analysis.

Insert Table 3 about here.

Table 3 summarizes the findings that we obtain when we use forward metal prices to measure the consensus forecast. In all cases, the forward metal prices match the forecasting horizon of the CEF survey data. The findings confirm those shown in Table 2. There is not a single case in which the null hypothesis of unbiased forecasts cannot be rejected. In all cases, we find strong evidence of forecaster anti-herding − across all nine metal prices and all four forecasting horizons.

Given the large price swings and sharp price reversals experienced by metal prices during our sample period, we also analyze the variation over time in the S-statistic. In economic terms, it may be the case that, if metal prices mushroom, demand for forecasts increases, providing strengthening incentives to anti-herd. Alternatively, demand for forecasts may increase in periods of falling metal prices, with implications for forecaster (anti-)herding. Fluctuations in the S-statistic should signal changes in the prevalence of forecaster anti-herding. In order to study changes in the prevalence of forecaster anti-herding and, thus, changes in the S-statistic, we use a rolling-window estimation approach. Every rolling-estimation window represents two years of data. When we move the rolling estimation window forward in time, we drop the data at the beginning of the rolling window and add new data at the end of the rolling window. We then continue this rolling estimation process until we reach the end of our sample period. Figure 2 plots the S-statistics and the resulting 99% confidence bands.

Figure 2: Rolling-estimation window of the (Anti-)Herding Statistic, S_t

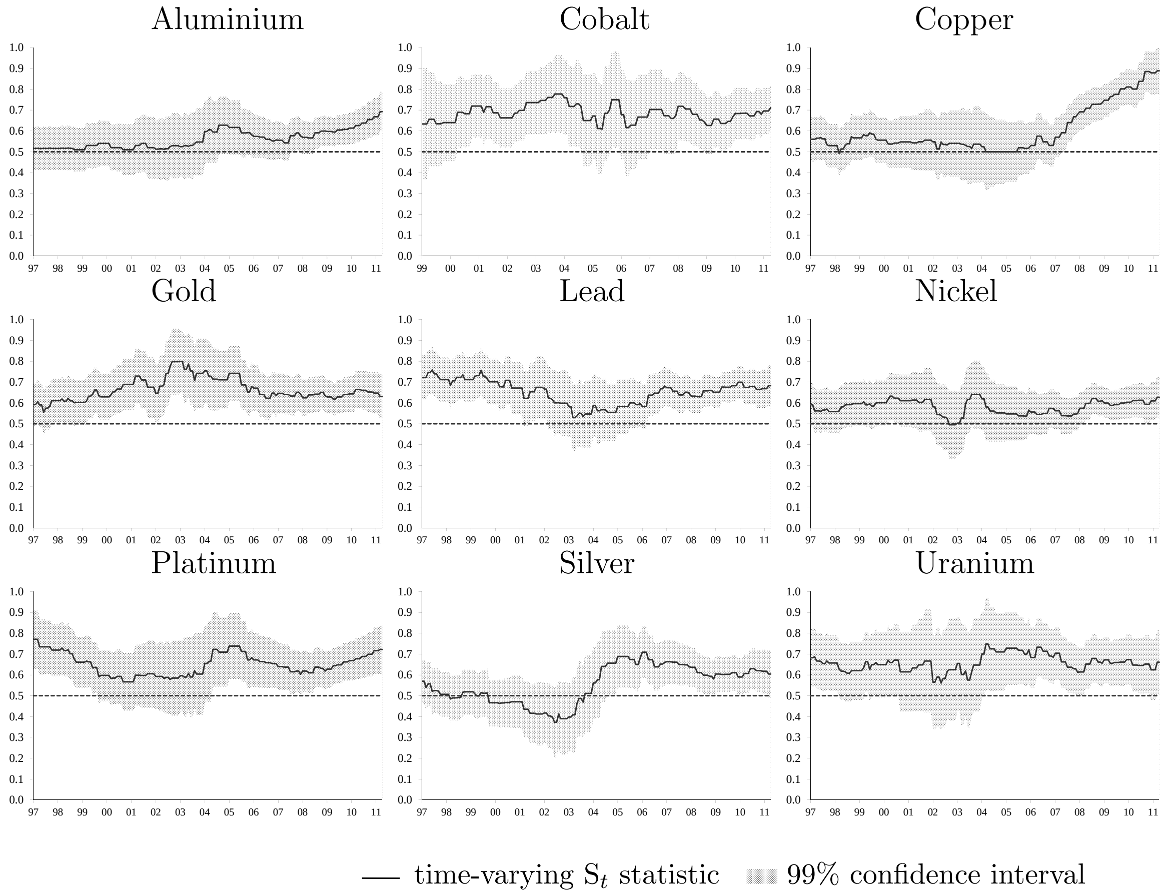

Note: Figure 2 shows the S-statistic (solid line) and the 99% confidence interval (shaded area) based on a two-year rolling-estimation window.

