

New Evidence of Anti-Herding of Oil-Price Forecasters

Christian Pierdzioch

Universität des Saarlandes
c.pierdzioch@mx.uni-saarland.de

Jan-Christoph Rülke

WHU-Otto Beisheim School
of Management
jan-c.ruelke@whu.edu

Georg Stadtmann

Europa-Universität Viadrina
stadtmann@whu.edu

Excellence in
Management
Education

WHU – Working Paper
Series in Economics

WP 10/04

New Evidence of Anti-Herding of Oil-Price Forecasters

Christian Pierdzioch, Jan Christoph Rülke and Georg Stadtmann*

May 2010

Abstract

We used the oil-price forecasts of the *Survey of Professional Forecasters* published by the European Central Bank to analyze whether oil-price forecasters herd or anti-herd. Oil-price forecasts are consistent with herding (anti-herding) of forecasters if forecasts are biased towards (away from) the consensus forecast. Based on a new empirical test developed by Bernhardt et al. (*J. Financ. Econ.* 80: 657–675, 2006), we found strong evidence of anti-herding among oil-price forecasters.

JEL classification: F31, D84, C33

Keywords: Oil price; Forecasting; Herding

Address:

Christian Pierdzioch, Universität des Saarlandes, Rechts- und Wirtschaftswissenschaftliche Fakultät, Postfach 15 11 50, 66041 Saarbrücken, c.pierdzioch@mx.uni-saarland.de.

* *Corresponding Author:* Georg Stadtmann, Europa-Universität Viadrina, Fakultät für Wirtschaftswissenschaften, Postfach 1786, 15207 Frankfurt (Oder), stadtmann@whu.edu

Jan-Christoph Rülke

Department of Economics, WHU – Otto Beisheim School of Management, Burgplatz 2, 56179 Vallendar, Germany, Tel.: +49-261-6509-286, Fax: +49-261-6509-289
E-mail address: jan-c.ruelke@whu.edu

We thank the editor Richard Tol and two anonymous referees for helpful comments and suggestions. We are grateful to the Euro Area Macroeconomic Developments Division of the European Central Bank for providing the data.

1 Introduction

The dynamics of the price of oil were characterized by large swings in 2008/2009. The large swings in the price of oil may reflect herding of traders. A natural question is whether such herding, to the extent that it occurred, was driven by herding in the forecasts of professional oil-price forecasters.

We implemented a robust empirical test developed by Bernhardt et al. (2006) to study whether professional oil-price forecasters did, in fact, herd. This test is easy to implement and delivers results that can be easily interpreted in economic terms. The test results do not provide evidence of herding. On the contrary, we find strong evidence of *anti*-herding. Evidence of anti-herding indicates that professional oil-price forecasters deliberately placed their forecasts away from the cross-sectional consensus forecast.

Evidence of anti-herding of professional oil-price forecasters is consistent with evidence of anti-herding of stock analysts (Naujoks et al. 2009) and macroeconomic forecasters (Batchelor and Dua 1990). Evidence of anti-herding, thus, is mounting, implying that it becomes an urgent issue for future research to analyze the consequences of anti-herding of forecasters for price dynamics in financial markets in general and in oil markets in particular. Anti-herding may also explain the significant extent of cross-sectional heterogeneity of forecasts documented in recent literature (e.g., Menkhoff et al. 2009).

In Section 2, we describe the test for (anti-)herding that we used in our empirical analysis. In Section 3, we describe our data. In Section 4, we present our empirical results. In Section 5, we offer some concluding remarks.

2 The Test

The economic intuition motivating the test developed by Bernhardt et al. (2006) can be developed by means of Figure 1. The horizontal line represents future realizations of the price of oil, and $\bar{E}_t[s_{t+1}]$ denotes the publicly known consensus forecast, computed as the average of the cross-section of all forecasts made in period t . Given their information set (and a potentially asymmetric posterior distribution over oil prices), forecasters form a median-unbiased private forecast of the price of oil, $\tilde{E}_{i,t}[s_{t+1}]$, where i denotes a forecaster index. The probability that a private forecast exceeds the subsequently realized oil price and the chance that a private forecast is less than the subsequently realized oil price, thus, should both be equal to $1/2$. In Figure 1, the private forecast exceeds the consensus forecast.

