Oil Price Forecasting under Asymmetric Loss

Christian Pierdzioch

Jan-Christoph Rülke

Georg Stadtmann

European University Viadrina Frankfurt (Oder)

Department of Business Administration and Economics

Discussion Paper No. 314

February 2012

ISSN 1860 0921

EUROPA-UNIVERSITÄT VIADRINA FRANKFURT (ODER)

Oil Price Forecasting under Asymmetric Loss

February 2012

Abstract

Based on the approach advanced by Elliott et al. (Rev. Ec. Studies. 72, 1197−1125, 2005), we found that the loss function of a sample of oil price forecasters is asymmetric in the forecast error. Our findings indicate that the loss oil price forecasters incurred when their forecasts exceeded the price of oil tended to be larger than the loss they incurred when their forecast fell short of the price of oil. Accounting for the asymmetry of the loss function does not necessarily make forecasts look rational.

JEL classification: F31, D84

Keywords: Oil price; Forecasting; Loss function; Rationality of forecasts

Addresses:

Christian Pierdzioch, Helmut-Schmidt-University, Department of Economics, Holstenhofweg 85, P.O.B. 700822, 22008 Hamburg, Germany.

Jan-Christoph Rülke, Department of Economics, WHU – Otto Beisheim School of Management, Burgplatz 2, 56179 Vallendar, Germany.

Georg Stadtmann*, University of Southern Denmark, Department of Business and Economics, Campusvej 55, 5230 Odense M, Denmark, and Europa-Universität Viadrina, Lehrstuhl für Volkswirtschaftslehre, insb. Makroökonomik, Postfach 1786, 15207 Frankfurt (Oder), Germany, Tel. +49 335 5534-2700, stadtmann@europa-uni.de

* Corresponding author.

We thank an anonymous referee for helpful comments and suggestions. We gratefully acknowledge the financial support received from the Deutsche Bundesbank through the foundation 'Geld and Währung' (S126/10081/11).

1 Introduction

The economically important and potentially complex links between oil prices and macroeconomic dynamics have been the subject of a large and growing empirical literature (Hamilton 2009, Kilian 2008). Because tracking the oil price is important to explain and forecast macroeconomic dynamics, interest in studying the properties of oil-price forecasts has mushroomed in recent years (Pierdzioch et al. 2010, Reitz et al. 2010). Our empirical study contributes to this recent literature. In contrast to earlier literature, we ask whether the loss function of oil-price forecasters is symmetric or asymmetric. Symmetry of the loss function implies that forecasters seek to minimize the mean-squared forecast error, an assumption on which traditional tests of unbiasedness and rationality of forecasts rest (Ito 1990, Elliott and Ito 1999). Such traditional tests, however, are misspecified if oil-price forecasters' loss function is asymmetric (Batchelor and Peel 1999, Elliott et al. 2008). Rejection of the null hypothesis of rationality, thus, may simply reflect that oil-price forecasters have an asymmetric loss function.

Because traditional tests for rationality of forecasts do not take into account the potential asymmetry of oil-price forecasters' loss function, we studied the shape of the loss function and the rationality of forecasts by means of an approach recently advanced by Elliott et al. (2005). Their approach is easy to implement, it informs about the type of a potential asymmetry in oil-price forecasters' loss function, and it allows the rationality of forecasts under an asymmetric loss function to be tested. Our application of this approach to the study of oil-price forecasts closes a gap in the literature. In fact, while much significant empirical research on asymmetric loss functions has been done in recent years (Batchelor and Peel 1999, Elliott et al. 2008, Döpke et al. 2010, to name just a few), the results of this research have not been applied, to the best of our knowledge, to the study of oil-price forecasts. A recent study of oil-price forecasts under asymmetric loss is Auffhammer (2007). The focus of his study, however, is the forecasts published by the United States Energy Information Administration (EIA).

In Section 2, we briefly outline the approach developed by Elliott et al. (2005). In Section 3, we describe our data while in Section 4, we report our results. Our results indicate that the loss function of a sample of oil price forecasters is asymmetric, and that the form of the asymmetric loss function is important for the question of whether accounting for the asymmetry of the loss function makes forecasts rational. In Section 5, we offer some concluding remarks.

2 Theoretical Background

The approach developed by Elliott et al. (2005) rests on the assumption that the loss function, $\mathcal{L}$, of oil-price forecasters can be described in terms of the following general functional form:

$$\mathcal{L} = [\alpha + (1-2\alpha)I(s_{t+1} - f_{t+1} < 0)]|s_{t+1} - f_{t+1}|^p, \tag{1}$$

where s_{t+1} denotes the oil price, f_{t+1} denotes the forecast of the oil price in period $t+1$ formed in period t, I denotes the indicator function, $p = 1$ for a linear-linear (lin-lin) loss function and $p = 2$ for a quadratic-quadratic (quad-quad) loss function, and $\alpha \in (0,1)$ governs the degree of asymmetry of the loss function. In the case of $\alpha = 0.5$, the loss function is symmetric. The standard symmetric quadratic loss function studied in earlier literature obtains for $\alpha = 0.5$ and $p = 2$. In this case, the loss forecasters incur increases in the squared forecast error. For $\alpha = 0.5$ and $p = 1$, the loss increases in the absolute forecast error.

Elliott et al. (2005) show that, for a given parameter p, which defines the general functional form of the loss function, the asymmetry parameter, α, can be consistently estimated as

$$\hat{\alpha} = \frac{\left[\frac{1}{T}\sum_{t=\tau}^{T+\tau-1} v_t|s_{t+1} - f_{t+1}|^{p-1}\right]' \hat{S}^{-1} \left[\frac{1}{T}\sum_{t=\tau}^{T+\tau-1} v_t I(s_{t+1} - f_{t+1} < 0)|s_{t+1} - f_{t+1}|^{p-1}\right]}{\left[\frac{1}{T}\sum_{t=\tau}^{T+\tau-1} v_t|s_{t+1} - f_{t+1}|^{p-1}\right]' \hat{S}^{-1} \left[\frac{1}{T}\sum_{t=\tau}^{T+\tau-1} v_t|s_{t+1} - f_{t+1}|^{p-1}\right]}, \tag{2}$$

