EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 999 / JANUARY 2009

RISK-ADJUSTED FORECASTS OF OIL PRICES

by Patrizio Pagano
and Massimiliano Pisani

EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 999 / JANUARY 2009

RISK-ADJUSTED FORECASTS OF OIL PRICES[1]

by Patrizio Pagano and Massimiliano Pisani[2]

In 2009 all ECB publications feature a motif taken from the €200 banknote.

1 We thank Stephane Dees, Brian Torgunrud and seminar participants at the Bank of Italy, at the Eurosystem Working Group on Forecasting and at the Bank of Canada Workshop on Commodity Price Issues for comments and stimulating discussions. We are indebted to several colleagues, and especially to Fabio Busetti, Paola Caselli, Alberto Locarno, Giuseppe Parigi, Roberto Piazza and Stefano Siviero for many constructive suggestions. Giovanna Poggi provided valuable research assistance. Part of this project was undertaken while Pisani was visiting the ECB Research Department, whose hospitality is gratefully acknowledged. We are solely responsible for any errors. The opinions expressed in this paper do not necessarily reflect those of the Bank of Italy, European Central Bank or the Eurosystem.

2 Both authors: Bank of Italy, via Nazionale 91, I - 00184 Rome, Italy; e-mail: patrizio.pagano@bancaditalia.it; massimiliano.pisani@bancaditalia.it

Address
Kaiserstrasse 29
60311 Frankfurt am Main, Germany

Postal address
Postfach 16 03 19
60066 Frankfurt am Main, Germany

Telephone
+49 69 1344 0

Website
http://www.ecb.europa.eu

Fax
+49 69 1344 6000

ISSN 1725-2806 (online)

CONTENTS

Abstract

This paper documents the existence of a significant forecast error on crude oil futures. We interpret it as a risk premium, which, in part, could have been explained by means of a real-time US business cycle indicator, such as the degree of capacity utilization in manufacturing. This result is robust to the specification of the estimating equation and to the considered business cycle indicator. An out-of-the-sample prediction exercise reveals that futures adjusted to take into account this time-varying component produce significantly better forecasts than those of unadjusted futures, of futures adjusted for the average forecast error and of the random walk, particularly at horizons of more than 6 months.

Keywords: Oil, Forecasting, Futures.
JEL classification: E37, E44, G13, Q4.

Non-technical summary

Oil price is a crucial variable for macroeconomic forecasts. A well established approach to predict the evolution of oil prices relies on oil futures. Alternatively, the last available spot price is commonly used (random walk assumption).

This work shows that oil futures price tend to significantly under-predict the subsequently realized spot price. The difference between the latter and the former is larger the longer the forecasting horizon. We interpret the forecast error as a measure of the oil futures risk premium. Furthermore, we show that this risk premium varies over the business cycle and that could be explained by means of a real-time US business cycle indicator, such as the degree of capacity utilization in the manufacturing sector. The outcome is robust to the specification of the estimating equation and to the considered business cycle indicator. We exploit this result to adjust futures-based forecasts for this time-varying component ("risk-adjusted forecasts").

An out-of-sample forecasting exercise reveals that risk-adjusted forecasts are more precise that those obtained with unadjusted futures, random walk or futures adjusted for a constant risk premium, particularly at horizons longer than 6 months. For example, at the twelve-month horizon the root mean squared forecast error of the risk-adjusted futures was 20 per cent lower than that obtained when using either the random walk assumption or the unadjusted futures. With respect to that obtained with the constant-adjusted futures, the gain was 10 per cent.

1 Introduction

Although the dependency of global economic activity on crude oil has fallen steadily over the last thirty years, the oil price baseline assumption remains an important variable for all macroeconomic forecasts. For example, forecast of future oil prices are crucial in central bank's monetary policy decisions, because they enter the construction of expected inflation and output-gap (Svensson, 2005). The increase in oil prices since mid-2003 (Figure 1), which has surprised most analysts by its rapidity and intensity, prompts a new call to investigate the validity of the forecasting assumptions.

Figure 1: Oil price

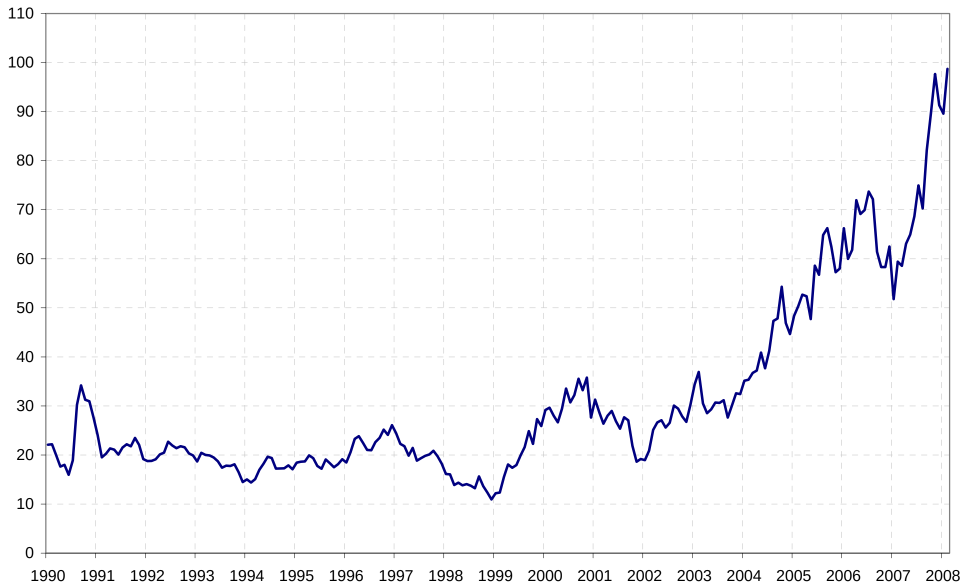

Notes: US dollars per barrel, West Texas Intermediate (WTI) grade. Monthly observations. Each observation is the simple average daily spot prices during the third week of the month.

A commonly used approach to forecast oil prices relies on futures contracts. The notion that the futures price is the optimal forecaster of the future spot price is a by-product of the expectations hypothesis, which assumes efficient (and rational) financial markets. It implies that the futures price should be equal to the expected future spot price and, as a consequence, the forecast error should be zero on average (unbiasedness property of the forecaster) and uncorrelated with any variable in the information set at the time the forecast is made.

However, there is a large and growing literature on financial markets (see e.g. Cochrane, 2005, for a survey) that has challenged the expectations hypothesis. In principle, market participants could be not rational or have a rational learning behavior. More importantly for the purpose of this paper, rejection of forecast unbiasedness could mean that, even if expectations are rational, futures prices contain a not negligible and possibly time-varying risk premium component. For instance, a number of studies (e.g. Cochrane and Piazzesi, 2005) show that excess returns on US treasuries are high in recessions and low in booms. In this paper we argue that excess returns on oil futures may be the outcome of time-varying risk premia, too.

Building on a methodology introduced by Piazzesi and Swanson (2008) to explain the excess return on federal funds futures, we show that the expectations hypothesis fails also for oil futures and that there is a systematic forecast error, interpretable as the negative of the risk premium.[1] The latter is on average positive and, because of the sensitivity of oil prices to the US business cycle, predictable by proxies of macroeconomic conditions, such as the level of capacity utilization in manufacturing.[2] Besides its widespread use as an indicator of the state of the business cycle, the degree of capacity utilization is known for its high complementarity with energy consumption, as emphasized by Finn (2000) in a study on the effect of energy price increases on economic activity. Our findings are robust to the specification used to estimate the sensitivity of oil prices to the business cycle and to the choice of indicators of macroeconomic or oil market conditions.

We assess the forecasting performance of our approach on the basis of an out-of-the-sample prediction exercise. Results show that forecasts adjusted to take into account the time-varying risk premium (that we dub "risk-adjusted forecasts") display lower mean and root-mean squared errors than those of the unadjusted futures, of the futures adjusted for the average forecast error ("constant-adjusted") and of the random walk hypothesis, particularly at horizons of over 6 months.

[1] As in Piazzesi and Swanson (2008) we use in the paper the label "risk-premium" quite loosely, referring to the part of the forecast error that could be predicted.

[2] Barsky and Kilian (2002) note that oil prices are endogenous with respect to macroeconomic variables.

