Resources Policy 51 (2017) 135–150

Contents lists available at ScienceDirect

Resources Policy

journal homepage: www.elsevier.com/locate/resourpol

A historical perspective of the informational content of commodity futures

Viviana Fernandez

Business School, Universidad Adolfo Ibañez, Avenida Diagonal Las Torres 2700, Office 512-C, Santiago, Chile

ARTICLE INFO

JEL codes:
Q02
L72
C13

Keywords:
Premium
Theory of storage
Seemingly unrelated equations

ABSTRACT

This article extends Chinn and Coibion (2014)'s work–*Journal of Futures Markets* 34– on predictive content of commodity futures by considering a more comprehensive database and a longer time span, ranging from 25 to 65 years, and by presenting two extensions: multi-equation estimation of risk premiums and testing for the theory of storage.

The empirical results show that futures-based forecasts for animal and agricultural products and industrial metals tend to be more efficient, in terms of mean absolute error, than random walk based-forecasts at a one-year horizon. On the other hand, based on robust rolling estimates, there is evidence of constant and time-varying risk premiums in agricultural and precious metals, but their statistical significance vary considerably along the sample period. In particular, gold and silver show evidence of a negative time-varying risk premium, as opposed to platinum.

Multi-equation estimation brings efficiency gains in premium gauging, which leads to reject that the futures price is an unbiased estimate of the spot price for all commodity classes. On the other hand, the sampled commodities lend only partial support to the theory of storage, and for the specific case of industrial metals, inventories seem to matter more than interest rates to explain the basis. Altogether, this article finds mixed support for the premium-based model and for the theory of storage.

1. Introduction

Literature on the assessment of the forecast performance of commodity futures prices dates back to the classic paper by Fama and French (1987). The authors examined two models of commodity futures: the theory of storage and expected premium. The theory of storage explains the difference between the futures and spot prices (basis)[1] in terms of interest rates, warehouse costs, and convenience yields. The expected premium model in turn splits a futures price into an expected premium and the current expectation of the change in the spot price.[2] Fama and French found more evidence in favor of futures prices responding to storage-cost factors than in favor of futures prices containing premiums or having the power to forecast spot prices.

More recent contributions in this area are the survey article by Chow et al. (2000), and various empirical applications by Chernenko et al. (2004), Movassagh and Modjtahedi (2005), Modjtahedi and Movassagh (2005), Watkins and McAleer (2006), Alquist and Killian (2010)[3], Reeve and Vigfusson (2011), Chinn and Coibion (2014),[4] Xiao et al. (2015), among others.[5] This article is connected to those of Reeve and Vigfusson (2011), and Chinn and Coibon (2014), in that it analyzes forecast accuracy and unbiasedness of futures prices with respect to spot prices (i.e., absence of risk premiums).

Specifically, based on a sample of 15 commodities belonging to the agricultural, energy, and industrial metals sectors, Reeve and Vigfusson

E-mail addresses: viviana.fernandez@uai.cl.

[1] The basis reflects the relationship between the cash (spot) price and the futures price, and it is obtained by subtracting the futures price from the cash price. Changes in the basis are generally driven by short-term demand and supply fluctuations. If demand is strong and current supply small, spot prices will rise relative to futures price, causing the basis to strengthen. On the other hand, if demand is weak and current supply large, cash prices will fall relative to the futures price, causing the basis to weaken (see, for instance, Hull (2014), chapter 3; http://www.theoptionsguide.com/futures-basis.aspx). Fama and French (1987) define the basis with the opposite sign, that is, $(F_{t,t+h} - S_t)$, where $F_{t,t+h}$ is the futures price at time t for delivery at (t+h), and S_t is the spot price at time t.

[2] That is, $(F_{t,t+h}-S_t)=E_t(P_{t,t+h}) +[E_t(S_{t+h})-S_t]$, where $F_{t,t+h}$ is the futures price today (t) for maturity at time (t+ h), S_t is the current spot price, $E_t(P_{t,\ t+h})$ is the expected premium between today and (t+h), and $[E_t(S_{t+h})-S_t]$, is the expected change in the spot price between today and (t+h).