where $\hat{S} = \frac{1}{T}\sum_{t=\tau}^{T+\tau-1} v_t v_t'(I(s_{t+1} - f_{t+1} < 0) - \hat{\alpha})^2 |s_{t+1} - f_{t+1}|^{2p-2}$ denotes a weighting matrix, v_t denotes a vector of instruments, T denotes the number of forecasts available, starting at $t = \tau + 1$. Because the weighting matrix depends on $\hat{\alpha}$, estimation is done iteratively. Following Elliott et al. (2005) and Döpke et al. (2010), we consider four alternative sets of instruments: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Testing whether $\hat{\alpha}$ differs from α_0 is done by using the following z-test $\sqrt{T}(\hat{\alpha} - \alpha_0) \rightarrow \mathcal{N}(0, (\hat{h}'\hat{S}^{-1}\hat{h})^{-1})$, where $\hat{h} = \frac{1}{T}\sum_{t=\tau}^{T+\tau-1} v_t|s_{t+1} - f_{t+1}|^{p-1}$. Elliott et al. (2005) further prove that a test for rationality of oil-price forecasts, given a loss function of the lin-lin or a quad-quad type ($p = 1, 2$), can be performed by computing

$$J(\hat{\alpha}) = \frac{1}{T}\left(x_t'\hat{S}^{-1}x_t\right) \sim \chi^2_{d-1}, \tag{3}$$

where $x_t = \sum_{t=\tau}^{T+\tau-1} v_t[I(s_{t+1} - f_{t+1} < 0) - \hat{\alpha}]|s_{t+1} - f_{t+1}|^{p-1}$ and d denotes the number of instruments. In the case of a symmetric loss function, the rationality test is given by $J(0.5) \sim \chi^2_d$. The statistic $J(0.5)$ answers the question of whether forecasters under the maintained assumption of a symmetric loss function form rational oil-price forecasts. The statistic $J(\hat{\alpha})$, answers the question of whether forecasters form rational oil-price forecasts, given an estimated asymmetric loss function (lin-lin or quad-quad). A comparison of $J(\hat{\alpha})$

with $J(0.5)$ shows whether an asymmetric loss function helps to remedy a potential failure of rationality of forecasts observed under a symmetric loss function.

3 The Data

Quarterly oil-price forecasts from the Survey of Professional Forecasts (SPF) data conducted by the European Central Bank (ECB) are available for the sample period 2002Q4−2010Q4.[1] As shown in Figure 1, this sample period witnessed a substantial swing in the oil price (solid line), where the oil price started in 2002 at around 26 dollars per barrel and peaked in 2008 at around 140 dollars per barrel.[2] In the third quarter of 2008, a large oil-price reversal occurred. After having slumped to a level of 44 dollar per barrel in late 2009, the oil price again gained momentum at the end of the sample period to reach a range between approximately 70 and 80 dollars per barrel in the second half of 2010.

– Please insert Figure 1 about here. –

The SPF data contain information on individual oil-price forecasts delivered by forecasters who work at financial or non-financial institutions based within the European Union (Bowles

[1] Empirical analyses of the SPF database are scarce because the ECB released the database only recently. The few available empirical studies of the SPF data focus on the accuracy of macroeconomic forecasts (Garcia and Manzanares 2007, Bowles et al. 2007).

[2] We used the oil price at the beginning of a quarter. The oil-price data are drawn from Thompson Financial Datastream. All figures and all computations were implemented using the software R (R Development Core Team 2009).

et al. 2007).[3] The forecasting horizon is three-months-ahead forecasts because the ECB publishes at the beginning of a quarter forecasts of the end-of-quarter oil price.[4] The SPF data are unbalanced because not all forecasters participated in all surveys. In order to avoid a nonresponse bias, we compiled forecasts for those 25 forecasters who participated in all questionnaire studies conducted during our sample period. In total, the SPF data contain forecasts from 83 oil-price forecasters. Approximately 30% of forecasters participated in all surveys. For every forecaster who always participated in the questionnaire study, we have 33 forecasts.

Figure 1 also shows the cross-sectional average (dashed line) across individual oil-price forecasts. In the majority of cases (∼70%), the cross-sectional mean forecast fell short of the actual oil price, suggesting that forecasters were more cautious with respect to underpredictions of the oil price than with respect to overpredictions. The tendency of the mean forecast to underpredict rather than to overpredict the oil price is important because for a lin-lin loss function the estimate of the asymmetry parameter, $\hat{\alpha}$, is simply equal to the proportion of negative forecast errors if one assumes that the only instrument is a constant. We, thus,

[3] A natural question is whether forecasters have an incentive to deliver accurate forecasts. In our view, at least two issues play a key role in this respect. First, it is important to note that forecasters do not necessarily trade on their forecasts. However, traders at their institution may do so, and this may give rise to a direct or indirect link between forecast errors, trading profits, and forecasters' income. This link should strengthen forecasters' incentive to deliver accurate forecasts. Second, forecasters' career prospects may depend on reputation, which, in turn, may be adversely affected by large forecast errors. Career prospects, however, may also depend on rare but spectacular relative forecast successes. For example, if all forecasters deliver an accurate forecast, an individual forecaster may have little to gain from delivering an accurate forecast. If, in contrast, a forecaster is the only one whose forecast is accurate, the effect on reputation may be significant. Anti-herding of forecasters may then be an equilibrium strategy. See Laster et al. (1999).

[4] The SPF data also contain longer-term forecasts. We do not present results for longer-term forecasts because our empirical approach requires computations of leads and lags of the forecast error. Accounting for leads and lags in the case of longer-term forecasts would substantially reduce the number of forecasts per forecaster available for the empirical analysis. For three-months-ahead forecasts, we could use in total 33 forecasting cycles to estimate the shape of forecasters' loss functions and to test for rationality of forecasts. As a robustness check, we also analyzed six-months-ahead forecasts. The results (available upon request) turned out qualitatively similar to the results for three-months-ahead forecasts.

expected estimates for the asymmetry parameter, $\hat{\alpha}$, smaller than 0.5, at least in the case of a lin-lin loss function. As we shall show in Section 4, our empirical findings are in line with this expectation.