We are not the first to have found that futures may yield biased forecast of oil prices. In the framework of the marginal convenience yield, on the basis of estimates of the oil risk-adjusted discount rate, Pindyck (2001) calculates that the 6-month futures contract should under-predict the realized spot price by around 3 to 4.5 per cent. Erb and Harvey (2006) and Gorton and Rouwenhorst (2006) provide evidence of the relationship between futures and risk premium. The latter work shows that commodity futures risk premium has been equal in size to the historical risk premium of stocks (the equity premium) and has exceeded the risk premium of bonds.[3] These works concentrate on the average risk premium, neglecting its time-variability. On the contrary, Moosa and Al-Lougani (1994), focusing on the properties of oil spot and futures prices in the context of co-integration, find that there is a time-varying risk premium that can be adequately modelled by a GARCH process. Consistent with this result, Considine and Larson (2001) suggest that crude oil assets contain a risk premium that rises sharply with higher price volatility. Other works relate the risk premium variation to macroeconomic factors. Bailey and Chan (1993) find that commodity risk premia co-vary with stock and bond market variables, reflecting economy-wide risk factors. Coimbra and Esteves (2004) provide evidence of a correlation between oil futures forecast errors and market expectation errors on world economic activity. Roache (2008) finds that commodity futures offer a hedge against lower interest rates and that investors are willing to accept lower expected returns for this position. Gorton *et al.* (2007) find that the state of inventories, that vary with business cycles (Fama and French, 1988), are informative about commodity futures risk premia. Consistently with this approach, this paper directly relates the risk premium to the current state of macroeconomic conditions and systematically analyzes the implications for the futures forecasting properties.

The rest of the paper is organized as follows. In the next Section we document the size of the ex post forecast errors on oil futures, showing that these display a non trivial cyclical component. In Section 3 we perform several robustness exercises both on the specification and of the business

[3] There are also studies on the efficiency of the oil futures market and on the forecasting properties of the futures that reached opposite conclusions. For example, Chinn et al. (2005) find that over the period January 1999-October 2004 futures prices are unbiased predictors of crude oil, even if futures typically explain only a small proportion of the variability in oil price movements. However, using the same methodology as Chinn et al. (2005), over the April 1989-December 2003 period Chernenko et al. (2004) find mixed evidence on the existence of risk premia associated with oil futures. More recently, Alquist and Kilian (2008) document that oil futures prices are less accurate than the random walk assumption in forecasting future spot prices.

cycle indicators considered. In Section 4 we propose a method to adjust the forecast based on oil futures and evaluate its performance with respect to the unadjusted futures and other alternatives. Section 5 contains some concluding remarks.

2 Risk premia on crude oil futures

This paper focuses on the question whether $f_t^{(n)}$ – the price at time t of a futures on an oil barrel for delivery n periods ahead – is the best predictor of the oil spot price realized at $t+n$ (denoted p_{t+n}).

Neglecting marking-to-market considerations, a long position on oil futures has no initial cost and a random payoff of $\left(f_t^{(n)} - p_{t+n}\right)$. Standard no-arbitrage condition requires that

$$E_t\left[m_{t+n}\left(f_t^{(n)} - p_{t+n}\right)\right] = 0, \tag{1}$$

where m_{t+n} is the stochastic pricing kernel and t subscripts throughout denote conditioning on the information set at time t. Rearranging gives

$$f_t^{(n)} = E_t\left(p_{t+n}\right) - \frac{Cov_t\left(p_{t+n}, m_{t+n}\right)}{E_t\left(m_{t+n}\right)} \tag{2}$$

which says that $f_t^{(n)}$ equals the expectation for the future realized value – the first term on the right hand side of (2) – minus the risk premium (the second term).

Our approach in this paper is a purely statistical one, using the standard tools of forecast evaluation. Since the main goal of the paper is to provide a method to forecast oil prices, we just check whether using futures prices as predictor of subsequent oil prices produces systematic errors. If futures are unbiased expectations of future prices, then the forecast errors must have mean zero and must be uncorrelated with any variable in the information set at the time that the forecast was made.

In the remaining of this section we first estimate the mean of the ex-post realized forecast error, i.e. the average of

$$fe_{t+n}^{(n)} = f_t^{(n)} - p_{t+n}. \tag{3}$$

and test whether this is significantly different from zero. We interpret a non-zero average as a risk premium, which is the negative of the forecast error. We then test whether the risk premium contains a component that varies over time. If there is a significant, possibly time-varying, risk premium, futures-based forecasts of oil prices should be opportunely adjusted.

2.1 Bias

In the following analysis we use oil price futures on the West Texas Intermediate (WTI) grade. They trade on the New York Mercantile Exchange (NYMEX) and are settled each month. The contract provides for the physical delivery of 1,000 barrels of oil at any point during the settlement month. We use monthly data for a sample period of over 17 years, from January 1990 to February 2007, for which futures data are available at all maturities from 2 to 12 months[4]. We neglect futures prices at 1-month maturity as they are almost always indistinguishable from the corresponding spot prices. As is apparent in Figure 1 the sample is long enough to display periods of relative price stability, periods of sharp price decreases and periods of prolonged price increases.

We first estimate the average forecast error and test whether it is statistically different from zero, that is we test if there is a non-zero risk premium.

To test whether futures are unbiased predictors of subsequently realized oil prices we run, for each horizon n, the following regression:

$$fe_{t+n}^{(n)} = \alpha^{(n)} + \varepsilon_{t+n}^{(n)}, \tag{4}$$

where α is a constant measuring the average ex post realized forecast error and ε is an error term.

To compute the dependent variable we take the simple average of futures daily quotations during the third week of each month t. The choice is suggested to avoid possible daily outliers. The week selected is the third because, as it will be clear below, it is the closest to the release of relevant business-cycle indicators. However, all the results also hold true if we sample the data on a particular day (the 15^{th}) of each month.

Given that futures contract overlap induces autocorrelation, we compute standard errors using Newey-West autocorrelation- and heteroskedasticity-consistent (HAC) standard errors, with a lag truncation parameter equal to $2(n-1)$, and throughout the paper we always report $t-$statistics based on those standard errors.[5]

Table 1 presents the results. Futures are not unbiased predictors of future oil spot prices: the mean forecast error at each forecast horizon n is

[4] Correspondingly, the last spot price considered for constructing the forecast errors (fe_{t+n}) is dated February 2008.

[5] We also computed standard errors using the heteroskedasticity- and autocorrelation-consistent procedure from Hodrick (1992), allowing for $n-1$ lags of excess returns to be serially correlated due to contract overlap, and obtained almost identical results.

Table 1: Average futures-based forecast errors

n	constant	
2	-0.43	(-1.3)
3	-0.73	(-1.6)
4	-1.09	(-1.8)
5	-1.49	(-2.0)
6	-1.89	(-2.0)
7	-2.32	(-2.1)
8	-2.73	(-2.2)
9	-3.15	(-2.2)
10	-3.55	(-2.3)
11	-3.93	(-2.3)
12	-4.37	(-2.3)

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - p_{t+n}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

significantly negative, ranging from 43 cents to more than \$4, and longer-horizon contracts display larger values. The values displayed in table 1 imply that a six-month contract under-predicts the realized spot by \$1.89, or 6.7 per cent if evaluated at the mean price of the sample.

