Journal of Banking & Finance 37 (2013) 150–158

Contents lists available at SciVerse ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

Forecasting metal prices: Do forecasters herd?

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

[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] University of Southern Denmark, Department of Business and Economics, Campusvej 55, 5230 Odense M, Denmark
[d] Deutsches Institut für Wirtschaftsforschung (DIW Berlin), Mohrenstraße 58, 10117 Berlin, Germany

ARTICLE INFO

Article history:
Received 11 December 2011
Accepted 19 August 2012
Available online 25 August 2012

JEL classification:
G17
C33
L61

Keywords:
Metal prices
Forecasting
Forecaster (anti-)herding

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.

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), 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

* 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), geo@sam.sdu.dk (G. Stadtmann).

[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.

0378-4266/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jbankfin.2012.08.016

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 Economics 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: 1 month, 1 quarter, 1 year, and 2 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.

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 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 $0 \leqslant \alpha \leqslant 1$ 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 Eq. (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 fore-*

[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.