It should be noted, however, that individual forecasts showed a substantial degree of cross-sectional dispersion. In order to shed light on the cross-sectional dispersion of forecasts, Figure 1 shows a shaded area which is defined as the cross-sectional range between the maximum and the minimum oil-price forecast. The shaded area illustrates that some forecasters overpredicted the oil price, while others underpredicted the oil price. In case overprediction and underprediction vary in a systematic way across forecasters, the dispersion of forecasts should result in a cross-sectional variation in the asymmetry parameter, $\hat{\alpha}$, across forecasters.[5]

The kind of dispersion of forecasts as illustrated in Figure 1 has been analyzed also in several earlier empirical studies (for exchange rates, see MacDonald and Marsh 1996 and Benassy-Quere et al. 2003), but the only researchers so far who have linked the cross-sectional dispersion of forecasts to the asymmetry of forecasters loss function are Capistrán and Timmermann (2009). Dispersion of oil-price forecasts has also been documented by Pierdzioch et al. (2010), who report that anti-herding of oil-price forecasters may be a source of the cross-sectional dispersion of oil-price forecasts. Anti-herding of forecasters may arise, for example, if forecasters strategically interact and, as a result, their loss functions is not of the simple traditional quadratic form (see, for example, Laster et al. 1999).

[5]In addition to accounting for the arguments put forward in Footnote 3, one could imagine that forecasters' loss function may be asymmetric in the forecast error because of the non-linear payoffs that arise in case of, for example, plain vanilla options or the kind of popular knock-out and barrier instruments.

4 Empirical Findings

Table 1 summarizes, for every forecaster, the estimates of the asymmetry parameter, $\hat{\alpha}$, the corresponding standard error, and the z-test of the null hypothesis $\hat{\alpha} = \alpha_0 = 0.5$. The loss function is of the lin-lin form. The table summarizes the results for the four alternative choices of instruments. Table 2 summarizes the results for the quad-quad loss function.

– Please include Tables 1 and 2 and Figure 2 about here. –

Our findings provide strong evidence of an asymmetry parameter, $\hat{\alpha}$, that is smaller than 0.5. Oil-price forecasters' loss functions, thus, seem to be asymmetric, where the loss in case of a negative forecast error (the oil price falls short the forecast) tends to be larger than the loss forecasters incurred in case of a positive forecast error (the oil price exceeds the forecast) of the same magnitude. This finding is in line with the observation (Figure 1) that the cross-sectional mean of oil-price forecasts was often below the actual oil price. Figure 2 plots the implications of our empirical findings for the shape of forecasters loss function, where we assumed for illustrative purposes that the loss function is of the lin-lin form. In order to draw Figure 2, we further assumed that the shape of the loss function is governed by the cross-sectional mean value of the estimated asymmetry parameter, $\hat{\alpha}$, estimated under Model 1.

Tables 1 and 2 also reveal some variation across forecasters with respect to the asymmetry parameter, $\hat{\alpha}$. This variation may account, at least in part, for the dispersion of forecasts shown in Figure 1. In the case of a lin-lin loss function, estimates of the asymmetry parameters varies approximately between 0.14 and 0.24, depending on the model that is being considered. The range of estimates in the case of a quad-quad loss function varies roughly between 0.04 and

0.38. In general, the standard errors of the estimates are larger in the cases of Model 1 and Model 2 for a quad-quad loss function than for a lin-lin function. Hence, the results of the z-test in case of Model 1 and Model 2 are somewhat smaller than in the case of Model 3 and Model 4 for the quad-quad loss function. As for Model 3 and Model 4, the results of the z-test are significant irrespective of whether one considers a lin-lin loss function or a quad-quad loss function.

– Please include Tables 3 and 4 about here. –

Tables 3 and 4 summarize the results of the J test of an asymmetric loss function and forecast rationality. We report results for $J(\hat{\alpha})$ and $J(0.5)$. Table 3 summarizes the results for a lin-lin loss function. Table 4 summarizes the results for a quad-quad loss function. Under the assumption of a lin-lin loss function, the vast majority of the $J(0.5)$ tests reject the hypothesis of forecast rationality. In contrast, the results for the $J(\hat{\alpha})$ tests do not lead to a rejection of the hypothesis of forecast rationality. Accounting for an asymmetric loss function of the lin-lin form, thus, helps to remedy the finding that, under a quadratic loss function, oil-price forecasters do not form rational forecasts. This finding, however, does not extend to the case of a quad-quad loss function. Assuming a quad-quad loss function leads to the result that both the $J(\hat{\alpha})$ tests and $J(0.5)$ tests yield significant results for Model 3 and Model 4, implying a rejection of the hypothesis of rational forecasts. It follows that, under the quad-quad loss function, the orthogonality condition between forecast errors, on the one hand side, and lagged forecast errors and the lagged oil price, on the other hand side, does not hold.

– Please include Tables 5 and 6 about here. –

As a robustness test, we studied pooled data. In the case of pooled data, we can use $33 \times 25 = 825$ forecasts. Table 5 summarizes the results for pooled data and a lin-lin loss function, and Table 6 summarizes the results for pooled data and a quad-quad loss function. As in the case of individual forecasters, the estimated asymmetry parameter, $\hat{\alpha}$, is significantly smaller than 0.5, a result that does not depend on the form of the loss function (lin-lin, quad-quad). The $J(0.5)$ test implies a rejection of the null hypothesis of rational oil-price forecasts. The $J(\hat{\alpha})$ test does not reject the null hypothesis of rationality of oil-price forecasts for Model 2, but rejects the null hypothesis of rationality of oil-price forecasts in all other models. Interestingly, rationality of forecasts can now be rejected also for Model 3 and Model 4 in the case of a lin-lin loss function, which is in contrast to the results we obtained for individual forecasts.

– Please include Table 7 and 8 about here. –

As yet another robustness test, we performed a subsample analysis. To this end, we excluded data from 2007Q1 onwards from our sample of data to check whether the large increase and eventual slump in the oil price that occurred at the end of 2008 has a significant effect on our results. For this subsample analysis, we used 450 forecasts. Tables 7 (lin-lin loss function) and 8 (quad-quad loss function) summarize the results of this robustness check for pooled data. As for the full sample of data, the asymmetry parameter, $\hat{\alpha}$, is smaller than 0.5 and the $J_3(\hat{\alpha})$ and $J_4(\hat{\alpha})$ for rationality yield significant results.[6]

[6] As yet another robustness check, we considered the possibility that oil-price forecasters' loss function is not of the lin-lin and quad-quad form considered by Elliott et al. (2005). To this end, we implemented the empirical test suggested by Batchelor and Peel (1998), which is based on the assumption that forecasters' loss function is of the linex form. Again, we found that accounting for asymmetry of the loss function often implies that rationality of forecasts can be rejected (results are available upon request).