2.2 Capacity utilization

Up to this point we have documented the presence of a significant forecast error in oil price futures. To investigate whether business cycle phases help explaining realized futures-based forecast errors, we run the following regression:

$$fe_{t+n}^{(n)} = \alpha^{(n)} + \beta^{(n)} UCap_{t-1} + \varepsilon_{t+n}^{(n)}, \tag{5}$$

where $UCap$ is the degree of capacity utilization in US manufacturing, which is a proxy of the US business cycle. We focus on the US since it is the largest world oil consumer. Furthermore, we want to use data known to market participants at the time future contacts are subscribed, that is in month t. Since business cycle indicators are subject to several backward revisions, we rely on real-time series. Given that US capacity utilization values are released by the Federal Reserve around the 15^{th} day of each month for the

Table 2: Futures-based forecast errors and real-time capacity utilization

n	constant	capacity$_{t-1}$			R^2	
2	-11.41	0.14	(1.5)	[0.178]	0.01	{50.5}
3	-19.06	0.23	(1.9)	[0.076]	0.03	{49.7}
4	-26.71	0.32	(2.2)	[0.051]	0.04	{51.3}
5	-35.16	0.42	(2.4)	[0.032]	0.06	{55.9}
6	-44.94	0.54	(2.6)	[0.027]	0.08	{61.4}
7	-54.15	0.65	(2.8)	[0.024]	0.10	{66.6}
8	-64.66	0.78	(2.9)	[0.028]	0.12	{69.5}
9	-74.19	0.90	(3.0)	[0.033]	0.14	{71.2}
10	-82.99	1.00	(3.0)	[0.041]	0.16	{71.8}
11	-91.87	1.11	(3.0)	[0.051]	0.18	{69.1}
12	-98.85	1.19	(3.0)	[0.059]	0.19	{68.7}

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - p_{t+n}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses. Bootstrapped p-values for the t-statistics in square brackets; percentages of bootstrapped R^2 values exceeding the actual ones in curly brackets.

previous month, we date our $UCap$ variable as $t-1$.[6] Table 2 reports the results. Risk premia (the negative of forecast errors) and capacity utilization are negatively related. The absolute value of the slope coefficient increases with the maturity of the contract and is statistically significant, at 5 per cent or lower level, from the 4-month horizon on.[7]

The HAC $t-$statistics we report may still be plagued by small-sample distributional properties. To tackle this problem, we do a bootstrap following Horowitz (2004): we resample observations with replacement to generate 50,000 synthetic samples of the same size as the original data set. To account for possible serial dependence of the data generating process, we resample the data in blocks, with block size equal to $n+1$ for each horizon n.[8] The bootstrap $p-$values, reported in square brackets in table 2, are

[6] See the appendix for further details.

[7] We also run predictability regressions including the own futures $f_t^{(n)}$, but results remain unchanged. When we run regression (5) using the last available revised vintage of the degree of capacity utilization point estimates of the slope coefficients are still positive, yet slightly smaller.

[8] To this end we adapted a routine by Eric Swanson, whose cooperation is gratefully

computed using the distribution of synthetically generated $t-$statistics centered at the actual $t-$statistic. Results support the choice of the degree of capacity utilization in manufacturing as predictor of realized excess-returns on oil futures.

According to the R^2 the percentage of the variance of the forecast error on oil futures explained by this specification is not trivial, especially at longer horizons. For instance, the model with capacity utilization explains almost 20 per cent of the forecast error on oil futures contract with 12-month maturity. One may argue that, especially for the longer-horizon forecasts, the forecast errors are bound to be quite persistent. Since the predictors are also persistent, this may give rise to a problem of near spurious regressions.[9] To check whether this is the case with our regressions, we calculate the percentages of bootstrapped R^2 values exceeding the actual ones. Results, reported in curly brackets, show that actual R^2 values are not unusually large compared with bootstrapped ones, as they remain below or close to the median.

In sum, the results in table 2 suggest that the adjustment to be made over the futures should be smaller during booms and higher during slacks: that is, to obtain risk-adjusted forecasts we should add to the futures a counter-cyclical term.

This time-varying risk premium arises from varying conditional covariance of the asset return with the pricing kernel (see equation (2)). What is the intuition? Cochrane (2001), referring to the evidence that long-run excess returns are quite predictable and that most of the variables forecasting excess returns are correlated with or forecast business cycles, suggests a natural explanation emphasized in Fama and French (1989): expected returns vary over business cycles. It takes a higher risk premium to hold risky assets at the bottom of a recession. Why? It is plausible that risk or risk adversion vary over the business cycle – for instance, it happens when the utility function of investors displays habit formation – and this is exactly the horizon at which we see predictable excess returns.

Further, as recently emphasized also by Gorton *et al.* (2007) there may

acknowledged. In choosing the block lengths there is a tradeoff: bigger block lengths reproduce the serial correlation better, but effectively reduce the number of independent draws from the sample, reducing sampling variation. We have tried to choose reasonable block lengths that capture as much serial correlation as possible without reducing the number of independent bootstrap draws too much. Results are not particularly affected if we change the block lengths.

[9] See Campbell, Lo and Mackinlay (1997) for a discussion of R^2 values in long-horizon predictability regressions.

be a link between risk premia and inventories. The level of inventories matters because, as in Deaton and Laroque (1992), future spot price variance is negatively related to the level of inventories. That is, when inventories are low, the buffer function of inventories to absorb shocks is diminished. In these circumstances, the risk of a stock-out increases which raises the volatility of the future spot price. Because commodity futures are used to insure price risk, inventory theory predicts an increase in the risk premium. Since inventories are relatively low during recessions (e.g. Fama and French, 1988), in turn, this produces a counter-cyclical risk premium.

In the next section we directly analyze the relationship between excess returns on oil futures and the level of inventories.

2.3 Inventories

Consistent with Gorton *et al.* (2007), we use inventories directly as explanatory variables for risk premia. Since we collect oil spot and futures prices as of the third week of each month, we select the level of US oil inventories as of the second week of each month to avoid endogeneity problems. We then estimate

$$fe_{t+n}^{(n)} = \alpha^{(n)} + \theta^{(n)} inventories_{t-1} + \varepsilon_{t+n}^{(n)}. \tag{6}$$

and find that $\theta^{(n)} > 0$ (table 3, panel A). That is, we do find that the ex-post forecast error is positively correlated with the level of inventories or, alternatively, that the realized risk premium is negatively correlated with the level of inventories.

However, when we also add as explanatory variable our preferred business cycle indicator (the degree of capacity utilization) the level of inventories is no longer significant on most horizons, while the sign, the size and the significance of capacity utilization remains almost intact (table 3, panel B). This result indicates that the information content of the level of US oil inventories is spanned by the business cycle indicator that we have already considered in Section 2.2.

Overall we interpret these findings as suggestive that oil futures risk premia co-vary with the level of inventories and across the business cycle. Indeed, as suggested by Gorton *et al.* (2007), problems related to the availability and the poor quality of inventory data, and issues regarding the appropriate definition of relevant inventories, may imply that business cycle indicators could be viewed as a better proxy for scarcity.

Still, as mentioned before, it is well plausible that the relevant explanation for the time variation of risk premia lies with the variability of risk

Table 3: Futures-based forecast errors, inventories and capacity utilization

n	constant	inventories$_{t-1}$		capacity$_{t-1}$		R^2
		(A)				
2	-2.15	0.01	(0.3)			0.00
3	-6.91	0.02	(0.8)			0.01
4	-12.22	0.03	(1.3)			0.03
5	-18.95	0.05	(1.8)			0.05
6	-26.14	0.08	(2.2)			0.09
7	-31.60	0.09	(2.4)			0.11
8	-36.70	0.11	(2.6)			0.13
9	-39.90	0.12	(2.6)			0.13
10	-41.22	0.12	(2.5)			0.13
11	-41.54	0.12	(2.2)			0.12
12	-43.23	0.12	(2.0)			0.11
		(B)				
2	-11.23	-0.00	(-0.2)	0.15	(1.8)	0.01
3	-19.62	0.01	(0.3)	0.21	(1.8)	0.03
4	-28.15	0.02	(0.8)	0.26	(1.8)	0.05
5	-37.79	0.04	(1.2)	0.30	(1.8)	0.08
6	-48.81	0.06	(1.6)	0.37	(2.0)	0.12
7	-58.84	0.07	(1.8)	0.44	(2.1)	0.15
8	-70.00	0.08	(2.0)	0.54	(2.3)	0.18
9	-79.73	0.08	(2.0)	0.64	(2.4)	0.19
10	-88.29	0.08	(1.8)	0.76	(2.5)	0.21
11	-96.67	0.07	(1.5)	0.89	(2.5)	0.22
12	-103.61	0.07	(1.4)	0.97	(2.5)	0.22

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - p_{t+n}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

or risk adversion over the business cycle. Identifying the exact nature of this correlation is beyond the scope of the current paper, which focuses on improving on the forecasting performance of oil futures. Yet, in the next section we provide several sensitivity exercises on the time-varying pattern of oil future risk premia.

3 Robustness analysis

To gauge the sensitivity of the results obtained in the previous section we perform several alternative estimations, changing either the specification of the estimating equation or the explanatory variable. First, we provide further evidence for a predictive relationship using one-month holding period returns on a futures contract as the dependent variable instead of futures forecast errors. Second, we show that a business cycle indicator is relevant even in the framework of the so-called "price spread" specification (see e.g. Chernenko *et al.* 2004). Third, we use alternative business cycle indicators. Finally, we estimate equation (5) allowing both constant and slope coefficients to differ according to the oil market being in backwardation or in contango. Overall, we find that results are robust across the various experiments.