[3] A more recent discussion on oil price fluctuations is offered by Baumeister and Kilian (2016).

[4] A previous version of this article is Chinn and Coibion (2010).

[5] Related literature on commodity markets include, among others, co-movement between spot and futures prices (e.g., Chang and Lee, 2015; Fernandez, 2015; Nicolau and Palomba, 2015); the impact of trading positions of hedgers and speculators on price formation (e.g., Chen and Chang, 2015), and portfolio strategies (e.g., Han et al., 2016).

http://dx.doi.org/10.1016/j.resourpol.2016.12.002
Received 24 July 2016; Received in revised form 29 November 2016; Accepted 2 December 2016

(2011) concluded that during 1990–2010 futures contracts generally outperformed a random walk without drift forecast-wise, although not by a large extent. However, when the futures price was well-above the spot price,[6] the forecast performance of futures improved in relation to a random walk without drift. In addition, Reeve and Vigfusson found that the futures price and a random walk without drift noticeably outperformed an extrapolation of recent trends (i. e, a random walk with drift).

Chinn and Coibon (2014) in turn analyzed a similar commodity set (plus precious metals) over 1990–2012. They concluded that futures prices of precious and industrial metals, as opposed to those of energy and agricultural products, were generally biased estimators of future spot prices, and found that futures prices were poor predictors of subsequent changes in spot prices. Chinn and Coibon attributed these differences among commodities to contract liquidity.

This article extends Chinn and Coibon (2014)'s work on predictive content of commodity futures by considering a more comprehensive database of 21 commodities and a longer time span, ranging from January 1991-December 2015 (e.g., lead) to January 1962–December 2015 (e.g., cocoa). In addition, this article presents two extensions: multi-equation estimation of risk premiums and testing for the theory of storage.

The empirical findings can be summarized as follows. First, based on robust-rolling estimates, forecast accuracy and the existence of a risk premium depend on the commodity in question and the sample period. In particular, animal and agricultural products and industrial metals tend to be more efficient, in terms of mean absolute error, than random walk based-forecasts at a one-year horizon. Moreover, and in contrast with the evidence reported by Reeve and Vigfusson (2011), there is no clear indication that the performance of a futures-based forecast improves when the futures price exceeds the spot price. In addition, the estimation results show that there is no statistical association between futures contract liquidity and the probability of rejecting the unbiasedness of futures prices with respect to future spot prices. This contrasts with the evidence provided by Chinn and Coibion op cit., who concluded that more illiquid futures contracts are more likely to disprove the unbiasedness of futures prices.

Furthermore, there is evidence of constant and time-varying risk premiums in agricultural and precious metals, but their statistical significance vary considerably over time. In particular, gold and silver show evidence of a negative time-varying risk premium, as opposed to platinum, which is not only an investment commodity but also an input to auto-catalyst and chemical applications.

Second, the use of a multi-equation model brings efficiency gains when estimating premiums. Such extra efficiency leads to reject that the futures price is an unbiased estimate of the spot price for all commodity classes. In this regard, the evidence against the unbiasedness of futures prices here presented is stronger that than reported by Chinn and Coibion op. cit, who only considered single-equation estimation. Third, the sampled commodities lend only partial support to the theory of storage. In particular, inventories seem more relevant than interest rates to explain the basis of industrial metals.

This article is organized as follows. Section 2 briefly presents the premium-based formulation and tests for forecast accuracy and predictive performance. Section 3 describes the characteristics of the futures contracts of the 21 commodities under analysis. Section 4 presents the estimation results, while Section 5 presents two extensions: multi-equation estimation of risk premiums (Section 5.1), and testing for the theory of storage (Section 5.2). Section 6 concludes by summarizing the main findings.

[6] Under the theory of storage, there exists a positive association between the futures/spot price ratio and the inventory level. Reeve and Vigfusson depict such an association for copper and crude oil.

