

FORWARD DISCOUNT BIAS: IS IT A RISK PREMIUM ON CRUDE OIL?

August 10, 2007

SHOKO NAKANO*

Graduate School of Economics, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo 186-8601, Japan

Ministry of Economy, Trade, and Industry, 3-1 Kasumigaseki, Chiyoda-ku, Tokyo 100-8901, Japan

Abstract

This paper investigates the forward discount bias in Light Sweet Crude Oil futures, also known as West Texas Intermediate (WTI) crude oil futures, listed on the New York Mercantile Exchange. We confirm that the current forward-spot differential is a biased forecast of future changes in the spot price. Using survey forecast data, we determined that the statistically significant causes of the forward discount bias were (i) the correlation between forecast errors and the current forward-spot differential, and (ii) the average level of forecast errors. Additionally, we found that, consistently after the first half of 2003, there existed a substantial average amount of risk premium—the current forward-spot differential was nine percentage points higher on average than what the average forecaster had expected.

JEL Classification: G13; G14; Q40

Keywords: Crude oil; Forward discount bias; Survey Forecast; Risk premium

* E-mail: nakano-shoko@meti.go.jp . Views expressed in this paper are those of authors and do not necessarily reflect those of Ministry of Economy, Trade, and Industry.

1. INTRODUCTION

Is the forward discount a biased predictor of the future change in the spot price?

This is one of the oldest questions in Financial Economics, especially in the International Finance literature. Most studies found it to be a biased predictor, and the result is referred to as the forward discount puzzle.

This paper conducts the unbiasedness test for crude oil futures, specifically for the Light Sweet Crude Oil futures listed on the New York Mercantile Exchange, known also as West Texas Intermediate (WTI) crude oil futures. This paper uses survey data on spot price expectations of WTI crude oil, since survey data allow us to decompose forward discount bias into components attributable to the risk premium and expectational errors. Without spot price expectations, one cannot determine whether the bias is evidence of a risk premium or of a violation of rational expectation. The literature on foreign exchange rates tends to disagree about the answer. Some studies, for example Bilson (1981), assume that investors are risk neutral, so that the systematic component of price changes in excess of the forward discount is interpreted as evidence of a failure of rational expectations. Others, however, such as Hsieh (1984), attribute the same systematic component to a time-varying risk premium that separates the forward discount from expected depreciation. In order to show the cause of the forward discount bias as clearly as possible, this paper follows Froot and Frankel (1989)

and uses survey data, specifically the one published by Consensus Economics Inc., which surveyed an estimate of the spot price of WTI for three months at a regular interval of one month since November 1995.

Additionally, we allow the average level of risk premium to change discretely and examine whether the change is caused by hedging pressure, a phenomenon first addressed by Keynes (1930). Froot and Frankel (1989) have conducted the unbiasedness test for foreign exchange rates using survey data and found a substantial average level of the risk premium. Their Figures I–IV also suggest that the average level of the risk premium changed discretely. To capture such discrete changes, this paper employs an estimation method developed by Bai and Perron (1998).

Our findings, stated below, are (1) that the WTI crude oil futures prices are a biased predictor of the future spot price; (2) that the bias is evidence of (i) the correlation between the forward discount and forecast errors, and (ii) the average level of forecast errors; (3) that a substantial risk premium emerged, that is, the forward discount was nine percentage points higher on average than what an average forecaster had expected after the early 2003; and (4) the substantial amount of the risk premium is not caused by hedging pressure.

The rest of this paper is organized as follows: we conduct the unbiasedness test in Section 2; give an overview of the data, including the survey data in Section

3; investigate whether the risk premium explains the forward discount, and also whether the average level of the risk premium is caused by hedging pressure in Section 4; examine whether forecast errors explain the forward discount in Section 5; and state our conclusions in Section 6.