3.1 One-month holding period returns

An alternative way to increase the number of independent observations in our regressions and check the robustness of our results is to consider the excess returns an investor would realize from holding an $n-$month ahead oil futures contract for just one month — by purchasing the contract and then selling it back as an $(n-1)-$month ahead contract in one month's time — rather than holding the contract all the way through to maturity. As suggested by Piazzesi and Swanson (2008), by considering one-month holding period returns we reduce potential problems of serial correlation and sample size for the longer-horizon contracts, and give ourselves more than 200 completely independent windows of data (under the null hypothesis of no predictability of the risk premium) for all contracts.

We thus consider regressions of the form

$$f_t^{(n)} - f_{t+1}^{(n-1)} = \alpha^{(n)} + \beta^{(n)} UCap_{t-1} + \varepsilon_{t+n}^{(n)}, \tag{7}$$

where $f_t^{(n)}$ denotes the $n-$month-ahead average contract price on the third week of month t, $f_{t+1}^{(n-1)}$ denotes the $(n-1)-$month-ahead contract price on

Table 4: One-month holding period returns and capacity utilization

n	constant	capacity$_{t-1}$		R^2
2	-6.73	0.08	(1.5)	0.01
3	-5.65	0.07	(1.3)	0.01
4	-6.34	0.08	(1.6)	0.01
5	-6.71	0.08	(1.8)	0.01
6	-6.72	0.08	(2.0)	0.01
7	-6.61	0.08	(2.0)	0.01
8	-6.44	0.08	(2.1)	0.01
9	-6.27	0.07	(2.2)	0.01
10	-6.18	0.07	(2.2)	0.01
11	-5.97	0.07	(2.3)	0.01
12	-5.91	0.07	(2.3)	0.01

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - f_{t+1}^{(n-1)}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

the third week of month $t+1$, and the difference between these two prices is the ex post realized one-month holding period return on the $n-$month-ahead contract. Using specification (7), the residuals are serially uncorrelated under the null hypothesis of no predictability of excess returns, because all variables in equation (7) are in financial markets' information set by the beginning of the next period.

Results, presented in table 4, show that the degree of capacity utilization is a significant predictor of such excess returns. As expected from quasi-differencing of the left-hand-side variable R^2 values are uniformly quite small.

3.2 Price spread specification

Previous studies on the forecasting ability of oil futures have applied the price-spread specification, also known as the Mincer-Zarnowitz regression (Mincer and Zarnowitz, 1969). This is a regression of the change of spot prices $(p_{t+n} - p_t)$ on the difference between the futures and current spot price (also known as "basis", $f_t^{(n)} - p_t$).

To compare our results with those works that adopted that framework (e.g. Chinn *et al.*, 2005 , Chernenko *et al.*, 2004) we run the same regression,

Table 5: Price spread , basis and capacity utilization

n	constant	basis		capacity$_{t-1}$		R^2
2	11.90	1.35	(2.4)	-0.14	(-1.6)	0.07
3	19.11	1.02	(2.1)	-0.23	(-1.9)	0.08
4	26.43	0.94	(2.0)	-0.32	(-2.2)	0.10
5	34.04	0.83	(1.8)	-0.41	(-2.5)	0.11
6	41.91	0.66	(1.7)	-0.51	(-2.7)	0.10
7	49.38	0.57	(1.6)	-0.60	(-2.8)	0.11
8	59.75	0.62	(1.9)	-0.72	(-3.0)	0.13
9	70.01	0.72	(2.2)	-0.85	(-3.1)	0.15
10	79.90	0.81	(2.4)	-0.96	(-3.1)	0.19
11	90.61	0.93	(2.6)	-1.09	(-3.1)	0.23
12	101.41	1.13	(2.8)	-1.22	(-3.1)	0.26

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $p_{t+n} - p_t$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

augmenting it with the variable intended to capture the time-varying risk premium, Specifically we run

$$p_{t+n} - p_t = \alpha^{(n)} + \beta^{(n)} UCap_{t-1} + \gamma^{(n)} \left(f_t^{(n)} - p_t \right) + \varepsilon_{t+n}^{(n)}. \tag{8}$$

Results are reported in table 5.

The degree capacity utilization is highly significant in explaining subsequent realized oil price changes. Such coefficients are negative and have absolute values increasing with the horizon. The negative sign implies that when capacity utilization is low oil prices are predicted to increase, while when capacity utilization is high oil prices are expected to decrease. This result also means that even with a more "classical" price-spread specification, the state of the business-cycle is informative on oil price developments. Overall, we interpret this evidence as highly supportive of the previous results.

3.3 Other business cycle indicators

To check the sensitivity of the results to the explanatory variable in regression (5) we use alternative business cycle indicators. In particular, we try to capture the cyclical variability of the oil futures risk premium using the bond yield curve in the US or a world leading indicator produced by the OECD.

Among leading indicators the bond yield curve is often used as a predictor of excess returns in the Treasury markets. We select three different term spreads based on the difference between 1, 2, 5 and 10 year constant maturity Treasury yields (annualized). As is evident in table 6 oil futures risk premia are significantly correlated with those spreads, with the sign of estimated coefficients always consistent across the different horizons. R^2 values range from 4 to 15 per cent.

Another possible driving factor of the forecast error of oil futures could be related to the oil demand originating not only from the US, but also from other areas, especially from some developing countries. In order to capture the cyclical conditions in the whole world, we use an indicator of global economic activity the composite leading indicator (CLI) constructed by the OECD for the aggregate of the member economies and the six major non-member economies (Brazil, China, India, Indonesia, Russia and South Africa).[10] Results in table 7 show that the coefficients of the composite leading indicator are positive and statistically significant, especially at horizons larger than 5-month. R^2 values are uniformly lower than those obtained using the US degree of capacity utilization, reaching a maximum of 10 per cent per cent at $n = 12$. Interestingly, the statistical significance of the latter indicator of global demand disappears when the degree of capacity utilization in the US manufacturing is included simultaneously in the regression[11].

3.4 Contango vs. backwardation

One may then ask whether the relationship between risk premia and the business cycle is non-linear: in fact, low inventory levels for a commodity are associated with an inverted ("backwardated") term structure of futures prices, while high levels of inventories are associated with an upward sloping futures curve ("contango").

We test whether our estimates are affected on the status of the oil market, that is we consider the possibility that risk premia may behave differ-

[10] We use the "ratio to trend" series. It refers to the deviation from the long-term trend of the series and focuses on the cyclical behavior of the indicator.

[11] In a previous working paper version we also report results with the year-on-year growth in US non-farm payrolls, whose data are also available in real time. Piazzesi and Swanson (2008) use such a variable as their preferred business cycle indicator to predict excess returns on the federal funds futures. Results are broadly consistent, even if estimated coefficients of employment growth display larger standard errors than those of capacity utilization.

Table 6: Futures-based forecast errors and Treasury spreads

n	constant	2yr-1yr		5yr-2yr		10yr-5yr		R^2
2	0.88	-6.19	(-2.2)	5.63	(1.8)	-5.38	(-1.5)	0.04
3	1.28	-8.88	(-2.5)	7.95	(2.2)	-7.86	(-1.9)	0.07
4	1.53	-11.08	(-2.8)	11.40	(3.1)	-12.19	(-2.8)	0.08
5	1.19	-10.63	(-2.3)	11.85	(2.7)	-13.41	(-2.6)	0.07
6	1.14	-11.06	(-2.1)	14.01	(2.5)	-16.95	(-2.7)	0.08
7	0.94	-10.88	(-1.9)	16.33	(2.4)	-20.99	(-2.7)	0.10
8	0.85	-11.41	(-1.8)	19.12	(2.7)	-25.25	(-3.0)	0.12
9	0.26	-9.41	(-1.3)	18.93	(2.5)	-26.47	(-2.9)	0.12
10	-0.30	-7.18	(-0.8)	18.29	(2.0)	-27.29	(-2.6)	0.14
11	-1.12	-3.55	(-0.4)	16.04	(1.5)	-26.47	(-2.1)	0.15
12	-2.06	-0.09	(-0.0)	13.45	(1.1)	-24.82	(-1.9)	0.15

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - p_{t+n}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