2. Methodology

This section briefly enunciates the hypotheses of interest of this article within the risk-premium framework discussed by Fama and French (1987), and later revisited by Chernenko et al. (2004), Reeve and Vigfusson (2011), and Chinn and Coibon (2014), among others. In addition, this section refers to how assessing forecast accuracy by means of Diebold-Mariano tests and measuring predictive performance by Pesaran and Timmermann (1992)'s non-parametric test.

2.1. Risk premium on the basis

Let us consider the regression of the change in the spot price between t and (t+h) on the current basis (see, for instance, Fama and French, 1987):

$$S_{t+h} - S_t = \alpha + \beta(F_{t,t+h} - S_t) + u_{t+h} \quad (1)$$

where u_{t+h} is a disturbance. The left-hand side variable may be interpreted as the error from a random-walk forecast, and the right-hand side variable as the difference between the forecasted prices from the futures market ($F_{t,t+h}$) and from a random-walk process (S_t).

Under the above formulation, if α=0 and β=1, the futures price is an unbiased estimate of the spot price, i.e., $E_t(S_{t+h})=F_{t,t+h}$, and there is no risk premium. This hypothesis is also referred to as the efficient markets hypothesis (e.g., Chinn and Coibion, 2014). Alternative hypotheses of interest are: (i) α=0: absence of a constant risk premium; (ii) β=0: the futures price provides no extra explanatory power relative to a random-walk process. Or, in other words, the basis is uninformative with respect to future changes in the spot price; and, (iii) β=1: absence of a time-varying risk premium on the basis (or unbiasedness hypothesis as referred to by Chinn and Coibon (2014). See also Chernenko et al. (2004) and Reeve and Vigfusson (2011) for a related discussion.

2.2. Forecast accuracy

Two well-known measures to assessing forecast accuracy in the time series literature are the root-mean squared error (RMSE) and the mean absolute error (MAE), e.g., Franses et al. (2014), chapter 3; Enders (2015), chapter 2:

$$\mathrm{RMSE}(h) = \sqrt{\frac{1}{N-h}\sum_{t=1}^{N-h}\hat{u}_{t+h|t}^2} \quad \mathrm{MAE}(h) = \sqrt{\frac{1}{N-h}\sum_{t=1}^{N-h}|\hat{u}_{t+h|t}|} \quad (2)$$

where h is the forecast-step length (i.e., time horizon) and N is the window length. Specifically, under model specification (1), $\hat{u}_{t+h|t} = \hat{\alpha} + \hat{\beta}(F_{t,t+h} - S_t)$, whereas for a random walk without drift, $\hat{u}_{t+h|t} = S_{t+h} - S_t$.

A fairly standard tool to examine the statistical discrepancy between the RMSEs and the MAEs of specification (1) and a pure random walk is Diebold-Mariano (1995) test. This is based on the following test statistics:

$$\mathrm{DM} = \frac{\bar{d}}{\hat{\mathrm{lrv}}(\bar{d})^{1/2}} \xrightarrow{d} N(0,1) \quad (3)$$

where $\bar{d} = \frac{1}{M}\sum_{t=1}^{M} d_t$, d_t=$L(u^1_{t+h|t})-L(u^2_{t+h|t})$, L(.) is a loss function for competing models 1 and 2, $\hat{\mathrm{lrv}}(\bar{d})$ is a consistent estimate of the long-run asymptotic variance of $\bar{d}$, and M =(N−h). In order to compute $\hat{\mathrm{lrv}}(\bar{d})$, Diebold and Mariano recommend using the Newey-West estimator with a rectangular kernel function and a lag truncation parameter equal to (h−1). For the RMSE and MAE measures, the corresponding loss functions are $L = (\hat{u}^i_{t+h|h})^2$ and $L=|\hat{u}^i_{t+h|h}|$, i=1, 2.

Under the null hypothesis, model specifications 1 and 2 have identical predictive accuracy in expected value, that is, $E(d_t)$=0. For a left-sided test, under the alternative hypothesis model 1 is more