The results suggest that the S-statistic never drops below the 0.5 reference line (unbiased forecasts) in a statistically significant way. It is also evident that there are fluctuations of the S-statistic over time. In particular, the results of the rolling-estimation window analysis indicate that forecaster anti-herding was somewhat less prevalent in 2001−2003, and that it became more prevalent since then.[5] It also seems that, when analyzed through the lens of a rolling-estimation window, anti-herding was strongest over time in the cases of Cobalt, Gold, Lead, Platinum, and Uranium.

Insert Tables 4 and 5 about here.

[5] Two issues should be taken into account when interpreting the result that anti-herding was less prevalent in 2001−2003. First, as evidenced by Figure 1, metal prices rose more or less steadily during this period of time. Second, Consensus Economics discontinued to survey metal prices between September 2002 and April 2004. Hence, there are no forecasts available for this period of time.

As yet another robustness check, we analyze whether optimism and pessimism among forecasters affect our finding of forecaster anti-herding. To this end, we define optimists (pessimists) as forecasters whose forecasts imply a positively (negatively) sloped term structure of metal prices, that is, forecasts of metal prices increase (decrease) in the forecast horizon.[6] The results reported in Tables 4 and 5 show that anti-herding is the predominant strategy among metal-price forecasters. The evidence of anti-herding is stronger for the pessimists than for the optimists, but we do not find a single case of significant forecaster herding. The interpretation of the result that the evidence of anti-herding is stronger for the pessimists than for the optimists, however, should not be stretched too far. Whether pessimists or optimists show stronger signs of anti-herding depends on how we define these two groups of forecasts. The key message reported by Tables 4 and 5 is that anti-herding is still the dominant strategy when we split the sample of forecasters into optimists and pessimists.

It is also interesting to explore whether forecasters that anti-herd with respect to forecasting one metal price also anti-herd when it comes to forecasting another metal price. Because the CEF survey data set contains forecasts of the prices of Gold and Silver published by the same institutions, we perform such a cross-check of forecaster anti-herding for these two metals. To this end, we compute a forecaster-specific S_i-statistic for forecasts of the prices of Gold and Silver. Figure 3 shows for each forecast horizon that, in fact, forecaster anti-herding with respect to the price of Gold is highly positively correlated with forecaster anti-herding with respect to the price of Silver. A natural question is whether the high positive correlation of the forecaster-specific S_i-statistic that we observed for Gold and Silver also is a characteristic feature of other pairs of metal-price forecasts. In order to explore this question, Figure 4 plots for the one-month forecast horizon the correlation of the individual (anti)-

[6] We also define the optimists (pessimists) as those forecasters who predict an increase (a decrease) in metal prices (results are not reported, but available upon request). This definition, however, results in a relatively small proportion of pessimists of only about 18%. Using the term structure of forecasts to define the groups of optimists and pessimists (Tables 4 and 5), we obtain 6,203 (12,613) optimistic (pessimistic) forecasts, which implies a proportion of pessimists of about 2/3.