Forecasters “herd” if they issue forecasts that are biased in the direction of the consensus forecast. Accordingly, in Figure 1, the eventually published forecast, $E_{i,t}[s_{t+1}]$, is smaller than the private forecast. As a result, the probability that the biased public forecast exceeds the subsequently realized oil price should be smaller than $1/2$. Similarly, if the biased published forecast is less than the consensus forecast, the probability that the published biased forecast is less than the subsequently realized oil price should also be smaller than $1/2$. When forecasters anti-herd, in contrast, the probabilities should exceed $1/2$ because, in this case, forecasters deliberately place public forecasts farther away from the consensus forecast than justified by private forecasts.

– Please insert Figure 1 about here. –

A test for (anti-)herding, thus, can be developed as follows (Bernhardt et al. 2006). Under the null hypothesis of no herding (or anti-herding), the conditional probability, P , that an unbiased published forecast of the oil

price overshoots (undershoots) the subsequently realized oil price should be $1/2$, regardless of the consensus forecast. Accordingly, the conditional probability of undershooting in case a forecast exceeds the consensus forecast should be

$$P(s_{t+1} < E_{i,t}[s_{t+1}] | E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}], s_{t+1} \neq E_{i,t}[s_{t+1}]) = 1/2, \quad (1)$$

and the conditional probability of overshooting in case a forecast is less than the consensus forecast should be

$$P(s_{t+1} > E_{i,t}[s_{t+1}] | E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}], s_{t+1} \neq E_{i,t}[s_{t+1}]) = 1/2. \quad (2)$$

The two conditional probabilities, thus, should average to $1/2$. This is not the case under the alternative hypothesis of (anti-)herding. If forecasters herd, as in Figure 1, the conditional probability that a biased published forecast exceeds the subsequently realized oil price, given that a published forecast exceeds the consensus forecast, should be less than $1/2$. In a similar vein, the probability that a published forecast is less than the subsequently realized oil price, given that the published forecast is less than the consensus forecast, should also be less than $1/2$. Averaging the two conditional probabilities should yield a value less than $1/2$. In contrast, if forecasters anti-herd, the two conditional probabilities, and the average thereof, should exceed $1/2$.

The test statistic, S , is defined as the average of the sample estimates of the two conditional probabilities given in Equations (1) and (2). If forecasters do not (anti-)herd (null hypothesis), the test statistic should assume the value $S = 0.5$. If forecasters herd, the test statistic should assume a value $S < 0.5$. If forecasters anti-herd, the test statistic should assume value a $S > 0.5$. The test statistic, S , has an asymptotic normal distribution.

Bernhardt et al. (2006) show that the test statistic, S , is robust to various problems arising in the case of, for example, correlated forecast errors,

market-wide shocks, and optimism or pessimism among forecasters. It is the averaging that makes the test robust under the null hypothesis. For example, large swings in oil prices may give rise to a preponderance of positive unforecasted market-wide shocks to oil prices. Such shocks raise (lower) the probability that the subsequently realized oil price exceeds (falls short of) forecasts, given any conditioning information, but leave the average of the conditional probabilities unaffected under the null hypothesis. As a result, market-wide shocks and the resulting positive cross-correlation of forecast errors do not bias the mean of the S statistic. Because the variance of the S statistic attains a maximum under the null hypothesis (zero cross-correlation of forecast errors), the test statistic S is conservative insofar as positive unforecasted shocks make it more difficult to reject the null hypothesis of unbiasedness when we should do so (Type II error).

Finally, it should be noted that the test statistic, S , is robust to outliers in the data and large disruptive events like a sharp trend reversal in the price of oil because such events should have a minor effect on the conditional probabilities (i.e., empirical frequencies of events).