2. UNBIASEDNESS TEST

In this section, we examine forward market unbiasedness for WTI crude oil futures prices using this regression:

$$S_{t+k} - S_t = \alpha + \beta (F_{t,t+k} - S_t) + \eta_{t+k}^k. \quad (1)$$

$S_{t+k} - S_t$ is the percentage increment of the spot price (the change in the log of the spot price of WTI Crude Oil) over k periods, $F_{t,t+k} - S_t$ is the current k -period forward discount (the log of the futures price minus the log of the spot price), and η_{t+k}^k is a random error (market participants' expectational errors). The null hypotheses to test whether the realized spot rate is equal to the forward rate plus a random error term are that $\beta=1$ and $\alpha=0$. Futures prices are computed by rolling over the prices of four-month-to-maturity contracts; therefore, the k in (1) indicates four months. As a proxy for the spot price, the settlement price at the day

of maturity is used. Therefore, the t in (1) indicates the maturity date.¹ The sample period is February 1996 through November 2006.

Eq. (1) is estimated using OLS. The results are presented in Table 1. The estimates α and β are 0.04 and 0.12, respectively. The β is significantly less than one, suggesting that the forward discount of WTI crude oil futures is a biased predictor of the future change in the spot price.

[Table 1 Here]

The result that $\beta \neq 1$ indicates either (a) that the market participants' expectational errors are correlated with the forward discount, or (b) that the risk premium is correlated with the forward discount (that is, the time-varying risk premium exists). In order to identify these factors, Froot and Frankel (1989) used the estimate of market expectation, \hat{S}_{t+k}^e , and decomposed the coefficient β into 1 (the null hypothesis) minus a term arising from any failure of rational expectations, b_{re} , minus another term arising from risk premium, b_{rp} :

$$\beta = 1 - b_{re} - b_{rp} \quad (2)$$

1 In order to avoid the information overlap problem caused by serial correlation in the errors, we use Andrew (1991)'s Heterogeneous Autocorrelation Consistent, HAC, estimators.

$$b_{re} = \frac{-\text{cov}(\eta_{t+k}^k, F_{t,t+k} - S_t)}{\text{var}(F_{t,t+k} - S_t)};$$

$$b_{rp} = \frac{\text{var}(rp_t^k) + \text{cov}(\hat{S}_{t+k}^e - S_t, rp_t^k)}{\text{var}(F_{t,t+k} - S_t)};$$

$$rp_t^k = F_{t,t+k} - \hat{S}_{t+k}^e.$$

For example, if there are no systematic prediction errors in the sample, then $b_{re} = 0$. On the other hand, if the risk premium is uncorrelated with the forward discount, $b_{rp} = 0$. Thus, (2) indicates that the estimate of market expectation makes it possible to clearly decompose the cause that $\beta \neq 1$ into forecast errors and time-varying risk premiums.

3. SURVEY DATA

This paper uses survey data collected by Consensus Economics Inc. Each month, Consensus Economics surveys several hundred economic and financial forecasters. Over fifty of them completed this survey question, giving estimates for the spot price in both three and twelve months of WTI (USD per barrel).² This paper uses the three-month survey forecast data.³

2 The forecast data of WTI crude oil is available at regular one-month intervals beginning November 1995.

3 For example, in the survey conducted on June 13, 2005, a forecaster provided his/her estimate for the spot price of WTI crude oil for the end of September 2005 as the estimate in three months. Trading of the WTI crude oil futures terminates at the close of business on the third business day prior to the 25th calendar day of the month preceding the delivery month. Thus, the WTI crude oil

Figure 1 shows the mean of the three-month survey forecasts (thin line), both the highest and lowest forecasts (dotted lines), the settlement prices of the fourth contract of WTI crude oil futures (thin line with diamond), and the realized spot prices of WTI crude oil (bold line). The date on the horizontal axis is appended based on the date of forecast. For example, the values shown on May 2003 are the mean of the survey forecasts made on May 12, 2003 (USD 24.9 per barrel), the futures price on the same date (USD 26.5 per barrel), and the realized spot price at the day of maturity in September 2003 (USD 26.9 per barrel). The reason why the realized spot price is shown on May 2003 is that the spot price was forecasted on May 12, 2003. Hence, the vertical difference between the realized spot price and the three-month survey forecast shows the survey-measured forecast errors of the forecast made on the appended date. Similarly, the difference between the futures price and the survey forecast shows the survey-measured risk premium paid on the appended date.