Figure 3: Individual (Anti-)Herding Statistic, S_i for Gold and Silver

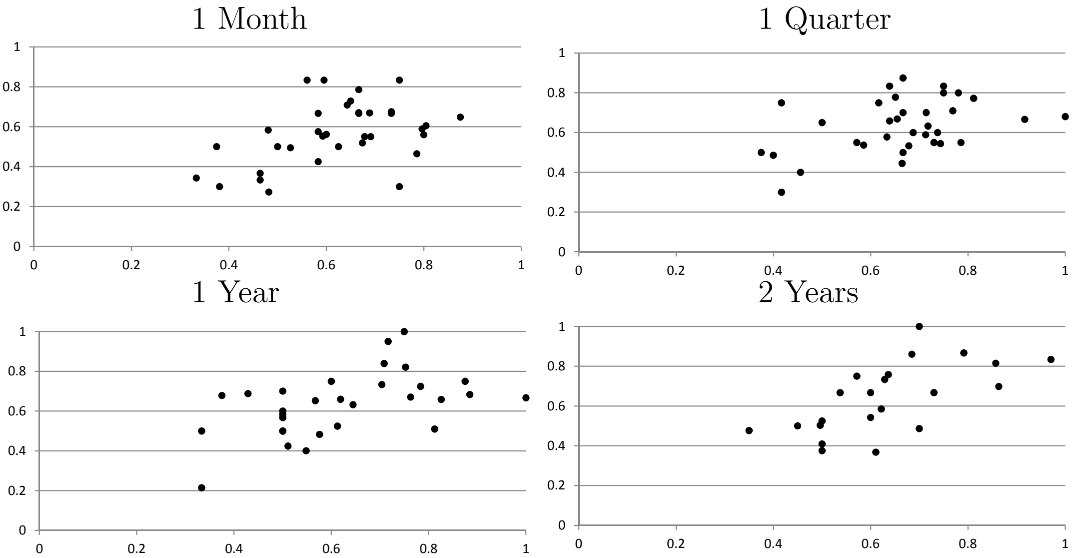

Note: Figure 3 shows the individual (anti-)herding statistics S_i for Gold (horizontal axis) and Silver (vertical axis) for each forecast horizon.

Figure 4: Individual (Anti-)Herding Statistic for Different Metal Pairs

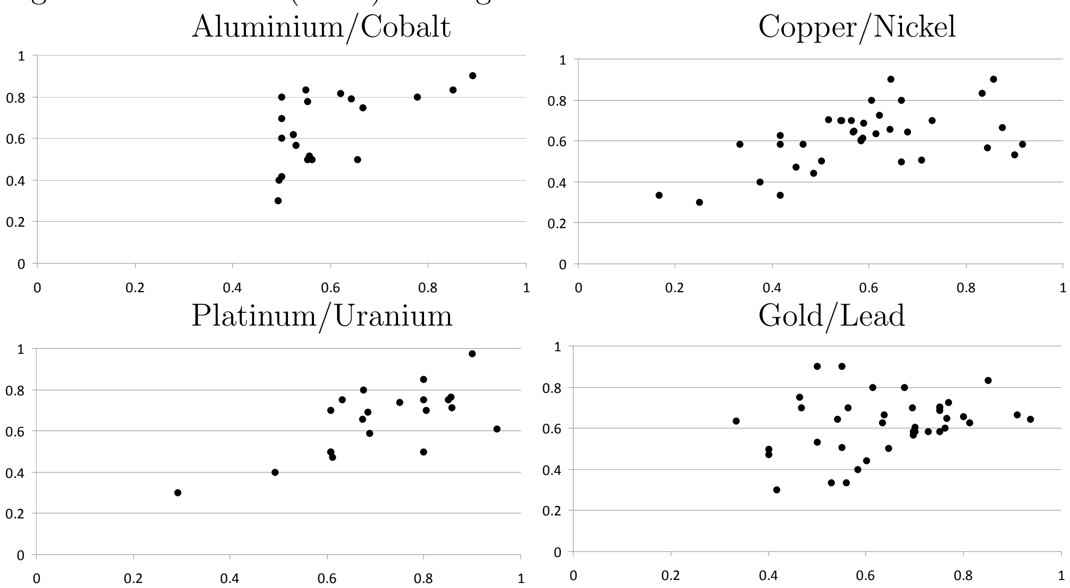

Note: Figure 4 shows the individual (anti-)herding statistics S_i, for different pairs of metal prices for one-month ahead forecasts. For example, the left-hand panel plots the individual (anti-)herding statistics for Aluminium on the vertical axis and the individual (anti-)herding statistics for Cobalt on the horizontal axis.

herding statistics for four pairs of metal prices: Aluminium/Cobalt, Copper/Nickel, Platinum/Uranium, and Gold/Lead. Two results emerge. First, anti-herding is the dominant forecasting strategy because most forecasters are located in the upper right quadrant of the four plotted panels. Second, there is a positive cross-metal link between the individual (anti-)herding statistics, thus, corroborating the results plotted in Figure 3.

