Rational Expectations and Risk Premia in Forward Markets: Primary Metals at the London Metals Exchange

David A. Hsieh, Nalin Kulatilaka

Rational Expectations and Risk Premia in Forward Markets: Primary Metals at the London Metals Exchange

DAVID A. HSIEH and NALIN KULATILAKA*

ABSTRACT

This paper tests whether forward prices equal the traders' expectations of the future spot prices at maturity, under two different models of expectations formation: full information rational expectations and incomplete information mechanical forecasting rule. The tests are performed, over the period January 1970 through September 1980, on the forward markets for the primary metals—copper, tin, lead, and zinc-traded in the London Metals Exchange. We find evidence consistent with the existence of time varying risk premia.

THE WAY IN WHICH traders form expectations about future price is of great interest to economists as well as to market participants, and forward prices have often been used as indicators of these otherwise unobservable expectations. Forward prices, however, include not only the expectation of future spot prices but also a component reflecting the riskiness of the contract. Therefore, the risk premium can be defined as the difference between the forward price and the expected future spot at the maturity date of the forward contract.[1]

There is a longstanding debate on the nature of this risk premium. The Hicks-Keynes view[2] is that in futures markets speculators provide insurance against the risk of fluctuations in the price of the spot commodity, and hedgers pay a premium for this risk transfer in the way of a forward price that is higher than the expected spot price. In other words the risk premium should be positive.[3] Hardy [11] has suggested, however, that a future market is merely a casino in which speculators can participate in a legalized form of gambling. For this privilege they have to pay a price in the form of a negative expected return, thus resulting in a negative risk premium.[4]

Recent empirical studies have not detected the presence of significant risk premia in forward markets. Dusak [7], Grauer [8], and Bodi and Rosinsky [4],

[1] Futures contracts are "marked to marketing" (i.e. gains and losses are settled every day) while forward contracts are not. Black [1] gives a detailed description of the differences between these contracts.

[2] Details of this debate can be found in a series of papers by Telser [13] and Cootner [5, 6].

[3] This is called "normal backwardation."

[4] This is called "contango."

used the Sharpe-Lintner capital asset pricing model and found that virtually all commodity futures prices examined had no systematic risk. Breeden [3] using the intertemporal asset pricing model in Breeden [2] found that consumption betas for most of the futures contracts examined were not significantly different from zero. Hence, futures prices do not appear to contain risk premia.

In this paper we use the three-month forward prices for the primary metals traded in the London Metals Exchange (LME), over the period January 1970 to September 1980, to test the null hypothesis that forward prices equal the traders' expectations of the spot prices at maturity. Since expectations are not observable, we test the null hypothesis jointly with two different models of expectations formation. In the first model, traders have full-information rational expectations. The testable implications are that the forecast errors of forward prices on future spot prices have zero mean and that these errors are uncorrelated with any variable included in the information set available to the market at the time the expectations are formed. In the second model, traders form expectations via a mechanical predictor such as the Kalman filter. If forward prices do not contain risk premia, then they should have a lower mean squared error in forecasting future spot prices than the mechanical predictor.

Furthermore, we can proxy the risk premia in forward prices as the difference between the forward prices and the corresponding Kalman filter forecasts, and analyze the time series properties of these derived risk premia.

This paper proceeds as follows: the next section outlines the tests assuming that traders have full information rational expectations. Section II briefly describes the data, and Section III gives the results. Section IV discusses the model where traders use a Kalman filter forecasting rule. Comparison of the forecasting performance of the Kalman filter and the forward prices are reported in Section V. Some concluding remarks are offered in the final section.

I. Full Information Rational Expectations

Let $S(t)$ be the spot price of a commodity at date t and $F(t, n)$ be the forward price contracted at date t for delivery and payment at date $t + n$. Under the full information rational expectations case, $I(t)$, the information set available to the market participants at time t is assumed to contain all past and current values of forward and spot prices, and the exact stochastic process by which prices are determined. We define the risk premium $r(t, n)$ as

$$r(t, n) = F(t, n) - E[S(t + n)| I(t)] \qquad (1)$$

Let $u(t, n)$, be the realized return of the forward contract[5]

$$u(t, n) = [F(t, n) - S(t + n)]/F(t, n) \qquad (2)$$

The null hypothesis of a zero risk premium implies that

$$E[u(t, n)| I(t)] = 0 \qquad (3)$$

---

[5] This expression can be thought of as the difference between holding yields from the spot contracts and the forward contracts.