5 Concluding Remarks

In terms of a suggested interpretation, our findings imply that oil-price forecasters' loss function may be asymmetric. The asymmetry of the loss function seems to reflect that the loss oil-price forecasters incurred when they overpredicted the price of oil exceeded the loss they incurred when they underpredicted the oil price. Irrespective of the link between forecast accuracy and forecasters' reputation, however, a lin-lin loss function and a quad-quad loss function do not necessarily suffice to make forecasts derived from such loss functions look rational. One possibility is that forecasters indeed form irrational forecasts that are not orthogonal to information in their information set. Another possibility is that forecasters form rational forecasts, but that the process of forecasting the oil price is more complex than implied by the lin-lin or quad-quad loss functions that we have considered in our empirical analysis. For example, strategic interactions among forecasters may lead forecasters to publish forecasts that intentionally deviate from the forecasts of others. Empirical evidence of such "anti-herding" of oil-price forecasters has been reported by Pierdzioch et al. (2010). If forecasters anti-herd, their loss function is likely to deviate from a simple quadratic loss function (Laster et al. 1999) and, thus, rational forecasts violate traditional rationality criteria, which are based on a quadratic loss function. If anti-herding, however, reflects deviations from a quadratic loss function, it is not necessarily the case that loss functions of the lin-lin or the quad-quad form suffice to fully account for such deviations from a quadratic loss function.[7]

[7] As a test of the anti-herding hypothesis, we used the S statistic suggested by Bernhardt et al. (2006) (see also Pierdzioch et al. 2010). We computed the S statistic for every single one of our 25 forecasters. In all cases, we found $S > 0.5$, a result that indicates that forecasters anti-herd. We then correlated the S statistic with the estimated asymmetry parameter, $\hat{\alpha}$. Irrespective of whether we used a lin-lin or a quad-quad loss function, the correlation was significantly negative. For example, in the case of a lin-lin loss function, we found a correlation of -0.51 (t-value = -2.76, p-value = 0.01). The asymmetry of forecasters' loss function, thus, is significantly inversely linked to the propensity of forecaster anti-herding. Because the correlation is not perfect, however, an asymmetric loss function, at least of the type considered in this research, does not fully capture forecaster anti-herding.

Given that the oil price substantially rose during our sample period, yet another possibility is that the persistent underestimation of the oil price reflects that forecasters expected a collapse of the ensuing "bubble". This aspect has been analyzed in recent literature by Reitz et al. (2010) by means of a regime-switching approach. As compared to a regime-switching approach, studying an asymmetric loss function renders it possible to make forecasters' fears of a collapse of the bubble "visible" by means of the shape of their loss function.[8] A further interesting feature of our analysis of forecasters' asymmetric loss function is that forecasts are forward-looking by nature, implying that asymmetries in forecasters' loss function may recover fears of a crashing bubble earlier than a regime-switching approach. Such a potential link between forecasts, asymmetric loss functions, and subsequent collapses of bubbles should be explored in detail in future research.

It is also interesting to compare our results with those reported by Auffhammer (2007) for the EIA. With regard to current-year forecasts of the price of oil, his estimated asymmetry parameter, α is larger than 0.5, where his data cover the sample period 1985−2003. Furthermore, he finds that forecasts of the price of oil are consistent with rational expectations. In contrast, our results indicate that the estimated asymmetry parameter, α, for professional economists is smaller than 0.5, and that their forecasts are not necessarily consistent with rationality. It, thus, seems that the loss function of the EIA, a government agency, markedly differs from the loss function of private agents. This difference in results should also be explored in future research.

[8] Another important point to note is that we found an asymmetric loss function not only for the full sample period, but also for the shorter subsample period. The asymmetric loss function, thus, most likely does not only reflect fears of a collapsing bubble (which gathered steam in the second half of the sample period). Rather it seems that the structure of forecasters' preferences led them to underestimate the oil price. An advantage of the asymmetric-loss-function approach is that we can find signs of such "deep preference structures" even when there are no regime shifts.

References

Auffhammer, M. (2007). The Rationality of EIA Forecasts Under Symmetric and Asymmetric Loss. *Resource and Energy Economics* 29: 102−112.

Batchelor, R. and D. Peel (1998). Rationality Testing under Asymmetric Loss. *Economics Letters* 61: 49−54.

Benassy-Quere, A., S. Larribeau, and R. MacDonald (2003). Models of Exchange Rate Expectations: How Much Heterogeneity? *Journal of International Financial Markets, Institutions and Money* 13: 113−136.

Bernhardt, D., M. Campello, and E. Kutsoati. (2006). Who Herds? *Journal of Financial Economics* 80: 657−675.

Bowles, C., R. Friz, V. Genre, G. Kenny, A. Meyler, and T. Rautanen (2007). The ECB Survey of Professional Forecasters (SPF) – A Review After Eight Years' Experience. European Central Bank Occasional Paper Series No. 59.

Capistrán, C., and A. Timmermann (2009). Disagreement and Biases in Inflation Expectations. *Journal of Money, Credit, and Banking* 41: 365−396.

Döpke, J., U. Fritsche, and B. Siliverstovs (2010). Evaluating German Business Cycle Forecasts under an Asymmetric Loss Function. *Journal of Business Cycle Measurement and Analysis* 2010/1: 1−18.

Elliott, G., and T. Ito (1999). Heterogeneous Expectations and Tests of Efficiency in the Yen/Dollar Forward Exchange Rate Market. *Journal of Monetary Economics* 43: 435−456.

Elliott,G., I. Komunjer, and A. Timmermann (2005). Estimation and Testing of Forecast Rationality under Flexible Loss. *Review of Economics Studies* 72: 1107−1125.

Elliott,G., I. Komunjer, and A. Timmermann (2008). Biases in Macroeconomic Forecasts: Irrationality or Asymmetric Loss? *Journal of the European Economic Association* 6: 122−157.

Garcia, J.A., and A. Manzanares (2007). Reporting Biases and Survey Results – Evidence from European Professional Forecasters. European Central Bank Working Paper Series No. 836.

Hamilton, J.D. (2009). Understanding Crude Oil Prices. *Energy Journal* 30: 179−206.

Ito, T. (1990). Foreign Exchange Expectations: Micro Survey Data. *American Economic Review* 80: 434−449.