[Figure 1 Here]

From 1995 through 1998, crude oil prices were relatively stable. Futures

delivered during October 2005 was terminated on September 20, 2005, and was the fourth contract on June 13, 2005. Therefore, we use the three-month survey forecast, which is consistent with the fourth contract in the WTI crude oil futures.

prices and survey forecasts were very close, suggesting that the average level of the risk premium was almost zero. The realized spot prices were oscillating around them within a very narrow range. Amid such circumstances, the Asian currency crisis in July 1997 brought a decline in oil prices. Alarmed by the decline, beginning in March 1998, both OPEC and non-OPEC countries started withholding output. This cooperative withholding caused oil prices to move slightly upward in the first quarter of 1999. After 2003, oil prices increased unprecedentedly because of uncertainty in Iraq, increasing demand in China, and other factors. During this period, the realized spot and futures prices were constantly higher than the mean of the survey forecasts. This suggests that there existed both the average level of forecast errors and the average level of risk premium.

4. DOES THE RISK PREMIUM EXPLAIN ANY OF THE FORWARD DISCOUNT BIAS?

This section examines whether the forward discount bias can be explained by the risk premium using the survey forecast data. We define S_{t+k}^e as the unobservable market-expected spot price and ε_t^e as the random measurement error in the surveys. We also assume that the mean of all survey responses is an estimated true

investor expectation, that is, $\hat{S}_{t+k}^e = S_{t+k}^e + \varepsilon_t^e$.⁴ Under this assumption, we regress the measure of expected change over k periods against the forward discount:

$$\hat{S}_{t+k}^e - S_t = a_j + b(F_{t,t+k} - S_t) + \varepsilon_t^e \quad \text{for } j = 1, \dots, m+1. \quad (3)$$

The probability limit of the coefficient b in (3) is $\frac{\text{Cov}(\hat{S}_{t+k}^e - S_t, F_{t,t+k} - S_t)}{\text{Var}(F_{t,t+k} - S_t)}$, and the

probability limit of the coefficient β in (1) is

$\frac{\text{Cov}(\hat{S}_{t+k}^e - S_t, F_{t,t+k} - S_t)}{\text{Var}(F_{t,t+k} - S_t)} + \frac{\text{Cov}(\eta_{t+k}, F_{t,t+k} - S_t)}{\text{Var}(F_{t,t+k} - S_t)}$. Using (2), we can express coefficient b as

$$b = \beta - \frac{\text{Cov}(\eta_{t+k}, F_{t,t+k} - S_t)}{\text{Var}(F_{t,t+k} - S_t)} = 1 - b_{rp}. \quad (4)$$

Hence, the null hypothesis that the risk premium is uncorrelated with the forward discount is that $b = 1$, that is, $b_{rp} = 0$. This means that there is no time-varying risk premium. The null hypothesis that there is no average level of the risk premium is that $a = 0$. The null hypothesis that the risk premium is identically zero ($\hat{S}_{t+k}^e = F_{t,t+k}$) is that both $a = 0$ and $b = 1$. To examine discrete changes in the average level of risk premium, we take into account the multiple structural breaks in the constant term ($j = 1, \dots, m+1$), and test the possible structural

4 This is analogous to the rational expectation approach, which uses ex-post price changes rather than survey data and assumes that the error in measuring true expected price change is random.

breaks using the method of Bai and Perron (1998).