Hence, the unconditional mean of $u(t, n)$ is zero, and that $u(t, n)$ is uncorrelated with any variable included in the information set $I(t)$. These properties can be empirically examined by performing the following tests.

T1: $u(t, n)$ has a zero mean

The sample mean of $u(t, n)$ can be calculated as the constant term of the regression:

$$u(t, n) = \alpha + \epsilon(t, n) \qquad (4)$$

Under the null hypothesis, when the length of the forward contract is n sample periods, we know that $\epsilon(t, n)$ is a moving average of order $n - 1$.[6] To obtain a consistent test of the hypothesis $\alpha = 0$, we compute a consistent estimate of variance of the sample mean using the method in Hansen and Hodrick [10]. A more efficient estimator can be found by purging the serial correlation entirely from $\epsilon(t, n)$ with spectral methods such as the "Hannan efficient" estimator (see Hannan [9]).

T2: $u(t, n)$ has a zero mean and is uncorrelated with $u(t - n - j, n)$ for $j \geq 0$

In this test, we run the regression

$$u(t, n) = \alpha + B(L)u(t - n, n) + \epsilon(t, n) \qquad (5)$$

where α is a constant term and $B(L) = \sum_{i>0} \beta_i L^i$ is a polynomial of the lag operator L. Under the null hypothesis, the coefficients α and β_i should be zero and residuals should be a moving average of order $n - 1$. Therefore, as in Hansen and Hodrick [10], Equation (5) can be estimated by OLS to get a consistent estimate of the covariance of the coefficients.

T2 is performed in single and multimarket frameworks. In the single market case, we look only for the serial correlation with the own error terms. In the multimarket case, we include errors in all four of the metals markets. The single market test makes the unrealistic assumption that the traders are ignorant about information from the other markets within the same commodity exchange. In the LME, traders deal in more than one commodity; thus, the multimarket test will be more appropriate.

T3: $u(t, n)$ has zero mean and is uncorrelated with $f(t - n - j, n)$ and $s(t - j)$ for $j > 0$, where $s(t) = [S(t) - S(t - 3)]/S(t - 3)$, and $f(t) = [F(t, n) - S(t)]/S(t)$

In this test we run the following regression

$$u(t, n) = \alpha + B(L)f(t - n - 1, n) + C(L)s(t - 1) + \epsilon(t, n) \qquad (6)$$

where α is a constant term, $B(L) = \sum_{i>0} \beta_i L^i$, and $C(L) = \sum_{i>0} \gamma_i L^i$. We test if all the right-hand side coefficients are zero.

---

[6] This is caused by the overlapping nature of the forward contracts. See Hansen and Hodrick [10] for a detailed description in the context of forward contracts in foreign exchange markets.

II. The Data

The three month forward contracts for copper, tin, zinc, and lead, traded in the London Metals Exchange (LME) were used in our empirical testing. We chose this data set because of the high trading volume in the LME for primary commodities (thereby reducing the "nontrading effect") and the convenient availability of the data. Furthermore, most of the previous empirical research into commodity markets have looked at agricultural commodities. The metals markets, on the other hand, have the added attraction that they do not suffer from the discontinuities of production patterns during the harvesting seasons.

We used the forward price on the first Friday of each month and the corresponding spot price on the delivery date as quoted in the *London Financial Times*. The market closing settlement prices were used in all cases. On days when the corresponding three month delivery date fell on a nontrading date, we used the prices from the first following trading day. The sample period was from January 1970 through September 1980, thus providing 129 observations. All contracts were denominated in pounds sterling per metric ton.

Several other caveats about the data need mentioning. We sample every month but the contracts are of three month duration; hence, the forecast errors will be correlated with information on realized spot prices of the three previous contracts. Furthermore, the unequal lengths of the contracts caused some overlapping even with a three month lag. Implications of these caveats are highlighted in later sections.

III. Empirical Results of the Full Information Tests

Table I gives the results from T1. Using OLS, we found only the tin forward price to have a systematic bias in predicting the future spot price. However, using Hannan efficient estimates, we found that the lead forward price also had a prediction bias.

Table II gives results of T2. In the single market tests, presented in Columns (1) and (2), we regressed the three period forecast errors on a constant term and

Table I
Mean Forecast Errors of Forward Prices: T1

	Copper	Tin	Zinc	Lead
OLS estimates				
mean	4.569	19.681	13.259	15.496
t-statistic[a]	0.395	2.312[b]	0.804	1.368
Hannan efficient estimates				
mean	2.781	17.010	12.733	16.061
t-statistic	0.334	2.545[b]	0.886	2.136[b]

Period: January 1971 to December 1980; 120 observations.