5. Concluding remarks

We analyze forecasts of nine metal prices at four different forecasting horizons using data for a sample period that covers fifteen years of data. Our main finding is that forecasters appear to anti-herd, where the prevalence of forecaster anti-herding has undergone changes over time. Our findings suggest that forecaster anti-herding is a source of the empirically observed cross-sectional heterogeneity of forecasts. As a result, forecasts of metal prices give, for an outside observer, a more dispersed and, thus, less precise account of expected future movements of metal prices than it would be the case if private sector forecasters delivered unbiased forecasts.

In future research, it is interesting to study whether forecaster anti-herding is linked to financial market volatility in general and the volatility of metal prices in particular. For example, Bewley and Fiebig (2002) study whether interest-rate forecasters (anti-)herd. They find that the prevalence of forecaster herding is positively correlated with the volatility of interest rates, that is, with the difficulty to predict interest-rate changes. Laster et al. (1999, p. 304) argue that if forecasters' loss function is stable over time, a change in the strength of forecaster anti-herding indicates a change in the volatility of the variable being forecasted. Our findings on the time-variation of forecaster anti-herding may be a useful starting point to analyze in detail whether a link exists between forecaster anti-herding and the volatility of metal prices.

Another avenue for future research is to examine in detail the links between forecaster anti-herding, the cross-sectional heterogeneity of forecasts, and macroeconomic determinants of metal prices. Such a study could draw, for example, on recent research by Menkhoff et al. (2009). They analyze the determinants of the cross-sectional heterogeneity of forecasts of exchange rates. Consistent with chartist-fundamentalist models of exchange-rate determination, they find, for example, that misalignments of the exchange rate explain cross-sectional heterogeneity of forecasts. According to chartist-fundamentalist models, cross-sectional heterogeneity should decrease as an asset price moves farther away from some "fundamental" value because a consensus should emerge among market participants that the asset price is not in line with fundamentals. Misalignments of metal prices, thus, may shift the benefits and costs of making "extreme" forecasts. An interesting question is whether a link between the cross-sectional heterogeneity of forecasts and the prevalence of forecaster anti-herding, on the one hand, and misalignments of metal prices, on the other hand, can be detected.

References

Ahti, V., 2009. Forecasting commodity prices with nonlinear models. Discussion paper no. 268, Helsinki Center of Economic Research.

Arends, B., 2010. Is gold the next bubble? Wall Street Journal online, May 25, 2010.

Bernhardt, D., Campello, M., Kutsoati, E., 2006. Who herds? Journal of Financial Economics 80, 657-675.

Bewley, R., Fiebig, D.G., 2002. On the herding instinct of interest rate forecasters. Empirical Economics 27, 403-425.

Brenner, R.J., Kroner, K.F., 1995. Arbitrage, cointegration, and testing the unbiasedness hypothesis in financial markets. Journal of Financial and Quantitative Analysis 30, 23-42.

Canarella, G., Pollard, S.K., 1986. The 'efficiency' of the London Metal Exchange. Journal of Banking and Finance 10, 575-593.

Chow, Y.F., 1998. Regime switching and cointegration tests of the efficiency of futures markets. Journal of Futures Markets 18, 871-901.

Dooley, G., Lenihan, H., 2005. An assessment of time series methods in metal price forecasting. Resources Policy 30, 208-217.

Froot, K.A., Scharfstein, D.S., Stein, J.C., 1992. Herd on the street: Informational inefficiencies in a market with short-term speculation. Journal of Finance 52, 1461-1484.

Hsieh, D. A., Kulatilaka, N., 1982. Rational expectations and risk premia in forward markets: Primary metals at the London Metals Exchange. Journal of Finance 37, 1199-1207.

Laster, D., Bennett, P., Geoum, I.S., 1999. Rational bias in macroeconomic forecasts. Quarterly Journal of Economics 114, 293-318.

Menkhoff, L., Rebitzky, R., Schröder, M., 2009. Heterogeneity in exchange rate expectations: Evidence on the chartist-fundamentalist approach. Journal of Economic Behavior and Organization 70, 241-252.

Monk, E., 2012. Warren Buffet warns gold prices now in a bubble similar to technology stocks. This is Money online, February 13, 2012.