Kilian, L., (2008). Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? *Review of Economics and Statistics* 90: 216−240.

Laster, D., P. Bennett, and I.S. Geoum (1999). Rational Bias in Macroeconomic Forecasts. *Quarterly Journal of Economics* 114: 293−318.

MacDonald, R., and I.W. Marsh (1996). Currency Forecasters are Heterogeneous: Confirmation and Consequences. *Journal of International Money and Finance* 15: 665−685.

Pierdzioch, C., J.C. Rülke, and G. Stadtmann (2010). New Evidence of Anti-Herding of Oil-Price Forecasters. *Energy Economics* 32: 1456−1459.

R Development Core Team (2009). R: A Language and Environment for Statistical Computing. Internet address: http://www.R-project.org.

Reitz, S., J.C. Rülke, and G. Stadtmann (2010). Are Oil-Price-Forecasters Finally Right? − Regressive Expectations Towards More Fundamental Values of the Oil Price. *Journal of Economics and Statistics* 230: 454−466.

Figure 1: The Data

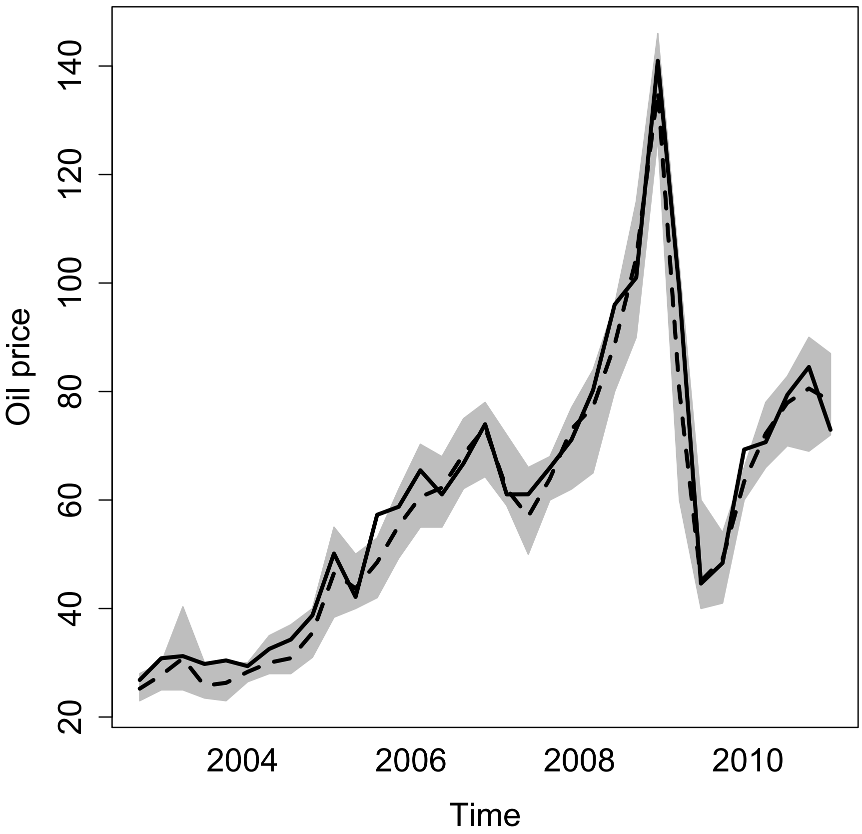

Note: The solid line shows the oil price. The dashed line shows the (lagged) cross-sectional mean forecast. The shaded area shows the range of forecasts.

Figure 2: The Loss Function

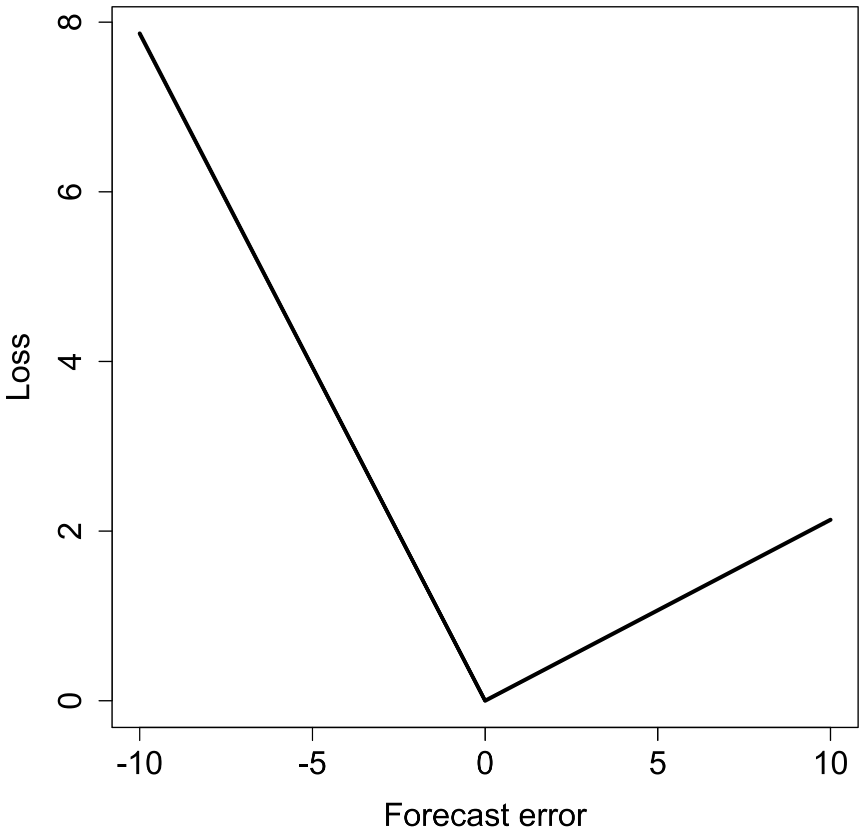

Note: This lin-lin loss function is based on the cross-sectional mean of the asymmetry parameter, $\hat{\alpha}$, estimated under Model 1. The forecast error is defined as the difference between the oil price and the oil-price forecast.