Equation: $u(t, 3) = f(t + 3) - s(t, 3) + \varepsilon(t, n)$

[a] OLS standard errors are corrected for serial correlation using the procedure in Hansen and Hodrick [10].

[b] Significant at the 5 percent level.

its own third lag, and rejected the hypothesis that $\alpha = \beta = 0$ at the five percent level for tin and zinc. For tin we were unable to reject the hypothesis that $\beta = 0$, while for zinc we were able to reject it. We accepted the null hypothesis for copper and lead. In the multimarket tests presented in Columns (3) and (4), we rejected the null hypothesis for copper, tin, and zinc, but not for lead. Table III gives the results of T3. All results from T2 were strengthened by T3.

Table II
Chi-Square Test for Serial Correlation
of Forecast Errors: T2

	Single market		Multimarket	
	(1)	(2)	(3)	(4)
Copper	0.84	0.66	12.40[a]	12.11[a]
Tin	8.82[a]	2.10	28.93[b]	16.58[b]
Zinc	9.12[a]	8.04[a]	20.09[b]	26.19[b]
Lead	1.86	0.02	7.62	5.27
Degrees of freedom	2	1	5	4

Period: January 1971 to December 1980; 120 observations.

[a] Significant at 5 percent level.
[b] Significant at 1 percent level.

(1) $u_{i,t} = \alpha + \beta u_{i,t-4}$ $HO: \alpha, \beta = 0$
(2) $u_{i,t} = \alpha + \beta u_{i,t-4}$ $HO: \beta = 0$
(3) $u_{i,t} = \alpha + \sum_j \beta_j u_{j,t-4}$ $HO: \alpha, \beta_j = 0$
(4) $u_{i,t} = \alpha + \sum_j B_j u_{j,t-4}$ $HO: \beta_j = 0$

Table III
Chi-Square Test of Correlation of Forecast Errors
with Available Information: T3

	Single market		Multimarket	
	(1)	(2)	(3)	(4)
Copper	1.32	1.12	20.54[a]	20.20[b]
Tin	8.92[a]	1.86	37.11[b]	23.18[b]
Zinc	15.04[b]	13.52[b]	48.59[b]	45.82[b]
Lead	2.91	0.63	9.92	7.08
Degree of freedom	4	3	13	12

Period: January 1971 to December 1980; 120 observations.

[a] Significant at 5 percent level.
[b] Significant at 1 percent level.

(1) $u_{i,t} = \alpha + \beta s_{i,t-3} + \gamma f_{i,t-3}$ $HO: \alpha, \beta, \gamma = 0$
(2) $u_{i,t} = \alpha + \beta s_{i,t-3} + \gamma f_{i,t-3}$ $HO: \beta, \gamma = 0$
(3) $u_{i,t} = \alpha + \sum_j \beta_j s_{j,t-3} + \sum_j \gamma_j f_{j,j-3}$ $HO: \alpha, \beta_j, \gamma_j = 0$
(4) $u_{i,t} = \alpha + \sum_j \beta_j s_{j,t-3} + \sum_j \gamma_j f_{j,j-3}$ $HO: \beta_j, \gamma_j = 0$

IV. Kalman Filter Forecasting Rule

It is possible that rational traders can make serially correlated forecast errors when they do not have full information. For example, suppose that the traders know the form of the true model but not the values of its parameters. They learn about these parameters over time. Then they can commit forecast errors which exhibit serial correlation (i.e. T2 is rejected) or which are correlated with the lags of forward and spot prices (i.e. T3 is rejected). However, there is no reason to conclude that they are acting in an "irrational" way.

To illustrate this point, consider the following simple example. Let $z(t)$ be a first-order Markov process:

$$z(t) = \delta z(t - 1) + \epsilon(t) \tag{7}$$

where δ is a real number between zero and one, and $\epsilon(t)$ has zero mean and is independently and identically distributed for all t.

When traders have full information, their forecast of $z(t)$ at time $t - 1$ is $\delta z(t - 1)$, and their one period forecast error is

$$u1(t, 1) = \epsilon(t) \tag{8}$$

which is uncorrelated with its own lags or with anything in the information set, e.g. past values of $z(t)$.