Naujoks, M., Aretz, K., Kerl, A., Walter, A., 2009. Do German security analysts herd? Financial Markets and Portfolio Management 23, 3-29.

Pierdzioch, C., Rülke, J.C., Stadtmann, G., 2010. New evidence of anti-herding of oil-price forecasters. Energy Economics 32, 1456-1459.

Rosen, S., 1981. The economics of superstars. American Economic Review 7, 845-858.

Scharfstein, D., Stein, J., 1990. Herd behavior and investment. American Economic Review 80, 465-479.

Schindler, M., 2011. Gold: Bubble or not? Forbes online, August 28, 2011.

Sephton, P.S., Cochrane, D.K., 1990. A note on the efficiency of the London Metal Exchange. Economics Letters 33, 341-345.

Stadtmann, G., C. Pierdzioch, Rülke, J.C., 2011. Scattered fiscal forecasts. Economics Bulletin 31, 2558-2568.

Watkins, C., McAleer, M., 2004. Econometric modelling of non-ferrous metal prices. Journal of Economic Surveys 18, 651-701.

United Nations, 2011. G20 Study Group on commodities. Contribution by the United Nations Secretariat, April 2011.

Table 1: Descriptive Statistics

Metal	Aluminium				Cobalt				Copper			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
Forecasts	1900.95	1915.02	1990.70	1998.63	19.48	18.70	16.06	13.85	4212.77	4184.20	4125.77	3981.99
Realizations	1860.12	1667.56	1680.01	1703.77	19.12	19.22	19.32	19.93	4313.23	4340.96	4410.53	4571.74
Cor	0.97	0.95	0.93	0.92	0.95	0.95	0.93	0.89	0.99	0.98	0.95	0.94
No. of forecasts	749	749	732	646	303	303	303	295	742	742	725	638
No. of forecasters	48	48	48	46	39	39	39	38	48	48	48	46
Sample period	08/1995 - 08/2011				08/1997 - 08/2011				08/1995 - 08/2011			

Metal	Gold				Lead				Nickel			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
Forecasts	613.99	620.90	622.59	605.51	1219.56	1201.00	1166.71	1144.13	14198.08	14086.56	13917.55	13139.21
Realizations	612.88	617.76	627.55	653.75	1249.86	1262.92	1285.93	1325.15	14835.24	14965.31	15240.93	15816.64
Cor	0.98	0.96	0.93	0.91	0.99	0.98	0.96	0.95	0.99	0.99	0.98	0.96
No. of forecasts	661	662	643	563	725	725	706	619	734	734	717	629
No. of forecasters	48	48	48	46	48	48	48	46	48	48	48	46
Sample period	08/1995 - 08/2011				08/1995 - 08/2011				08/1995 - 08/2011			

Metal	Platinum				Silver				Uranium			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
Forecasts	939.05	940.26	947.98	931.09	10.56	10.65	10.55	10.13	35.16	36.60	39.53	40.41
Realizations	946.55	946.55	946.55	987.17	10.99	11.12	11.36	11.82	32.01	32.44	33.33	34.61
Cor	0.99	0.98	0.97	0.97	0.99	0.98	0.97	0.97	0.98	0.97	0.95	0.94
No. of forecasts	485	485	473	431	598	598	576	508	358	358	348	332
No. of forecasters	45	45	45	44	48	48	48	46	48	48	48	46
Sample period	08/1996 - 08/2011				08/1995 - 08/2011				08/1995 - 08/2011			

Notes: Table 1 reports the average of the consensus forecasts and the realizations for the different forecast horizons. The realizations are the realized values for each forecast horizon, which is why the realizations differ across the forecast horizons. Cor= $Cor(\tilde{E}_t[s_{t+k}], f_{t,t+k})$ = correlation between the consensus forecast and the respective forward rate. The forecasts and realized values for Gold, Silver, and Platinum refer to $ per ounce, for Cobalt and Uranium they refer to $ per pound, and for the remaining metal prices the values refer to $ per tonne.