The results of (3) are presented in the first row in Table 2. We first examine the results for the time-varying risk premium. The estimate b is 1.17, and is not statistically different from one. This indicates that the reason why the hypothesis that $\beta = 1$ in (1) was rejected is not because of the time-varying risk premium. Next, we examine the average level of the risk premium. From the constant term in (3), one structural break is detected in May 2003. Before May 2003, the constant term was -0.02 (a_1); and after June 2003, it became -0.09 (a_2). While the constant term a_1 is not significantly different from zero, the constant term a_2 is different with a 1% significance level. It indicates that after the break date the average level of the risk premium existed. Specifically, the forward discount was nine percentage points higher on average than what an average forecaster expected.

[Table 2 Here]

To determine who paid the average level of risk premium, we add a variable that can capture hedging pressure, q , in (3). (See Appendix for the concept of hedging pressure, the definition of hedging pressure in this paper, and the regression equation with the hedging pressure variable.) The results after

controlling the effects from hedging pressure are shown in the second row in Table 2. A structural break is detected in January 2003. Before January 2003, the constant term was -0.02 (a_1); after the break date it became -0.09 (a_2). These estimates are significant and almost the same as the estimates in (3). Therefore, we can conclude that the average level of the risk premium is not caused by the hedging pressure.⁵

5. DO EXPECTATIONAL ERRORS EXPLAIN ANY OF THE FORWARD DISCOUNT BIAS?

In this section, we test whether there are systematic expectational errors that can explain the finding of bias in the forward discount. Among factors that induced the rejection in Table 1, the portion arising from the correlation between the forecast errors and the forward discount (the coefficient b_{re} in (2)) is obtained by regressing the expectational prediction error, $(\hat{S}_{t+k}^e - S_{t+k})$, on the forward discount:⁶

5 Considering (A2) in Appendix as 2SLS, at the first stage, we regress the variable $F_{t,t+k} - S_t$ on the variable q_t . The estimated coefficient for q_t at the first stage is -0.25 (its standard deviation, 0.08). This significant negative coefficient means that, when there is selling (buying) hedging pressure, the WTI crude oil futures tend to significantly decline (increase). Combining this result with the results from (A2), we can consider that the hedging pressure is significant but is not the main cause of the average of the risk premium.

6 The error term in (5) is the measurement error in the surveys less the unexpected change in the spot rate, that is, $v_{t+k}^k = \varepsilon_t^k - \eta_{t+k}^k$.

$$(\hat{S}_{t+k}^e - S_{t+k}) = d_j + b_{re} (F_{t,t+k} - S_t) + v_{t+k}^k \quad \text{for } j = 1, \dots, m+1. \quad (5)$$

The null hypothesis that the forecast error is uncorrelated with the forward discount is that $b_{re} = 0$. A finding that $b_{re} > 0$, for example, means that an investor could have made excess profit by betting lower in absolute magnitude than the forward discount. The null hypothesis that there is no average level of the forecast error is that $d = 0$. Additionally, the null hypothesis that there is no systematic forecast error is that both $d = 0$ and $b_{re} = 0$.

The estimation results are presented in Table 3. The estimate of b_{re} is 0.89, which is significantly different from zero. This suggests that the forward discount bias is caused by the correlation between the forecast error and the forward discount. A structural break is not detected from the constant term. The estimate of the constant term is -0.09 throughout all the sample period, which is significantly different from zero. These estimation results suggest that systematic forecast errors arose due to (i) the correlation between the forecast error and the forward discount ($b_{re} > 0$), and (ii) the average level of underestimation ($d < 0$).

[Table 3 Here]

We cannot reject the hypothesis that there is no systematic prediction

error. However, this finding does not necessarily mean that investors are irrational, as a systematic prediction error can emerge while rational investors incorporate uncertainty about shifts in the distribution of economic shocks into their forecasts, that is, while learning. Since 2003, suspect events in such shifts have occurred consecutively. These are, for example, the uncertainty in Iraq and the more-than-expected growth of the Chinese economy. In the light of learning, it is reasonable to assume that these consecutively occurring suspect events are sufficient to induce the systematic prediction errors.