Alternatively, suppose traders do not know the value of δ. At time $t - 1$, their forecast of $z(t)$ is $\delta(t - 1) z(t - 1)$, where $\delta(t - 1)$ is the OLS estimate of δ at time $t - 1$. Then, their forecast error, $u2(t, 1)$, has two components:

$$u2(t, 1) = [\delta - \delta(t - 1)] z(t - 1) + \epsilon(t) \tag{9}$$

where the first component represents the error in estimating δ, and the second component the innovation at time t. In a finite sample, $u2(t, 1)$ may be correlated with its own lags and with lags of $z(t)$,[7] and so may fail the tests T2 and T3 more often than the forecast error $u1(t, 1)$.

To allow for learning, we replace the full information assumption by the assumption that market participants use a mechanical forecasting rule to form forecasts of the spot price three months ahead.[8] The simplest time series forecasting technique is a vector autoregression of the spot prices:

$$s(t + 3) = \alpha + B(L) s(t) + \epsilon(t, 3) \tag{10}$$

where $B(L) = \sum_{i>0} \beta_i L^i$. At every period, the traders update their estimates with the most current data (this is called a "rolling regression"). They then use the results to forecast the next period's forward price.

A more general method, of which the "rolling regression" is a special case, is the Kalman Filter. In the Kalman Filter technique, the coefficients of (10) are allowed to vary over time in a specific manner. Rewrite (10) as

$$s(t + 3) = Z(t) b(t) + \eta(t) \tag{11}$$

---

[7] This is because $b(t - 1)$ can be written as a distributed lag of the innovation $e(t)$. Hence $u2(t, 1)$ is correlated with its own past and with lags of $z(t)$.

[8] Hsieh [12] used a similar procedure in studying foreign exchange forward markets.

where $Z(t)$ is a matrix containing the right-hand side variables including the constant term in (10).

The coefficient vector $b(t)$ is allowed to vary over time in the following way:

$$b(t) = A(t)\, b(t - 1) + v(t) \qquad (12)$$

$A(t)$ is a transition matrix. The random errors $\eta(t)$ and $v(t)$ have zero mean and covariance matrices $R(t)$ and $Q(t)$, respectively. $A(t)$, $Q(t)$, and $R(t)$ are assumed to be known for every period t. The only unknown is the vector $b(0)$, which has a prior mean b_0 and a prior covariance matrix P_0. Furthermore, $b(0)$, $\eta(s)$, and $v(t)$ are assumed to be stochastically independent for all t and s.[9]

The choice of $A(t)$, $Q(t)$, and $R(t)$ are left to the traders' judgement. For example, the "random walk" specification is $A(t) = I$ and $Q(t) = Q$ where I is the identity matrix and Q is some positive definite matrix. The "rolling regression" specification is $A(t) = I$ and $Q(t) = 0$, where 0 is a matrix of zeros. These specifications are extremely restrictive since the coefficients $b(t)$ are not allowed to interact with each other except through the covariance Q. Instead, we allowed the coefficients within each equation to interact with one another, but allowed them no interaction across equations, as this drastically increases the dimensionality of the problem. We used the following specifications:

$$A(t) = c1 \qquad Q(t) = I \qquad R(t) = R \qquad (13)$$

where 1 is matrix of ones, c a normalization constant, and R a positive number. We used six lags of the log of the spot price, so that the normalization constant c is $1/7$.

In order to find a "reasonable" prior mean of $b(0)$, we divided the sample into two segments: January 1971 to December 1974, and January 1975 to December 1980 and used the coefficients of Equation (11), estimated over the first half of the sample, as prior means and the identity matrix as the prior covariance. The sum of squared residuals of this regression was used as R. The Kalman Filter was then used to compute the posterior distribution for the entire sample, updating the posterior mean at every point in time and generating a three-period-ahead forecast of the spot price.

V. Comparison of Forecasts

Table IV presents the mean squared error (MSE) for three forecasting methods over the second half of the sample period. (The "no change" forecast assumes that the next period's spot price will be identical to this period's spot price). In terms of MSE, the Kalman filter (KF) outperformed the forward price (FP) and the no change (NC) forecast in all four markets. Furthermore, we devised a statistical test to compare the MSE from various statistical forecasts. The general testing method is as follows; let $u1(t, n)$ and $u2(t, n)$ be the errors from two alternative forecasts made at time t of the spot price at $(t + n)$. The MSE of Method 1 is said to be significantly lower (higher) than that of Method 2 if the mean of the following series is significantly less (larger) than zero:[10]

---

[9] A more general version of the Kalman filter allows $\eta(t)$ and $v(t)$ to be correlated with each other.

[10] This is simply a test of the mean difference between the MSE.