Table 2: Result of the Test of Forecaster (Anti-)Herding

Metal	Aluminium				Cobalt				Copper			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.56*	0.58*	0.60*	0.59*	0.71*	0.71*	0.74*	0.70*	0.59*	0.64*	0.64*	0.65*
Stand. Dev.	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.03	0.02	0.02	0.02	0.02
Lower 99%	0.52	0.53	0.55	0.54	0.63	0.64	0.67	0.63	0.55	0.59	0.59	0.59
Upper 99%	0.61	0.62	0.65	0.64	0.79	0.79	0.82	0.78	0.64	0.68	0.69	0.70
Observations	748	749	732	646	299	295	296	294	741	741	725	638

Metal	Gold				Lead				Nickel			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.63*	0.62*	0.57*	0.52	0.67*	0.68*	0.70*	0.68*	0.58*	0.62*	0.64*	0.66*
Stand. Dev.	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Lower 99%	0.58	0.57	0.51	0.46	0.62	0.63	0.65	0.62	0.54	0.57	0.59	0.60
Upper 99%	0.68	0.67	0.62	0.57	0.71	0.73	0.74	0.73	0.63	0.67	0.69	0.71
Observations	657	660	635	555	715	712	701	615	734	734	715	629

Metal	Platinum				Silver				Uranium			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.67*	0.70*	0.67*	0.67*	0.58*	0.61*	0.58*	0.53	0.65*	0.66*	0.59*	0.59*
Stand. Dev.	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.03
Lower 99%	0.61	0.64	0.61	0.61	0.52	0.56	0.53	0.48	0.58	0.59	0.52	0.52
Upper 99%	0.73	0.76	0.73	0.74	0.63	0.67	0.63	0.59	0.72	0.73	0.67	0.66
Observations	480	478	470	422	595	596	575	508	354	352	341	327

Notes: Table 2 reports the herding statistic, S, its standard deviation and the upper/lower 99% bound. The results are based on Equations (2) – (7). * indicates whether the S statistic is significantly different from 0.5 at the one percent significance level.

Table 3: Result of the Test of Forecaster (Anti-)Herding (Forward as Consensus)

Metal	Aluminium				Cobalt				Copper			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.65*	0.65*	0.70*	0.75*	0.80*	0.84*	0.88*	0.89*	0.79*	0.82*	0.86*	0.91*
Stand. Dev.	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.04	0.02	0.02	0.02	0.02
Lower 99%	0.60	0.60	0.65	0.70	0.72	0.76	0.79	0.79	0.75	0.77	0.81	0.86
Upper 99%	0.70	0.70	0.75	0.81	0.88	0.92	0.97	0.99	0.84	0.87	0.91	0.96
Observations	747	748	732	645	299	297	301	295	731	735	723	638

Metal	Gold				Lead				Nickel			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.83*	0.74*	0.68*	0.60*	0.82*	0.84*	0.87*	0.90*	0.74*	0.77*	0.82*	0.83*
Stand. Dev.	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
Lower 99%	0.78	0.69	0.63	0.55	0.77	0.79	0.83	0.85	0.69	0.73	0.77	0.78
Upper 99%	0.89	0.79	0.74	0.66	0.86	0.89	0.92	0.96	0.79	0.82	0.87	0.88
Observations	657	658	639	555	714	714	699	615	732	733	714	629

Metal	Platinum				Silver				Uranium			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.81*	0.83*	0.86*	0.90*	0.78*	0.63*	0.61*	0.58*	0.70*	0.65*	0.72*	0.71*
Stand. Dev.	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.03	0.03	0.03	0.03
Lower 99%	0.75	0.77	0.80	0.84	0.72	0.58	0.55	0.52	0.62	0.57	0.64	0.63
Upper 99%	0.87	0.89	0.92	0.96	0.83	0.69	0.66	0.64	0.77	0.73	0.80	0.79
Observations	478	477	469	423	597	598	575	508	353	352	343	330

Notes: Table 3 reports the herding statistic, S, its standard deviation and the upper/lower 99% bound. The results are based on Equations (2) – (7). * indicates whether the S statistic is significantly different from 0.5 at the one percent significance level.