6. CONCLUSION

This paper has investigated the forward discount bias in the Light Sweet Crude Oil futures, also known as the WTI crude oil futures, listed in New York Mercantile Exchange. We confirmed that (1) the forward discount was a biased forecast of the future change in the spot price. We then found that (2) the statistically significant causes of the forward discount bias were (i) the correlation between forecast errors and forward discount, and (ii) the average level of forecast errors. We could not find the time-varying risk premium with statistical significance. However, we found that (3) there existed a substantial average amount of risk premium—after 2003, the forward discount was consistently nine percentage points higher on average than what the average forecaster expected. (4) The substantial amount of the risk

premium is not caused by hedging pressure.

Appendix. WHO PAID THE AVERAGE LEVEL OF RISK PREMIUM?

Here, we try to determine who paid the average level of risk premium by adding, in (3), a variable that can capture hedging pressure, that is, hedge demand from hedgers in the futures market. Hedging pressure was first expressed by Keynes (1930) to explain the futures price bias and was developed by Hirshleifer (1990), De Roon et al. (2000).

In many papers, the hedgers' net short position is used as a proxy for hedging pressure.⁷ Following this approach, we compose a hedging pressure variable, q , using the large trader's positions published by the Commodity Futures Trading Commission (CFTC). Specifically, it is composed of the short and long positions at the WTI crude oil market in NYMEX:

$$q = \frac{\text{number of short hedge positions} - \text{number of long hedge positions}}{\text{total number of hedge positions}}. \quad (\text{A1})$$

Hedging pressure at date t affects a futures price quoted at date t . Therefore, the specification including hedging pressure is as follows.

⁷ For example, Chang, et. al. (1985) and Sanders, et. al. (2004) have used the hedgers' net short position as a proxy for hedging pressure. Additionally, there are rare papers that have used the disaggregated (non-public) version of the Commitments of Traders (COT) data of the Commodity Futures Trading Commission (CFTC), e.g., Haigh, et. al. (2007).

$$\hat{S}_{t+k}^e - S_t = a_j + b(F_{t,t+k} - S_t) + cq_t + \varepsilon_t^e \text{ for } j=1, \dots, m+1. \quad (\text{A2})$$

In this specification, the effects from hedging pressure are removed from the coefficients a_j and b . Therefore, if we find structural breaks in (A2) similar to those that we found in (3), we may conclude that the hedging pressure is not the cause of the average level of the risk premium.

Reference

- Bai, J., and Perron, P, 1998. Estimating and testing linear models with multiple structural changes. *Econometrica* 66, 47-78.
- Bilson, J., 1981. The speculative efficiency hypothesis. *Journal of Business* LIV, 435-51.
- Chang, E., 1985. Returns to speculators and the theory of normal backwardation. *The Journal of Finance* Vol. 40 No. 1, 193-208.
- Commodity Futures Trading Commission, 1962-(ongoing). *The Commitments of Traders*. Washington, DC.
- Consensus Economics Inc., 1984-(ongoing). *Consensus Forecasts*. London.
- De Roon, F., Nijman, T., and Veld, C., 2000. Hedging pressure effects in futures markets. *The Journal of Finance* Vol. LV No. 3, 1437-1456.
- Froot, K., and Frankel, J., 1989. Forward discount bias: Is it an exchange risk premium? *The Quarterly Journal of Economics* Vol. 104 No. 1, 139-161.
- Haigh, M., Harris, J., Overdahl, J., and Robe, M., 2007. *Market Growth, Trader Participation and Pricing in Energy Futures Markets*. Manuscript, Commodity Futures Trading Commission, Washington, D.C.
- Hirshleifer, D., 1990. Hedging pressure and futures price movements in a general equilibrium model. *Econometrica* Vol. 58 No. 2, 411-428.
- Hsieh, D., 1984. Tests of rational expectations and no risk premium in forward exchange markets. *Journal of International Economics* XVII, 173-84.
- Keynes, J., 1930. *A Treatise on Money*, Vol. 2. London, MacMillan.
- Sanders, D., Boris, W., and Manfredo, M., 2004. Hedging, funds, and small speculators in the energy futures markets: An analysis of the CFTC's Commitments of Traders reports. *Energy Economics* 26, 425-445.