3 The Data

Our database is the Survey of Professional Forecasters (SPF), which is conducted and compiled by the European Central Bank (ECB) on a quarterly basis. Empirical analyses of the SPF database are scarce because the ECB released the database only recently. Furthermore, the few available empirical studies of the SPF database focus on macro forecasts. Garcia and Manzanares (2007) and Bowles et al. (2009) analyze the forecast accuracy of SPF forecasts and find that SPF forecasters tend to be too “optimistic” with respect to high growth rates of GDP and low inflation rates. Their

results indicate that it may be important to account for behavioral aspects of the forecasting process to fully understand SPF forecasts. One such behavioral aspect, which has not been analyzed in earlier literature, is the (anti-)herding of forecasters.

The database covers the sample period 2002Q1–2009Q2, and it contains information on individual oil-price forecasts issued a large number of forecasters. We can thus analyze both the cross-sectional and the time-series dimension of the data. The database contains three-month-ahead forecasts because the ECB publishes at the beginning of a quarter forecasts of the end-of-quarter oil price. Forecasts are available for 30 forecasting cycles, where participating forecasters work for institutions such as investment banks, large international corporations, economic research institutes, and at universities. In total, forecasts from 83 forecasters are available. The database, however, is unbalanced because not all forecasters participated in all surveys. In total, more than 1,400 forecasts are available for implementation of the herding test.¹ Table 1 summarizes descriptive statistics of the SPF database.

– Please include Table 1 about here. –

An important feature of the SPF database is that, unlike the stock analysts analyzed by Bernhardt et al. (2006), professional oil-price forecasters do not issue their forecasts sequentially. A natural question, thus, is whether individual forecasters can account for the consensus forecast when delivering their forecasts. One answer to this question is that the SPF survey is one survey among others, such that forecasters can infer information on the

¹In order to test whether the unbalancedness of the database distorts our empirical results, we studied, as a robustness check, the forecasts of those 21 forecasters who always made forecasts (roughly 630 observations). We obtained the following test statistics: one-quarter-ahead forecasts – $S = 0.5469$, two-quarter-ahead forecasts – $S = 0.5723$, three-quarter-ahead forecasts – $S = 0.5512$, four-quarter-ahead forecasts is $S = 0.5439$. All test statistics are significant at the 95% level. Detailed results are available upon request.

prevailing consensus forecast from the results of other surveys. Another answer is that many forecasts are being published in newsletters and/or on the internet, such that forecasters can easily make inferences about recent trends in the forecasting industry. Notwithstanding, we shall analyze in Section 4 the robustness of our results to this sequencing issue.

– Please insert Figure 2 about here. –

Figure 2 plots the oil price (solid line) and the cross-sectional mean of the lagged one-quarter-ahead oil-price forecasts (dotted lines). The vertical difference between the oil price and the cross-sectional mean is the aggregate forecast error. The forecast error was positive during the period of time before 2008Q3, when forecasters observed significant increases in the oil price. In other words, forecasts fell short of the oil price. The forecast error turned positive after the trend reversal in the oil price that occurred in 2008Q3. The shaded area shown in Figure 2 captures the cross-sectional range of oil price forecasts and indicates a fairly high degree of forecast heterogeneity.

4 Empirical Results

Table 2 summarizes the results for one-quarter-ahead forecasts. We estimated a conditional probability of undershooting (given that a forecast exceeded the consensus forecast) of 0.264, and a conditional probability of overshooting (given that a forecast was less than the consensus forecast) of 0.837, implying a test statistic of $S = 0.551$. The standard deviation is 0.013, such that the test statistic significantly exceeds its unbiased-forecasts value of 0.5. In other words, we find significant evidence of anti-herding of professional oil-price forecasters. For two-quarter ahead forecasts, we obtained $S = 0.546$, for three-quarter ahead forecasts the test statistic is $S = 0.545$, and for four-quarter ahead forecasts the test yields the result $S = 0.531$. All results are significant (Table 3). Irrespective of the forecast horizon, the

estimates of the conditional probability of undershooting are much smaller than the estimates of the conditional probability of overshooting. These estimated conditional probabilities are in line with the result shown in Figure 2 that forecasts tended to fall short of (exceed) prices when prices rose (fell).