Table IV

**Mean Squared Error of Alternative
Three-month-ahead Forecasts**

	Kalman Filter	Forward Price	No Change
Copper	3325	3349	3474
Tin	1543	1642	1898
Zinc	2206	2384	2223
Lead	5600	5840	6307

Period: January 1975 to December 1980; 72 observations.

Table V

Hannan Efficient Tests of MSE t-statistics

	(1) FP vs NC	(2) KF vs NC	(3) KF vs FP
Copper	−2.494	−3.078[b]	−2.098[a]
Tin	−2.288[b]	−2.502[b]	−1.999[a]
Zinc	2.626[b]	−0.149	−2.203[a]
Lead	−1.037	−2.707[b]	−1 040

Period: January 1975 to December 1980; 72 observations.
[a] Significant at 5 percent level.
[b] Significant at 1 percent level.
(1) $[s(t + 3) - f(t, 3)]^2 - [s(t + 3) - s(t)]^2 = u1(t)$ HO: $E[u1(t)] = 0$
(2) $[s(t + 3) - kf(t, 3)]^2 - [s(t + 3) - s(t)]^2 = u2(t)$ HO: $E[u2(t)] = 0$
(3) $[s(t + 3) - kf(t, 3)]^2 - [s(t + 3) - f(t, 3)]^2 = u3(t)$ HO: $E[u3(t)] = 0$

$$u(t, n) = [u1(t, n)]^2 - [u2(t, n)]^2 \qquad (14)$$

Table V provides a pairwise comparison of the MSE of the three different forecasts. Since the forecasts are overlapping, we remove the serial correlation using the Hannan efficient estimator to compute the t-statistic of $u(t, n)$ in (14). Column 1 shows that FP had a significantly lower MSE than the NC for copper, tin, and lead but the reverse was true of zinc. The KF forecast had MSE significantly lower than NC for copper, tin, and lead (in Column 2), and also significantly lower than FP for copper, tin, and zinc (in Column 3). This finding is consistent with the view that forward prices contain risk premia, and so are not optimal predictions of the future spot price.

VI. Concluding Remarks

Our results were not able to support the joint null hypothesis of rational expectations and no risk premium under the assumption of full information. We found that the forecast errors of the forward prices have nonzero means, have serial correlation, and also have some correlation with other available market information.

We then replaced the assumption of full information with one that allows market participants to use a mechanical forecasting rule, in this case a Kalman filter. Future spot prices were better forecasted with this rule than with the forward prices. These results are consistent with the hypothesis that the forward prices contain nonzero risk premia. In our future research, we plan to compute these risk premia as the difference between the forward prices and the Kalman filter forecasts, and to study the correlations of these risk premia with other market variables.

REFERENCES

1. F. Black. "The Pricing of Commodity Contracts." *Journal of Financial Economics* 3 (1976), 167–79.
2. D. T. Breeden. "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities." *Journal of Financial Economics* 7 (September 1979), 265–96.
3. ———. "Consumption Risks in Futures Markets." *Journal of Finance* 35 (May/June 1980) 503–20.
4. Z. Bodi and V. Rozansky. "Risk and Return in Commodity Futures." *Financial Analysts Journal* (1980).
5. P. Cootner. "Returns to Speculators: Telser versus Keynes." *Journal of Political Economy* 68 (August 1960), 396–404.
6. ———. "Rejoinder," *Journal of Political Economy* 68 (August 1960), 415–18.
7. K. Dusak. "Futures Trading and Investor Returns: An Investigation of Commodity Risk Premium." *Journal of Political Economy* 81 (October 1973).
8. F. L. A. Grauer. "A Test of the Hypothesis that Backwardation is a Function of the Real Social Risk of Commodity Futures Contracts." Dissertation, Stanford University, Graduate School of Business, 1981, Ch. IV.
9. E. J. Hannan. "Regression for Time Series." In M. Rosenblatt, ed., *Proceedings of a Symposium on Time Series Analysis.* Wiley: New York.
10. L. Hansen and R. Hodrick. "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis." *Journal of Political Economy* 88 (October 1980), 829–53.
11. C. Hardy. *Risk and Risk Bearing.* Chicago: University of Chicago Press.
12. D. A. Hsieh. "Expectations and Efficiencies in Foreign Exchange Markets." Dissertation, Massachusetts Institute of Technology, 1981.
13. L. Telser. "Returns to Speculators: Telser versus Keynes: Reply." *Journal of Political Economy* 66 (August 1960), 404–15.