Table 7: Futures-based forecast errors and world leading indicator

n	constant	CLI_{t-1}		R^2
2	-26.87	0.26	(1.3)	0.01
3	-43.27	0.43	(1.5)	0.02
4	-56.54	0.55	(1.7)	0.03
5	-70.36	0.69	(1.8)	0.04
6	-90.24	0.88	(2.0)	0.06
7	-99.70	0.97	(2.0)	0.06
8	-115.24	1.13	(2.1)	0.07
9	-127.84	1.25	(2.1)	0.07
10	-145.15	1.42	(2.2)	0.09
11	-159.26	1.55	(2.2)	0.10
12	-169.84	1.65	(2.1)	0.10

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - p_{t+n}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

Table 8: Futures-based forecast errors and capacity utilization: contango vs. backwardation

n	constant	Dcontango		capacity$_{t-1}$		Dcon×cap$_{t-1}$		R^2
2	-10.71	-2.88	(-0.2)	0.13	(1.3)	0.04	(0.2)	0.01
3	-18.03	-0.88	(-0.0)	0.22	(1.7)	0.02	(0.1)	0.03
4	-24.69	-4.57	(-0.2)	0.29	(1.8)	0.06	(0.2)	0.04
5	-33.42	-2.31	(-0.1)	0.40	(2.0)	0.04	(0.1)	0.06
6	-40.96	-3.61	(-0.1)	0.48	(2.1)	0.07	(0.1)	0.10
7	-48.22	-5.40	(-0.2)	0.56	(2.3)	0.10	(0.2)	0.14
8	-59.69	1.29	(0.0)	0.70	(2.5)	0.02	(0.1)	0.16
9	-70.92	8.25	(0.3)	0.84	(2.7)	-0.06	(-0.2)	0.18
10	-81.68	13.98	(0.4)	0.97	(2.8)	-0.14	(-0.3)	0.19
11	-91.72	13.73	(0.4)	1.10	(2.8)	-0.14	(-0.3)	0.20
12	-102.61	22.66	(0.6)	1.23	(2.8)	-0.27	(-0.5)	0.19

Notes: n denotes the horizon of the oil futures contract in months. The dependent variable is $f_t^{(n)} - p_{t+n}$. Sample 1990:1-2007:2. Estimation by OLS, t-statistic based on HAC standard error in parentheses.

ently when the market is commonly described as in backwardation or as in contango. To this end we construct a contango dummy (D_c), equal to one whenever $f_t^{(6)} > p_t$, that is when the futures price with maturity 6 months is larger than the current spot price, and zero otherwise. In our sample it happens 72 times, representing 43 per cent of the total.

We estimate the following equation

$$fe_{t+n}^{(n)} = \alpha^{(n)} + \theta^{(n)} D_c + \beta^{(n)} UCap_{t-1} + \gamma^{(n)} \left(D_c \times UCap_{t-1}\right) + \varepsilon_{t+n}^{(n)}. \quad (9)$$

Results, reported in table 8, show that the contango dummy is never significant, neither by itself nor interacted with the capacity utilization. This means that both the constants and the slopes estimated at any horizon are stable across states of the oil market.

We can thus turn to evaluating the forecasting performance of our model.

4 Predictability of oil prices in real time

Having documented the presence of a significant and time-varying risk premium on oil futures, in this section we evaluate the forecast of oil prices ($\hat{p}_{t+n}$) comparing four alternative methodologies:

Figure 2: oil price forecasts and realized spot prices on two dates

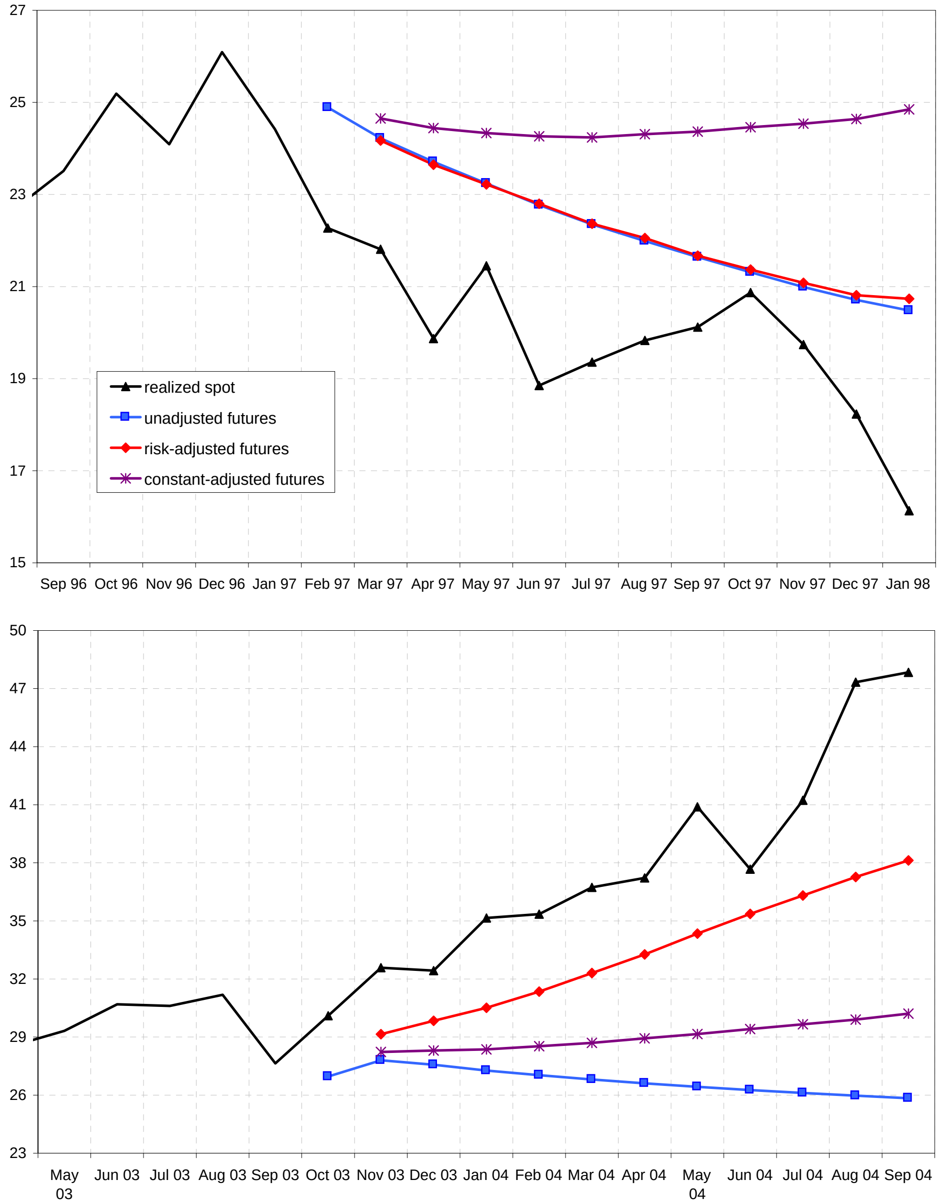

Notes: US dollars per barrel. Risk-adjusted forecasts are computed using estimated coefficients as in table 2. Constant-adjusted forecasts are computed from the estimated coefficient as in table 1.

1. random walk, which implies $\hat{p}_{t+n} = p_t$;
2. unadjusted futures: $\hat{p}_{t+n} = f_t^{(n)}$;
3. constant-adjusted futures, based on regression (4): $\hat{p}_{t+n} = f_t^{(n)} - \hat{\alpha}^{(n)}$;
4. risk-adjusted futures, based on regression (5): $\hat{p}_{t+n} = f_t^{(n)} - \hat{\alpha}^{(n)} - \hat{\beta}^{(n)} UCap_{t-1}$.

A first assessment of the forecasting performance of these different methodologies can be obtained by looking at figure 2. It shows forecasts of oil prices in two illustrative months, January 1997 and September 2003. In January 1997 (upper panel) the oil spot price was around \$26 and, according to futures, oil prices were expected to decline to just over \$20 by the following January. Demand was very high and capacity utilization in manufacturing was running well above the historical average, at almost 83 per cent. The risk-adjusted procedure predicts that the risk premium required over the futures would have been very low. In fact, risk-adjusted futures were virtually indistinguishable from unadjusted futures. By contrast, the constant-adjusted forecast would have signaled roughly constant prices. Indeed, by January 1998 the oil price declined, to \$16.3.