Table 1. Unbiasedness Test for WTI Crude Oil Futures Prices

$$(1) \quad S_{t+k} - S_t = \alpha + \beta (F_{t,t+k} - S_t) + \eta_{t+k}^k$$

	α	β
(1)	0.04	0.12
	<i>0.05</i>	<i>0.06 +++</i>

Notes:

Italic figure is Andrew (1991)'s Heterogeneous Autocorrelation Consistent, HAC, standard deviation.

*, **, and *** represent significance at the 10, 5, and 1 percent levels, respectively, with H0: $x = 0$.

+, ++, and +++ represent significance at the 10, 5, and 1 percent levels, respectively, with H0: $x = 1$.

Table 2. Whether the Risk Premium Explains any of the Forward Discounts Bias

$$(3) \quad \hat{S}_{t+k}^e - S_t = a_j + b(F_{t,t+k} - S_t) + \varepsilon_t^e \quad j = 1, \dots, m + 1$$

$$(A2) \quad \hat{S}_{t+k}^e - S_t = a_j + b(F_{t,t+k} - S_t) + c\hat{q}_t + \varepsilon_t^e \quad j = 1, \dots, m + 1$$

	a_1	a_2	b	c	Break Date	Adj. R ²
(3)	-0.02 <i>0.04</i>	-0.09 <i>0.01 ***</i>	1.17 <i>2.17</i>		May 2003	0.63
(A2)	-0.02 <i>0.06</i>	-0.09 <i>0.03 ***</i>	1.04 <i>3.56</i>	-0.24 <i>2.59</i>	January 2003	0.65

Notes:

Italic figure is Andrew (1991)'s Heterogeneous Autocorrelation Consistent, HAC, standard deviation.

*, **, and *** represent significance at the 10, 5, and 1 percent levels, respectively, with H0: x = 0.

+, ++, and +++ represent significance at the 10, 5, and 1 percent levels, respectively, with H0: x = 1.

In Bai and Perron (1998)'s estimation, two structural breaks are allowed as maximum, and .25 is set as the value of the trimming.

Table 3. Whether Expectationl Errors Explains any of the Forward Discount Bias

$$(5) \quad (\hat{S}_{t+k}^e - S_{t+k}) = d_j + b_{re} (F_{t, t+k} - S_t) + v_{t+k}^k \quad j = 1, \dots, m + 1$$

	d	b_{re}	Break Date	Adj. R ²
(5)	-0.09 <i>0.01 ***</i>	0.89 <i>0.03 ***</i>	none	0.05

Notes:


Italic figure is Andrew (1991)'s Heterogeneous Autocorrelation Consistent, HAC, standard deviation.

*, **, and *** represent significance at the 10, 5, and 1 percent levels, respectively, with H0: x = 0.

+, ++, and +++ represent significance at the 10, 5, and 1 percent levels, respectively, with H0: x = 1.

In Bai and Perron (1998)'s estimation, two structural breaks are allowed as maximum, and .25 is set as the value of the trimming.

Figure 1. The WTI Crude Oil Prices and its Forecasts



The date on the horizontal axis is appended based on the date of forecast. For example, the values shown on May 2003 are the mean of the survey forecasts made on May 12, 2003 (USD 24.9 per barrel), the futures price on the same date (USD 26.5 per barrel), and the realized spot price at the day of maturity in September 2003 (USD 26.9 per barrel). The reason why the realized spot price is shown on May 2003 is that the spot price was forecasted on May 12, 2003. Hence, the vertical difference between the realized spot price and the three-month survey forecast shows the survey-measured forecast errors of the forecast made on the appended