Table 4: (Anti-)Herding of Optimists (Positive Term Structure of Forecasts)

Metal	Aluminium				Cobalt				Copper			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.62*	0.60*	0.54	0.52	0.71*	0.76*	0.71*	0.70	0.55	0.57	0.51	0.54
Stand. Dev.	0.03	0.03	0.03	0.04	0.07	0.07	0.07	0.08	0.03	0.04	0.04	0.05
Lower 99%	0.55	0.51	0.45	0.42	0.52	0.57	0.54	0.48	0.47	0.48	0.40	0.41
Upper 99%	0.70	0.68	0.62	0.62	0.89	0.96	0.89	0.92	0.64	0.67	0.62	0.66
Observations	278	268	225	168	51	51	56	42	238	221	168	112

Metal	Gold				Lead				Nickel			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.63*	0.57	0.53	0.47	0.62*	0.61*	0.57	0.49	0.57	0.64*	0.55	0.53
Stand. Dev.	0.03	0.03	0.03	0.04	0.03	0.04	0.04	0.05	0.03	0.03	0.03	0.04
Lower 99%	0.55	0.49	0.44	0.37	0.53	0.51	0.46	0.36	0.49	0.55	0.45	0.42
Upper 99%	0.71	0.66	0.62	0.57	0.71	0.71	0.68	0.62	0.65	0.72	0.64	0.64
Observations	277	271	230	175	215	196	147	106	268	253	213	156

Metal	Platinum				Silver				Uranium			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.70*	0.69*	0.54	0.46	0.59*	0.63*	0.62*	0.56	0.65*	0.55	0.54	0.55
Stand. Dev.	0.04	0.04	0.05	0.06	0.03	0.04	0.04	0.04	0.04	0.04	0.04	0.05
Lower 99%	0.59	0.58	0.41	0.30	0.50	0.53	0.52	0.44	0.55	0.44	0.42	0.42
Upper 99%	0.80	0.80	0.67	0.62	0.67	0.72	0.72	0.67	0.76	0.66	0.65	0.69
Observations	168	159	124	75	228	221	186	141	148	144	126	98

Notes: Table 4 reports the herding statistic, S, its standard deviation and the upper/lower 99% bound. The results are based on Equations (2) – (7). * indicates whether the S statistic is significantly different from 0.5 at the one percent significance level.

Table 5: (Anti-)Herding of Pessimists (Negative Term Structure of Forecasts)

Metal	Aluminium				Cobalt				Copper			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.61*	0.63*	0.59*	0.57*	0.71*	0.66*	0.64*	0.68*	0.61*	0.58*	0.60*	0.60*
Stand. Dev.	0.02	0.02	0.03	0.03	0.03	0.03	0.03	0.04	0.02	0.02	0.02	0.03
Lower 99%	0.55	0.56	0.52	0.50	0.63	0.57	0.54	0.59	0.56	0.52	0.54	0.54
Upper 99%	0.67	0.69	0.65	0.64	0.79	0.74	0.73	0.78	0.67	0.64	0.67	0.67
Observations	469	457	417	331	248	238	217	189	503	496	466	383

Metal	Gold				Lead				Nickel			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.64*	0.66*	0.64*	0.61*	0.68*	0.60*	0.59*	0.56	0.59*	0.61*	0.59*	0.61*
Stand. Dev.	0.03	0.03	0.03	0.03	0.02	0.02	0.02	0.03	0.02	0.02	0.03	0.03
Lower 99%	0.58	0.59	0.57	0.53	0.63	0.54	0.53	0.49	0.53	0.55	0.52	0.54
Upper 99%	0.71	0.73	0.71	0.69	0.74	0.66	0.65	0.63	0.65	0.67	0.65	0.69
Observations	380	371	339	276	500	498	458	376	466	457	413	330

Metal	Platinum				Silver				Uranium			
Horizon	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years	1 Month	1 Quarter	1 Year	2 Years
S-statistic	0.66*	0.63*	0.60*	0.54	0.57*	0.66*	0.57*	0.66*	0.65*	0.68*	0.60	0.58
Stand. Dev.	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.03	0.04	0.04	0.04	0.04
Lower 99%	0.58	0.56	0.52	0.45	0.51	0.59	0.51	0.59	0.56	0.59	0.50	0.47
Upper 99%	0.73	0.71	0.68	0.63	0.64	0.73	0.64	0.73	0.74	0.78	0.70	0.68
Observations	312	308	286	242	367	358	367	358	206	195	179	157

Notes: Table 5 reports the herding statistic, S, its standard deviation and the upper/lower 99% bound. The results are based on Equations (2) – (7). * indicates whether the S statistic is significantly different from 0.5 at the one percent significance level.