– Please include Tables 2 and 3 about here. –

Because of the significant run-up and eventual collapse of the oil price shown in Figure 2, we analyzed the temporal stability of our results by means of rolling-window estimates of the test statistic, S . Figure 3 shows the results that we obtained when we used a rolling window of four years length.² The first data point shown in Figure 3 represents the results for 2002Q1–2005Q4. We then rolled the estimation window one step forward in time, and dropped (added) the data for 2002Q1 (2006Q1). We continued this process of dropping and adding data until we reached the end of the sample period. While the rolling-window estimates of the test statistic show some fluctuations over time, the confidence bands also shown in Figure 3 reveal that the test statistic always significantly exceeds its unbiased-forecasts value of 0.5.

– Please insert Figure 3 about here.–

Because forecasters simultaneously issue their forecasts, they may not know the consensus forecasts when forecasting the price of oil. For this reason, we combined short-term forecasts with long-term forecasts to analyze whether potential uncertainty regarding the consensus forecast affects our results. For example, we used the two-quarter-ahead forecast made in period $t - 1$ of the oil price in period $t + 1$ as consensus forecast when computing the S -statistic for one-quarter ahead forecasts made in period t . The consensus forecast derived from the two-quarter-ahead forecast is in the information

²A rolling window thus contains forecasts from 16 different forecasting cycles. Results for other rolling windows (not reported, but available upon request) are qualitatively similar.

set of forecasters in period t , and the forecasting horizon as of period t is identical for the lagged two-quarter-ahead forecasts and the one-quarter-ahead forecasts.³ The results summarized in Table 4 provide strong evidence of anti-herding of oil-price forecasters.

– Please insert Table 4 about here.–

5 Concluding Remarks

Our empirical results indicate that anti-herding prevails among SPF oil-price forecasters. Anti-herding of forecasters may provide a behavioral explanation for the wide range of forecasts observed in virtually every forecasting cycle. While we have reported evidence of anti-herding of oil-price forecasters, we have not analyzed reasons for why forecasters anti-herd. Given that evidence of anti-herding of forecasters is mounting, it would be interesting to develop in future research empirical tests that can discriminate between competing explanations of anti-herding. Anti-herding may reflect, for example, the shape of forecasters' utility function. Laster et al. (1999) construct a model economy populated by forecasters endowed with the same information set, the same beliefs about the "correct" forecasting model, and the same utility function. Forecasts are made for two types of customers. The first group of customers regularly consumes forecasts and is interested in an accurate forecast. The second group of customers only occasionally use forecasts. They take into consideration only the performance of a forecaster in the last forecasting cycle. The larger is the influence of the occasional users, the stronger is the incentive to deviate from the consensus forecast because if *'...forecasters are paid according to relative ability, they might scatter, since it is hard to win when making a forecast similar to others'* (Lamont 2002, page 268). The

³As a robustness check, we replaced the consensus forecast with the forward rate corresponding to the forecasting horizon. The results based on the forward rate (not reported, but available upon request) corroborated our main finding of anti-herding of oil-price forecasters.

question of whether this argument is applicable to SPF forecasts is left for future research.

References

- Batchelor, R.A. and Dua, P., 1990, Product Differentiation in the Economic Forecasting Industry, *International Journal of Forecasting* 6, 311 – 316.
- Bernhardt, D., Campello, M., and E. Kutsoati, 2006, Who Herds? *Journal of Financial Economics* 80, 657 – 675.
- Bowles, C., Friz, R., Genre, V., Kenny, G., Meyler A. and T. Rautanen, 2009, The ECB Survey of Professional Forecasters (SPF) – A Review After Eight Years' Experience, *Occasional Paper Series No. 59*.
- Garcia, J.A. and A. Manzanares, 2007, Reporting Biases and Survey Results – Evidence from European Professional Forecasters, *Working Paper Series No. 836*.
- Lamont, O.A., 2002, Macroeconomic Forecasts and Microeconomic Forecasters, *Journal of Economic Behavior and Organization* 48, 265 – 280.
- Laster, D., P. Bennett and I.S. Geoum, 1999, Rational Bias in Macroeconomic Forecasts, *Quarterly Journal of Economics* 114, 293 – 318.
- Menkhoff, L., R. Rebitzky, and M. Schröder, 2009, Heterogeneity in Exchange Rate Expectations: Evidence on the Chartist-Fundamentalist Approach, *Journal of Economic Behavior and Organization* 70, 241 – 252.
- Naujoks, M., Aretz, K., Kerl, A.G. and A. Walter, 2009, Do German Security Analysts Herd? *Financial Markets and Portfolio Management* 23, 3 – 29.