In the Summer of 2003 (figure 2, lower panel) oil prices were stable at around \$30. In September of that year futures pointed to a decline in oil price to just below \$26 in the following 12 months. The recovery out of the recession in 2001 was not yet firmly established, the capacity utilization index was still relatively low, at around 73 per cent, and the risk premium was correspondingly sizeable. In fact, the risk-adjusted forecast would have signaled an oil price as high as more than \$38 by September 2004. Note that not taking into account the cyclical factor – as the constant-adjusted forecast does – would have yielded just slightly increasing oil prices. Indeed oil prices did rise and at the end of the horizon were at around \$47. In order to perform a more formal comparison of the forecasting ability of the different methodologies we run a set of rolling "out-of-sample" regressions. First of all we calculate rolling-endpoint (or expanding window) real-time forecasts, initializing our estimates by using the first 30 observations. We then compute one-step-ahead forecasts for each maturity of the futures in our sample.

To gauge a quantitative measure of how different these four forecasting methodologies are, in table 9, panel (A) we report some summary statistics

Table 9: Test statistics for different forecasts of oil price: expanding window

	benchmark random walk		futures based unadjusted		constant-adjusted		risk-adjusted	
n	ME	RMSE	ME	RMSE	ME	RMSE	ME	RMSE
(A) expanding window								
2	0.44	3.85	0.49	3.81	0.33	3.79	0.24	3.78
3	0.69	4.50	0.80	4.49	0.43	4.44	0.31	4.39
4	0.96	5.07	1.16	5.14	0.58	5.02	0.39	4.92
5	1.27	5.73	1.59	5.92	0.81	5.73	0.53	5.56
6	1.55	6.29	2.00	6.63	1.00	6.35	0.62	6.09
7	1.91	6.80	2.48	7.25	1.30	6.85	0.79	6.50
8	2.31	7.64	2.99	8.08	1.77	7.56	1.15	7.09
9	2.75	8.48	3.54	8.88	2.34	8.23	1.60	7.65
10	3.15	8.96	4.06	9.33	2.84	8.54	1.99	7.82
11	3.55	9.52	4.56	9.85	3.34	8.92	2.37	8.07
12	4.01	10.42	5.11	10.62	3.85	9.54	2.78	8.62
(B) moving window								
2	0.44	3.85	0.49	3.81	-0.06	3.89	0.02	4.00
3	0.69	4.50	0.80	4.49	-0.12	4.61	0.07	4.58
4	0.96	5.07	1.16	5.14	-0.13	5.26	0.02	5.07
5	1.27	5.73	1.59	5.92	-0.09	6.01	0.00	5.65
6	1.55	6.29	2.00	6.63	-0.08	6.66	0.05	6.11
7	1.91	6.80	2.48	7.25	0.02	7.17	0.01	6.47
8	2.31	7.64	2.99	8.08	0.16	7.84	0.09	7.06
9	2.75	8.48	3.54	8.88	0.33	8.42	0.13	7.60
10	3.15	8.96	4.06	9.33	0.47	8.60	0.12	7.72
11	3.55	9.52	4.56	9.85	0.59	8.82	0.13	7.89
12	4.01	10.42	5.11	10.62	0.77	9.25	0.10	8.35

Notes: n is the forecasting horizon. Sample 1990:1-2007:2. ME is the mean error and RMSE is the root mean squared error (both in US dollars). For each model, forecast errors are derived from an expanding window in panel (A), with an initial window of 30 observations, and from a moving (rolling) window of 30 monthly observations in panel (B).

on forecast errors, namely the mean error – defined as the difference between the realized and the forecast price $(p_{t+n} - \hat{p}_{t+n})$ – and the root mean squared error of the forecast. The mean error of the risk-adjusted forecast is the lowest one, both at short and long horizons. For instance at the 6-month horizon the mean error of risk-adjusted futures is around 60 cents, compared with $1 for the constant-adjusted, $1.55 for the random walk and $2 for the unadjusted futures. At the 12-month horizon the mean forecast error committed by the risk-adjusted futures is $2.78, compared with $3.85 for the constant-adjusted futures, $4 for the random walk, and more than $5 for the unadjusted futures. A similar conclusion can be drawn on the basis of root mean squared errors, which for the risk-adjusted forecast are always below those implied by the other three forecasting techniques considered. In particular, at the 12-month horizon risk-adjusted forecast mean squared errors are 10 per cent below the constant adjusted, 17 per cent below the random walk and 19 per cent below the unadjusted futures. Interestingly, the random walk assumption seems to produce better forecasts than the unadjusted futures, in terms of both mean and root mean squared errors, a result consistent with Alquist and Kilian (2008). The constant-adjusted futures performs better than the random walk in terms of mean errors, while in terms of root mean squared errors results are mixed.

To evaluate whether forecast obtained with the risk-adjusted futures are also statistically significantly more accurate than those produced by the other three methodologies we first repeat the above calculations on a moving (rolling) window of 30 observations, which corresponds to roughly one sixth of the sample. We then perform a battery of tests. Details of mean and root mean squared statistics of the four forecasting methodologies are reported in table 9, panel (B). First of all one must notice that mean errors of risk-adjusted futures are very low, never larger than 13 cents. Mean errors of constant-adjusted futures are negative at short horizons – i.e. adjusting futures for constant risk premia tends to over-predict oil prices up to $n = 6$ – and positive at longer ones. In absolute terms they are relatively low, yet larger than those obtained by adjusting futures for time-varying risk premia. In terms of root mean squared errors the gains obtained with risk-adjusted futures are similar to those derived with the expanding window.

In table 10 we present results of the unconditional test of predictive ability proposed by Giacomini and White (2006).[12] It asks which forecast

[12] Viewing the difference in forecast performance (e.g., squared prediction error) as the dependent variable in a regression containing only a constant, it is like a test for whether the regression intercept is zero

Table 10: Giacomini and White (2006) test of predictive accuracy: various methods vs. risk-adjusted futures

n	random walk	unadj. futures	const–adj. futures
2	0.61	1.66	1.07
3	0.07	0.01	0.06
4	0.00	0.10	0.64
5	0.10	0.97	1.46
6	0.41	2.61	2.23
7	1.08	4.18**	2.95*
8	2.35	5.31**	2.74*
9	4.85**	7.51***	2.82*
10	9.15***	10.62***	2.74*
11	11.42***	12.77***	2.87*
12	12.95***	14.27***	2.70*

Notes: n is the forecasting horizon. Values reported in the table are test statistic of the equal unconditional accuracy of forecasts: the null hypothesis is that the forecasts of the method indicated in each column are as accurate as those of the risk-adjusted futures; * denotes significance at 10 per cent; ** denotes significance at 5 per cent; *** denotes significance at 1 per cent. Forecast errors are derived from moving window (30 monthly observations) estimates of each model.

was more accurate, on average, in the past; it may thus be appropriate for making recommendations about which forecast may be better for an unspecified future date. With respect to the classical Diebold and Mariano (1995) statistic, this unconditional test of predictive ability provides primitive conditions that ensure its validity and extends it to an environment permitting parameter estimation. It turns out that risk-adjusted forecasts are more accurate – at the 5 per cent or lower level – than those of the unadjusted futures at all horizons larger than 6-month. The risk-adjusted method outperforms the random walk at all horizons larger than 8-month. With respect to the constant-adjusted futures the evidence is weaker, as risk-adjusted futures are more accurate at all horizon larger than 6-month, but only at the 10 per cent level.

To further investigate on the relative performance of constant-adjusted and risk-adjusted forecast we also run an encompassing test. In fact, since the model with the constant-adjusted futures can be viewed as a particular case (with $\beta = 0$) of the model including also the utilized capacity, in

Table 11: Clark and McCracken (2001) tests of predictive accuracy and encompassing: constant-adjusted vs. risk-adjusted futures

n	MSE-F	ENC-F
2	-9.46	3.99*
3	2.55***	15.57***
4	12.85***	25.44***
5	22.83***	33.48***
6	33.28***	42.02***
7	40.11***	47.83***
8	41.08***	50.54***
9	40.32***	50.73***
10	42.68***	55.50***
11	44.09***	60.40***
12	39.78***	57.37***

Notes: n is the forecasting horizon. MSE-F is the value of the equal accuracy of forecasts: the null hypothesis is that the root mean squared errors of the two models are equal. ENC-F is the value of the test statistic for encompassing of forecasts: under the null the forecasts of the constant-adjusted futures encompass that of the risk-adjusted; * denotes significance at 10 per cent; ** denotes significance at 5 per cent; *** denotes significance at 1 per cent. Forecast errors are derived from moving window (30 monthly observations) estimates of each model.

table 11 we compare root mean squared errors produced by the two nested models using two tests, described in Clark and McCracken (2001).[13] Since all the forecasts under consideration are one-step-ahead, the critical values reported in Clark and McCracken (2001) can be safely applied. According to the test of equal accuracy of prediction, risk-adjusted forecasts are more accurate than the constant-adjusted counterparts from $n = 3$ on at the 1 per cent level. Furthermore, the ENC-F statistics indicate that from the $3-$month horizon the degree of capacity utilization has predictive content for the futures-based forecast errors.