Table 1: Asymmetry parameter, lin-lin loss function

No.	$\hat{\alpha}_{Model1}$	se	z-test	$\hat{\alpha}_{Model2}$	se	z-test	$\hat{\alpha}_{Model3}$	se	z-test	$\hat{\alpha}_{Model4}$	se	z-test
1	0.242	0.075	-3.453	0.242	0.075	-3.466	0.164	0.064	-5.222	0.157	0.063	-5.403
2	0.182	0.067	-4.739	0.182	0.067	-4.740	0.153	0.063	-5.536	0.145	0.061	-5.776
3	0.182	0.067	-4.739	0.182	0.067	-4.741	0.153	0.063	-5.536	0.145	0.061	-5.801
4	0.182	0.067	-4.739	0.182	0.067	-4.745	0.153	0.063	-5.536	0.148	0.062	-5.710
5	0.273	0.078	-2.932	0.272	0.078	-2.936	0.245	0.075	-3.412	0.240	0.074	-3.501
6	0.242	0.075	-3.453	0.241	0.074	-3.484	0.213	0.071	-4.021	0.199	0.070	-4.321
7	0.242	0.075	-3.453	0.242	0.075	-3.453	0.209	0.071	-4.118	0.203	0.070	-4.244
8	0.212	0.071	-4.045	0.209	0.071	-4.107	0.176	0.066	-4.881	0.176	0.066	-4.893
9	0.182	0.067	-4.739	0.182	0.067	-4.740	0.118	0.056	-6.812	0.110	0.054	-7.175
10	0.182	0.067	-4.739	0.182	0.067	-4.745	0.153	0.063	-5.536	0.145	0.061	-5.800
11	0.242	0.075	-3.453	0.231	0.073	-3.658	0.187	0.068	-4.621	0.187	0.068	-4.622
12	0.182	0.067	-4.739	0.182	0.067	-4.740	0.153	0.063	-5.536	0.146	0.061	-5.765
13	0.273	0.078	-2.932	0.268	0.077	-3.008	0.220	0.072	-3.881	0.219	0.072	-3.896
14	0.242	0.075	-3.453	0.242	0.075	-3.457	0.239	0.074	-3.521	0.237	0.074	-3.551
15	0.212	0.071	-4.045	0.188	0.068	-4.590	0.167	0.065	-5.119	0.166	0.065	-5.162
16	0.242	0.075	-3.453	0.241	0.074	-3.483	0.218	0.072	-3.920	0.216	0.072	-3.965
17	0.212	0.071	-4.045	0.212	0.071	-4.052	0.176	0.066	-4.881	0.170	0.065	-5.039
18	0.212	0.071	-4.045	0.211	0.071	-4.059	0.171	0.065	-5.030	0.165	0.065	-5.174
19	0.152	0.062	-5.583	0.148	0.062	-5.705	0.061	0.042	-10.522	0.049	0.038	-11.969
20	0.212	0.071	-4.045	0.205	0.070	-4.195	0.185	0.068	-4.664	0.160	0.064	-5.326
21	0.182	0.067	-4.739	0.181	0.067	-4.759	0.153	0.063	-5.536	0.144	0.061	-5.813
22	0.273	0.078	-2.932	0.269	0.077	-2.997	0.227	0.073	-3.748	0.224	0.073	-3.807
23	0.182	0.067	-4.739	0.181	0.067	-4.759	0.153	0.063	-5.536	0.150	0.062	-5.617
24	0.182	0.067	-4.739	0.182	0.067	-4.739	0.128	0.058	-6.395	0.113	0.055	-7.007
25	0.212	0.071	-4.045	0.211	0.071	-4.066	0.184	0.067	-4.696	0.173	0.066	-4.977

Note: se = standard error, z-test = test of the null hypothesis that $\hat{\alpha} = 0.5$. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 2: Asymmetry parameter, quad-quad loss function

No.	$\hat{\alpha}_{Model1}$	se	z-test	$\hat{\alpha}_{Model2}$	se	z-test	$\hat{\alpha}_{Model3}$	se	z-test	$\hat{\alpha}_{Model4}$	se	z-test
1	0.292	0.126	-1.644	0.288	0.125	-1.688	0.117	0.066	-5.828	0.086	0.040	-10.432
2	0.265	0.117	-2.005	0.238	0.102	-2.556	0.110	0.072	-5.385	0.109	0.056	-6.938
3	0.307	0.125	-1.541	0.308	0.125	-1.539	0.136	0.084	-4.350	0.109	0.063	-6.232
4	0.314	0.121	-1.534	0.317	0.120	-1.530	0.184	0.092	-3.447	0.178	0.082	-3.939
5	0.390	0.135	-0.816	0.378	0.134	-0.911	0.232	0.100	-2.692	0.202	0.087	-3.430
6	0.389	0.142	-0.781	0.365	0.135	-1.002	0.204	0.087	-3.411	0.160	0.069	-4.952
7	0.355	0.129	-1.122	0.355	0.129	-1.120	0.198	0.089	-3.392	0.170	0.071	-4.680
8	0.237	0.114	-2.303	0.238	0.114	-2.299	0.086	0.066	-6.260	0.065	0.046	-9.454
9	0.290	0.121	-1.741	0.287	0.120	-1.770	0.128	0.073	-5.072	0.106	0.060	-6.547
10	0.285	0.122	-1.760	0.285	0.122	-1.757	0.116	0.076	-5.038	0.089	0.055	-7.509
11	0.342	0.126	-1.254	0.336	0.125	-1.320	0.186	0.081	-3.879	0.161	0.067	-5.046
12	0.285	0.122	-1.754	0.286	0.122	-1.753	0.121	0.074	-5.149	0.094	0.051	-7.936
13	0.326	0.118	-1.469	0.323	0.118	-1.500	0.182	0.076	-4.209	0.160	0.062	-5.504
14	0.315	0.121	-1.526	0.312	0.121	-1.555	0.186	0.082	-3.809	0.161	0.068	-4.984
15	0.276	0.119	-1.881	0.233	0.112	-2.377	0.094	0.059	-6.885	0.081	0.045	-9.267
16	0.298	0.130	-1.551	0.289	0.129	-1.638	0.120	0.073	-5.179	0.089	0.045	-9.148
17	0.349	0.139	-1.089	0.324	0.134	-1.311	0.149	0.079	-4.425	0.103	0.052	-7.710
18	0.312	0.133	-1.418	0.308	0.131	-1.468	0.153	0.079	-4.423	0.119	0.063	-6.016
19	0.212	0.104	-2.778	0.214	0.103	-2.773	0.054	0.055	-8.122	0.040	0.034	-13.671
20	0.306	0.128	-1.517	0.288	0.121	-1.744	0.135	0.070	-5.241	0.096	0.051	-7.945
21	0.275	0.120	-1.875	0.274	0.120	-1.882	0.118	0.071	-5.338	0.089	0.051	-8.118
22	0.361	0.129	-1.077	0.358	0.129	-1.097	0.189	0.086	-3.600	0.165	0.066	-5.107
23	0.243	0.124	-2.076	0.211	0.101	-2.860	0.059	0.062	-7.154	0.068	0.042	-10.302
24	0.309	0.119	-1.604	0.310	0.119	-1.602	0.147	0.081	-4.379	0.128	0.067	-5.590
25	0.295	0.119	-1.716	0.293	0.119	-1.742	0.141	0.078	-4.620	0.108	0.059	-6.597