Figure 1: Herding of Forecasters






Figure 2: Actual Oil Price and Mean Forecast



Notes: The dotted line reflects the mean of the one-quarter-ahead oil price forecast at the time of its realization while the solid line shows the actual oil price. The vertical difference between the actual oil price and the oil price forecast is therefore the forecast error.

Figure 3: Rolling-Window Estimates of the Test Statistic (One-Quarter-Ahead Forecasts)



Note: The solid line represents the test statistic, S . The dashed lines represent the corresponding 95 % confidence band.

Table 1: Descriptive Statistics

	Average
Actual oil price	54.78
Expected oil price	54.08
Standard deviation of oil price forecasts	0.69
Sample period	2002Q1 – 2009Q2
Number of forecasters	83
Number of observations	1,434

Table 2: Empirical Results

One-quarter-ahead oil price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	120 / 0.163	184 / 0.264
$s_{t+1} > E_{i,t}[s_{t+1}]$	617 / 0.837	513 / 0.736
Sum	737 / 1	697 / 1
S-statistic	0.5506	
Standard deviation	0.0132	
Lower 95 %	0.5247	
Upper 95 %	0.5765	

Table 3: Results for Longer-Term Forecasts

Two-quarter-ahead oil-price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	101 / 0.139	149 / 0.230
$s_{t+1} > E_{i,t}[s_{t+1}]$	628 / 0.861	499 / 0.770 %
Sum	729 / 1	648 / 1
S-statistic	0.5457	
Standard deviation	0.0135	
Lower 95 %	0.5192	
Upper 95 %	0.5722	

Three-quarter-ahead oil-price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	63 / 0.096	124 / 0.185
$s_{t+1} > E_{i,t}[s_{t+1}]$	595 / 0.904	547 / 0.815
Sum	658 / 1	671 / 1
S-statistic	0.5445	
Standard deviation	0.0137	
Lower 95 %	0.5176	
Upper 95 %	0.5714	

Four-quarter-ahead oil-price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	54 / 0.084	91 / 0.145
$s_{t+1} > E_{i,t}[s_{t+1}]$	590 / 0.916	536 / 0.855
Sum	644 / 1	627 / 1
S-statistic	0.5306	
Standard deviation	0.0140	
Lower 95 %	0.5032	
Upper 95 %	0.5581	

Table 4: Using the Average Lagged Longer-Term Forecast as the Consensus Forecast

Two-quarter-ahead oil-price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	121 / 12.7 %	159 / 33.4 %
$s_{t+1} > E_{i,t}[s_{t+1}]$	834 / 87.3 %	317 / 66.6 %
Sum	955 / 100.0 %	476 / 100.0 %
S-Stat	0.6037	
Stand. Dev.	0.0140	
Lower 95 %	0.5762	
Upper 95 %	0.6312	

Three-quarter-ahead oil-price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	100 / 11.4 %	146 / 26.5 %
$s_{t+1} > E_{i,t}[s_{t+1}]$	781 / 88.6 %	405 / 73.5 %
Sum	881 / 100.0 %	551 / 100.0 %
S-Stat	0.5757	
Stand. Dev.	0.0136	
Lower 95 %	0.5491	
Upper 95 %	0.6023	

Four-quarter-ahead oil-price forecasts		
	$E_{i,t}[s_{t+1}] < \bar{E}_t[s_{t+1}]$	$E_{i,t}[s_{t+1}] > \bar{E}_t[s_{t+1}]$
$s_{t+1} < E_{i,t}[s_{t+1}]$	71 / 9.2 %	172 / 25.9 %
$s_{t+1} > E_{i,t}[s_{t+1}]$	700 / 90.8 %	491 / 74.1 %
Sum	771 / 100.0 %	663 / 100.0 %
S-Stat	0.5837	
Stand. Dev.	0.0132	
Lower 95 %	0.5577	
Upper 95 %	0.6096	