[13] In fact, with nested models, properties of test such as the Diebold-Mariano are likely to differ because, under the null, the forecast errors are asymptotically the same and therefore perfectly correlated.

5 Concluding remarks

This paper documents that crude oil futures display a significant ex post forecast error, which is negative on average. We also show that this forecast error has a non trivial cyclical component which can be, in part, explained by means of real-time US business cycle indicators, such as the degree of utilized capacity in manufacturing. Results appear robust to various checks such as the use of alternative specifications of the estimating equation and the consideration of different business cycle indicators.

Adjusting the oil price forecast embedded into futures to take account of this time-varying risk premium yields "risk-adjusted" forecasts which perform extremely well in periods both of "bear" and of "bull" oil markets. More formally, with an out-of-the-sample prediction exercise we show that the forecast adjusted for a time-varying risk premium – linked to the US business cycle – performs significantly better than the unadjusted futures, the simple constant-adjusted futures and the random walk, particularly at horizons longer than 6 months.

If the forecast error could have been significantly reduced by investors exploiting available information on the US business cycle, as we have shown, the question that naturally arises is why they did not do so. A thorough analysis of such issue is beyond the scope of the current work and is left for future research. Yet, an inspection of net long positions, reported by the US Commodity Futures Trading Commission, held by non-commercial traders, usually referred to as "speculators", reveals that both in 1999-2000 and in late 2003 these positions were largely positive, signaling expectations of rising oil prices, which effectively were realized in the following months.[14] Therefore, it is possible that this category of market participants was aware of this risk premium and provided an insurance to (hedging) commercial market participants. This notwithstanding, a significant part of the premium was not competed away. A possible explanation is that non-commercial traders represent a small percentage (just a little over 10 per cent) of all open interest, since they trade mainly in over-the-counter markets. Alternatively, as suggested by Piazzesi and Swanson (2008) in the context of futures on federal funds, the futures market may not be perfectly competitive or non-commercial traders may themselves be risk adverse.

[14] Formal analysis of energy futures markets (Sanders et al., 2004) reveals that positive futures returns Granger-cause increases in the net long positions held by reporting non-commercial traders. There is no consistent evidence that traders' net long positions contain any general predictive information about market returns, that is net long positions do not generally lead market returns (see also Gorton et al., 2007).

As a final remark we note that our results have crucial implications for policy analysis and economic modelling. First, they point out that futures should be appropriately adjusted for predicting oil prices which, in turn, affect inflation and output gap forecasts, the two variables that, according to modern economic theory, are crucial for monetary policy decisions. Second, our results show that the identification of unexpected oil price changes ("shocks"), which are often used in the context of dynamic macro analyses, as the difference between futures and subsequently realized prices should be taken with caution, since the resulting series may be contaminated by risk premia. Futures-based forecasts of oil prices, adjusted for time-varying risk premia, may be exploited to identify such shocks, but we leave this as an other interesting topic for future research.

Appendix: Data Sources

WTI oil spot (CRUDOIL), and futures prices (NCL...): *Thomson Financial Datastream.*

Real-time US indicators: *Federal Reserve Bank of Philadelphia (www.phil.frb.org/econ/forecast/reaindex.html).*

Term spreads: *Thomson Financial Datastream,* calculated from FRTCM1Y, FRTCM2Y, FRTCM5Y, FRTCM10.

Composite leading indicator: *OECD (http://www.oecd.org/document/34/0,3343,en_2649_34349_38368994_1_1_1_1,00.html)*

Oil inventories: *U.S. Department of Energy, Energy Information Administration (www.eia.doe.gov/dnav/pet/pet_stoc_wstk_dcu_nus_w.htm).*

References

[1] Alquist, Ron and Lutz Kilian (2008), "What Do We Learn from the Price of Crude Oil Futures?", *mimeo.*

[2] Bailey, Warren and K. C. Chan (1993), "Macroeconomic Influences and the Variability of the Commodity Futures Basis", *Journal of Finance*, vol. 48, pp. 555-573.

[3] Barsky, Robert B., and Lutz Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: *NBER*

Macroeconomics Annual 2001, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.

[4] Campbell, J.Y., A.W. Lo and A.C. Mackinlay (1997), *The Econometrics of Financial Markets*, Princeton University Press, Princeton, New Jersey.

[5] Chernenko, Sergey V., Krista B. Schwarz and Jonathan H. Wright (2004) "The Information Content of Forward and Futures Prices: Market Expectations and the Price of Risk", *Board of Governors of the Federal Reserve System, International Finance Discussion Papers, No. 808.*

[6] Chinn, Menzie, Michael LeBlanc, and Olivier Coibion (2005) "The Predictive Content of Energy Futures: An update on Petroleum, Natural Gas, Heating Oil and Gasoline", *NBER working paper No. 11033.*

[7] Clark, Todd E. and Michael W. McCracken (2001), "Tests of equal forecast accuracy and encompassing for nested models", *Journal of Econometrics*, vol. 105, pp. 85–110.

[8] Cochrane, John H. (2001), *Asset Pricing*, Princeton University Press: Princeton, New Jersey.

[9] Cochrane, John H. (2005), "Financial Markets and the Real Economy", *NBER working paper No. 11193.*

[10] Cochrane, John H and Monika Piazzesi. (2005), "Bond Risk Premia", *American Economic Review,* vol. 95, pp. 138-160.

[11] Coimbra, Carlos and Paulo Soares Esteves (2004), "Oil Price Assumptions in Macroeconomic Forecasts: Should We Follow Futures Market Expectations?", *OPEC Review,* vol. 28, pp. 87-106.

[12] Considine, Timothy J. and Donald F. Larson (2001) "Risk Premiums on Inventory Assets: The Case of Crude Oil and Natural Gas", *The Journal of Futures Markets*, vol. 21, pp. 109-126.

[13] Deaton, Angus and Guy Laroque (1992), "On the Behavior of Commodity Prices", *Review of Economic Studies,* vol. 59, pp. 1-23.

[14] Diebold, Francis X. and Roberto S. Mariano (1995), "Comparing predictive accuracy", *Journal of Business and Economic Statistics,* vol. 13, pp. 253-263.

[15] Erb, Claude and Campbell Harvey (2006), "The Strategic and Tactical Value of Commodity Futures", *Financial Analysts Journal,* vol. 62. pp. 69-97.

[16] Fama, Eugene F. and Kenneth R. French (1988), "Business Cycles and the Behavior of Metals Prices", *Journal of Finance,* vol. 43, pp. 1075-1093.

[17] Fama, Eugene F. and Kenneth R. French (1989), "Business conditions and expected returns on stocks and bonds", *Journal of Financial Economics,* vol. 25, pp. 23-49.

[18] Finn, Mary G. (2000), "Perfect Competition and the Effects of Energy Price Increases on Economic Activity", *Journal of Money Credit and Banking,* vol. 32, pp. 400-416.

[19] Giacomini, Raffaella and Halbert White (2006), "Tests of conditional predictive ability", *Econometrica* vol. 74, pp. 1545-1578.

[20] Gorton, Gary and K. Geert Rouwenhorst (2006), "Facts and Fantasies about Commodity Futures", *Financial Analysts Journal,* vol. 62, pp. 47- 68.

[21] Gorton, Gary, Fumio Hayashi and K. Geert Rouwenhorst (2007), "The fundamentals of commodity futures returns", *NBER working paper No. 13249.*

[22] Hodrick, Robert (1992), "Dividend Yields And Expected Stock Returns: Alternative Procedures For Influence And Measurement", *Review of Financial Studies*, vol. 5, pp. 357-386.

[23] Horowitz, Joel L., (2004), "The bootstrap", in Heckman, J.J., Leamer, E. (Eds.), *Handbook of Econometrics*, vol. 5. Elsevier, New York, pp. 3160–3228.

[24] Mincer, Jacob and Victor Zarnowitz (1969), "The Evaluation of Economic Forecasts" in J. Mincer (ed.), *Economic Forecasts and Expectations*, New York: NBER.