Note: se = standard error, z-test = test of the null hypothesis that $\hat{\alpha} = 0.5$. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 3: J-test, lin-lin loss function

No.	$J_2(0.5)$	p	$J_3(0.5)$	p	$J_4(0.5)$	p	$J_2(\hat{\alpha})$	p	$J_3(\hat{\alpha})$	p	$J_4(\hat{\alpha})$	p
1	8.820	0.012	13.890	0.001	14.039	0.003	0.047	0.828	5.359	0.021	5.891	0.053
2	13.367	0.001	15.210	0.000	15.502	0.001	0.002	0.968	1.578	0.209	2.041	0.360
3	13.371	0.001	15.210	0.000	15.578	0.001	0.005	0.944	1.578	0.209	2.091	0.352
4	13.381	0.001	15.210	0.000	15.317	0.002	0.011	0.915	1.578	0.209	1.914	0.384
5	6.839	0.033	8.934	0.011	9.133	0.028	0.020	0.889	1.953	0.162	2.285	0.319
6	8.861	0.012	10.809	0.004	11.567	0.009	0.108	0.743	1.843	0.175	2.742	0.254
7	8.759	0.013	11.148	0.004	11.415	0.010	0.001	0.974	2.138	0.144	2.514	0.285
8	11.174	0.004	13.355	0.001	13.365	0.004	0.167	0.682	2.122	0.145	2.151	0.341
9	13.367	0.001	17.114	0.000	17.533	0.001	0.002	0.963	4.068	0.044	4.807	0.090
10	13.383	0.001	15.210	0.000	15.612	0.001	0.013	0.909	1.578	0.209	2.088	0.352
11	10.075	0.006	13.079	0.001	13.400	0.004	0.698	0.404	3.618	0.057	3.618	0.164
12	13.367	0.001	15.210	0.000	15.508	0.001	0.002	0.962	1.578	0.209	2.019	0.364
13	7.384	0.025	11.106	0.004	11.170	0.011	0.338	0.561	3.634	0.057	3.688	0.158
14	8.772	0.012	9.072	0.011	9.152	0.027	0.013	0.908	0.237	0.626	0.338	0.845
15	13.490	0.001	14.218	0.001	15.457	0.001	1.409	0.235	2.695	0.101	2.798	0.247
16	8.897	0.012	10.448	0.005	10.462	0.015	0.107	0.744	1.532	0.216	1.673	0.433
17	10.960	0.004	13.355	0.001	13.578	0.004	0.018	0.892	2.122	0.145	2.501	0.286
18	10.987	0.004	13.815	0.001	14.216	0.003	0.037	0.848	2.480	0.115	2.828	0.243
19	16.280	0.000	20.928	0.000	20.972	0.000	0.184	0.668	8.092	0.004	11.120	0.004
20	11.362	0.003	12.763	0.002	14.083	0.003	0.400	0.527	1.590	0.207	3.192	0.203
21	13.421	0.001	15.210	0.000	15.611	0.001	0.040	0.841	1.578	0.209	2.112	0.348
22	7.197	0.027	10.364	0.006	10.369	0.016	0.290	0.590	3.173	0.075	3.379	0.185
23	13.446	0.001	15.210	0.000	15.229	0.002	0.042	0.837	1.578	0.209	1.736	0.420
24	13.364	0.001	16.473	0.000	17.162	0.001	0.000	0.984	3.240	0.072	4.463	0.107
25	11.009	0.004	12.847	0.002	13.321	0.004	0.056	0.812	1.669	0.196	2.354	0.308

Note: p = p-value. J_i, $i = 2, 3, 4$ denotes the J-test for Model i. $J(0.5)$ denotes the J-test for a symmetric lin-lin loss function. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 4: J-test, quad-quad loss function

No.	$J_2(0.5)$	p	$J_3(0.5)$	p	$J_4(0.5)$	p	$J_2(\hat{\alpha})$	p	$J_3(\hat{\alpha})$	p	$J_4(\hat{\alpha})$	p
1	2.858	0.239	14.220	0.001	15.791	0.001	0.069	0.792	9.056	0.003	29.252	0.000
2	5.716	0.057	13.145	0.001	14.112	0.003	0.266	0.606	6.626	0.010	9.715	0.008
3	2.612	0.271	12.901	0.002	14.590	0.002	0.010	0.919	6.335	0.012	12.228	0.002
4	2.602	0.272	10.730	0.005	11.213	0.011	0.025	0.875	3.752	0.053	4.519	0.104
5	1.193	0.551	8.304	0.016	9.649	0.022	0.200	0.655	4.242	0.039	6.604	0.037
6	1.366	0.505	9.770	0.008	12.501	0.006	0.243	0.622	6.159	0.013	13.193	0.001
7	1.372	0.504	9.780	0.008	11.147	0.011	0.009	0.925	4.988	0.026	8.904	0.012
8	4.666	0.097	13.595	0.001	14.208	0.003	0.024	0.876	7.206	0.007	16.036	0.000
9	3.221	0.200	13.599	0.001	14.699	0.002	0.053	0.818	6.862	0.009	11.506	0.003
10	3.336	0.189	14.733	0.001	16.346	0.001	0.014	0.906	6.993	0.008	15.021	0.001
11	1.998	0.368	10.801	0.005	11.846	0.008	0.116	0.733	5.572	0.018	9.339	0.009
12	3.234	0.198	13.362	0.001	14.768	0.002	0.019	0.890	6.953	0.008	16.066	0.000
13	2.316	0.314	10.901	0.004	11.623	0.009	0.081	0.775	5.545	0.019	9.324	0.009
14	2.595	0.273	9.756	0.008	10.860	0.013	0.052	0.820	3.993	0.046	6.857	0.032
15	6.779	0.034	16.433	0.000	16.516	0.001	0.944	0.331	11.664	0.001	20.824	0.000
16	3.342	0.188	14.243	0.001	16.033	0.001	0.190	0.663	7.713	0.005	23.605	0.000
17	2.277	0.320	12.555	0.002	15.032	0.002	0.307	0.579	8.187	0.004	24.991	0.000
18	2.105	0.349	11.580	0.003	13.464	0.004	0.026	0.871	5.722	0.017	11.243	0.004
19	6.860	0.032	17.351	0.000	18.075	0.000	0.041	0.839	10.732	0.001	28.672	0.000
20	3.258	0.196	12.010	0.002	13.550	0.004	0.180	0.671	7.802	0.005	19.265	0.000
21	3.604	0.165	14.444	0.001	16.071	0.001	0.021	0.884	6.664	0.010	15.516	0.000
22	1.560	0.458	9.949	0.007	10.972	0.012	0.202	0.653	5.972	0.015	10.967	0.004
23	6.810	0.033	16.124	0.000	16.996	0.001	0.267	0.605	10.880	0.001	19.363	0.000
24	2.772	0.250	12.768	0.002	13.738	0.003	0.008	0.928	6.332	0.012	9.781	0.008
25	3.198	0.202	12.949	0.002	14.345	0.002	0.061	0.805	5.944	0.015	12.184	0.002