[25] Moosa, Imad A., and Nabeel E. Al.Loughani (1994) "Unbiasedness and Time Varying Risk Premia in the Crude Oil Futures Market". *Energy Economics*, vol. 16, pp. 99-105.

[26] Piazzesi, Monika, and Eric Swanson (2008) “Future Prices as Risk-Adjusted Forecasts of Monetary Policy”, *Journal of Monetary Economics*, vol. 55, pp. 677-691.

[27] Pindyck, Robert (2001) “The dynamics of commodity spot and futures markets: a primer”, *The Energy Journal*, vol. 22, pp. 1-29.

[28] Roache, Shaun K. (2008) “Commodities and the Market Price of Risk”, *IMF Working paper No. 08/221.*

[29] Sanders, Dwight R., Keith Boris and Mark Manfredo (2004) “Hedgers, funds, and small speculators in the energy futures markets: an analysis of the CFTCs Commitments of Traders reports”, *Energy Economics,* vol. 26, pp. 425-445

[30] Svensson, Lars E. O. (2005) “Oil prices and ECB Monetary Policy”, *manuscript*, www.princeton.edu/∼svensson/papers/ep501.pdf

European Central Bank Working Paper Series

For a complete list of Working Papers published by the ECB, please visit the ECB's website (http://www.ecb.europa.eu).

944 The New Area-Wide Model of the euro area: a micro-founded open-economy model for forecasting and policy analysis" by K. Christoffel, G. Coenen and A. Warne, October 2008.

945 "Wage and price dynamics in Portugal" by C. Robalo Marques, October 2008.

946 "Macroeconomic adjustment to monetary union" by G. Fagan and V. Gaspar, October 2008.

947 "Foreign-currency bonds: currency choice and the role of uncovered and covered interest parity" by M. M. Habib and M. Joy, October 2008.

948 "Clustering techniques applied to outlier detection of financial market series using a moving window filtering algorithm" by J. M. Puigvert Gutiérrez and J. Fortiana Gregori, October 2008.

949 "Short-term forecasts of euro area GDP growth" by E. Angelini, G. Camba-Méndez, D. Giannone, L. Reichlin and G. Rünstler, October 2008.

950 "Is forecasting with large models informative? Assessing the role of judgement in macroeconomic forecasts" by R. Mestre and P. McAdam, October 2008.

951 "Exchange rate pass-through in the global economy: the role of emerging market economies" by M. Bussière and T. Peltonen, October 2008.

952 "How successful is the G7 in managing exchange rates?" by M. Fratzscher, October 2008.

953 "Estimating and forecasting the euro area monthly national accounts from a dynamic factor model" by E. Angelini, M. Bańbura and G. Rünstler, October 2008.

954 "Fiscal policy responsiveness, persistence and discretion" by A. Afonso, L. Agnello and D. Furceri, October 2008.

955 "Monetary policy and stock market boom-bust cycles" by L. Christiano, C. Ilut, R. Motto and M. Rostagno, October 2008.

956 "The political economy under monetary union: has the euro made a difference?" by M. Fratzscher and L. Stracca, November 2008.

957 "Modeling autoregressive conditional skewness and kurtosis with multi-quantile CAViaR" by H. White, T.-H. Kim, and S. Manganelli, November 2008.

958 "Oil exporters: in search of an external anchor" by M. M. Habib and J. Stráský, November 2008.

959 "What drives U.S. current account fluctuations?" by A. Barnett and R. Straub, November 2008.

960 "On implications of micro price data for macro models" by B. Maćkowiak and F. Smets, November 2008.

961 "Budgetary and external imbalances relationship: a panel data diagnostic" by A. Afonso and C. Rault, November 2008.

962 "Optimal monetary policy and the transmission of oil-supply shocks to the euro area under rational expectations" by S. Adjemian and M. Darracq Pariès, November 2008.

963 "Public and private sector wages: co-movement and causality" by A. Lamo, J. J. Pérez and L. Schuknecht, November 2008.

964 "Do firms provide wage insurance against shocks? Evidence from Hungary" by G. Kátay, November 2008.

965 "IMF lending and geopolitics" by J. Reynaud and J. Vauday, November 2008.

966 "Large Bayesian VARs" by M. Bańbura, D. Giannone and L. Reichlin, November 2008.

967 "Central bank misperceptions and the role of money in interest rate rules" by V. Wieland and G. W. Beck, November 2008.

968 "A value at risk analysis of credit default swaps" by B. Raunig and M. Scheicher, November 2008.

969 "Comparing and evaluating Bayesian predictive distributions of asset returns" by J. Geweke and G. Amisano, November 2008.

970 "Responses to monetary policy shocks in the east and west of Europe" by M. Jarociński, November 2008.

971 "Interactions between private and public sector wages" by A. Afonso and P. Gomes, November 2008.

972 "Monetary policy and housing prices in an estimated DSGE for the US and the euro area" by M. Darracq Pariès and A. Notarpietro, November 2008.

973 "Do China and oil exporters influence major currency configurations?" by M. Fratzscher and A. Mehl, December 2008.

974 "Institutional features of wage bargaining in 23 European countries, the US and Japan" by P. Du Caju, E. Gautier, D. Momferatou and M. Ward-Warmedinger, December 2008.

975 "Early estimates of euro area real GDP growth: a bottom up approach from the production side" by E. Hahn and F. Skudelny, December 2008.

976 "The term structure of interest rates across frequencies" by K. Assenmacher-Wesche and S. Gerlach, December 2008.

977 "Predictions of short-term rates and the expectations hypothesis of the term structure of interest rates" by M. Guidolin and D. L. Thornton, December 2008.

978 "Measuring monetary policy expectations from financial market instruments" by M. Joyce, J. Relleen and S. Sorensen, December 2008.

979 "Futures contract rates as monetary policy forecasts" by G. Ferrero and A. Nobili, December 2008.

980 "Extracting market expectations from yield curves augmented by money market interest rates: the case of Japan" by T. Nagano and N. Baba, December 2008.

981 "Why the effective price for money exceeds the policy rate in the ECB tenders?" by T. Välimäki, December 2008.

982 "Modelling short-term interest rate spreads in the euro money market" by N. Cassola and C. Morana, December 2008.

983 "What explains the spread between the euro overnight rate and the ECB's policy rate?" by T. Linzert and S. Schmidt, December 2008.

984 "The daily and policy-relevant liquidity effects" by D. L. Thornton, December 2008.

985 "Portuguese banks in the euro area market for daily funds" by L. Farinha and V. Gaspar, December 2008.

986 "The topology of the federal funds market" by M. L. Bech and E. Atalay, December 2008.

987 "Probability of informed trading on the euro overnight market rate: an update" by J. Idier and S. Nardelli, December 2008.

988 "The interday and intraday patterns of the overnight market: evidence from an electronic platform" by R. Beaupain and A. Durré, December 2008.

989 "Modelling loans to non-financial corporations in the euro area" by C. Kok Sørensen, D. Marqués Ibáñez and C. Rossi, January 2009

990 "Fiscal policy, housing and stock prices" by A. Afonso and R. M. Sousa, January 2009.

991 "The macroeconomic effects of fiscal policy" by A. Afonso and R. M. Sousa, January 2009.

992 "FDI and productivity convergence in central and eastern Europe: an industry-level investigation" by M. Bijsterbosch and M. Kolasa, January 2009.

993 "Has emerging Asia decoupled? An analysis of production and trade linkages using the Asian international input-output table" by G. Pula and T. A. Peltonen, January 2009.

994 "Fiscal sustainability and policy implications for the euro area" by F. Balassone, J. Cunha, G. Langenus, B. Manzke, J. Pavot, D. Prammer and P. Tommasino, January 2009.

995 "Current account benchmarks for central and eastern Europe: a desperate search?" by M. Ca' Zorzi, A. Chudik and A. Dieppe, January 2009.

996 "What drives euro area break-even inflation rates?" by M. Ciccarelli and J. A. García, January 2009.

997 "Financing obstacles and growth: an analysis for euro area non-financial corporations" by C. Coluzzi, A. Ferrando and C. Martinez-Carrascal, January 2009.

998 "Infinite-dimensional VARs and factor models" by A. Chudik and M. H. Pesaran, January 2009.

999 "Risk-adjusted forecasts of oil prices" by P. Pagano and M. Pisani, January 2009.