Note: p = p-value. J_i, $i = 2, 3, 4$ denotes the J-test for Model i. $J(0.5)$ denotes the J-test for a symmetric quad-quad loss function. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 5: Results for pooled data, lin-lin loss function

Panel A

No.	$\hat{\alpha}_{Model1}$	se	z-test	$\hat{\alpha}_{Model2}$	se	z-test	$\hat{\alpha}_{Model3}$	se	z-test	$\hat{\alpha}_{Model4}$	se	z-test
All	0.213	0.014	-20.099	0.213	0.014	-20.15	0.177	0.013	-24.354	0.173	0.013	-24.892

Panel B

No.	$J_2(0.5)$	p	$J_3(0.5)$	p	$J_4(0.5)$	p	$J_2(\hat{\alpha})$	p	$J_3(\hat{\alpha})$	p	$J_4(\hat{\alpha})$	p
All	272.164	0.000	333.161	0.000	335.654	0.000	0.692	0.406	54.44	0.000	60.998	0.000

Note: Panel A: se = standard error, z-test = test of the null hypothesis that $\hat{\alpha} = 0.5$. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4). Panel B: p = p-value. J_i, $i = 2, 3, 4$ denotes the J-test for Model i. $J(0.5)$ denotes the J-test for a symmetric lin-lin loss function. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 6: Results for pooled data, quad-quad loss function

Panel A

No.	$\hat{\alpha}_{Model1}$	se	z-test	$\hat{\alpha}_{Model2}$	se	z-test	$\hat{\alpha}_{Model3}$	se	z-test	$\hat{\alpha}_{Model4}$	se	z-test
All	0.304	0.025	-7.763	0.301	0.025	-7.907	0.144	0.016	-22.805	0.117	0.012	-32.089

Panel B

No.	$J_2(0.5)$	p	$J_3(0.5)$	p	$J_4(0.5)$	p	$J_2(\hat{\alpha})$	p	$J_3(\hat{\alpha})$	p	$J_4(\hat{\alpha})$	p
All	68.309	0.000	309.472	0.000	343.003	0.000	1.598	0.206	150.832	0.000	296.838	0.000

Note: Panel A: se = standard error, z-test = test of the null hypothesis that $\hat{\alpha} = 0.5$. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4). Panel B: p = p-value. J_i, $i = 2, 3, 4$ denotes the J-test for Model i. $J(0.5)$ denotes the J-test for a symmetric quad-quad loss function. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 7: Results for pooled data, lin-lin loss function, subsample analysis

Panel A

No.	$\hat{\alpha}_{Model1}$	se	z-test	$\hat{\alpha}_{Model2}$	se	z-test	$\hat{\alpha}_{Model3}$	se	z-test	$\hat{\alpha}_{Model4}$	se	z-test
All	0.213	0.019	-14.844	0.21	0.019	-15.119	0.183	0.018	-17.348	0.165	0.018	-19.103

Panel B

No.	$J_2(0.5)$	p	$J_3(0.5)$	p	$J_4(0.5)$	p	$J_2(\hat{\alpha})$	p	$J_3(\hat{\alpha})$	p	$J_4(\hat{\alpha})$	p
All	150.431	0.000	166.248	0.000	173.289	0.000	2.769	0.096	23.904	0.000	39.715	0.000

Note: Panel A: se = standard error, z-test = test of the null hypothesis that $\hat{\alpha} = 0.5$. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4). Panel B: p = p-value. J_i, $i = 2, 3, 4$ denotes the J-test for Model i. $J(0.5)$ denotes the J-test for a symmetric lin-lin loss function. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).

Table 8: Results for pooled data, quad-quad loss function, subsample analysis

Panel A

No.	$\hat{\alpha}_{Model1}$	se	z-test	$\hat{\alpha}_{Model2}$	se	z-test	$\hat{\alpha}_{Model3}$	se	z-test	$\hat{\alpha}_{Model4}$	se	z-test
All	0.184	0.023	-13.986	0.182	0.023	-14.095	0.088	0.014	-28.855	0.075	0.013	-32.878

Panel B

No.	$J_2(0.5)$	p	$J_3(0.5)$	p	$J_4(0.5)$	p	$J_2(\hat{\alpha})$	p	$J_3(\hat{\alpha})$	p	$J_4(\hat{\alpha})$	p
All	122.875	0.000	178.9	0.000	184.693	0.000	1.199	0.274	71.192	0.000	101.06	0.000

Note: Panel A: se = standard error, z-test = test of the null hypothesis that $\hat{\alpha} = 0.5$. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4). Panel B: p = p-value. J_i, $i = 2, 3, 4$ denotes the J-test for Model i. $J(0.5)$ denotes the J-test for a symmetric quad-quad loss function. The instruments used are the following: a constant (Model 1), a constant and the lagged forecast error (Model 2), a constant and the lagged oil price (Model 3), and, a constant, the lagged forecast error, and the lagged oil price (Model 4).