

NBER WORKING PAPER SERIES

THE PREDICTIVE CONTENT OF COMMODITY FUTURES

Menzie D. Chinn
Olivier Coibion

Working Paper 15830
<http://www.nber.org/papers/w15830>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2010

We thank Jonathan McBride for assistance with data collection. Chinn acknowledges the financial support of the University of Wisconsin Center for World Affairs and the Global Economy. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Menzie D. Chinn and Olivier Coibion. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Predictive Content of Commodity Futures
Menzie D. Chinn and Olivier Coibion
NBER Working Paper No. 15830
March 2010
JEL No. G13,Q43

ABSTRACT

This paper examines the relationship between spot and futures prices for a broad range of commodities, including energy, precious and base metals, and agricultural commodities. In particular, we examine whether futures prices are (1) an unbiased and/or (2) accurate predictor of subsequent spot prices. While energy futures prices are generally unbiased predictors of future spot prices, there is much stronger evidence against the null for other commodity markets. This difference appears to be driven in part by the depth of each market. We find that over the last five years, it is much harder to reject the null of futures prices being unbiased predictors of future spot prices than in earlier periods for almost all commodities. In addition, futures prices do approximately as well as a random walk in forecasting future spot prices, and vastly outperform a reduced form empirical model.

Menzie D. Chinn
Dept. of Economics
University of Wisconsin
1180 Observatory Drive
Madison, WI 53706
and NBER
mchinn@lafollette.wisc.edu

Olivier Coibion
Department of Economics
College of William and Mary
115 Morton Hall
Williamsburg, VA 23188
ocoibion@wm.edu

“Policymakers and other analysts have often relied on quotes from commodity futures markets to derive forecasts of the prices of key commodities... The poor recent record of commodity futures markets in forecasting the course of prices raises the question of whether policymakers should continue to use this source of information and, if so, how.” Ben Bernanke, June 9, 2008

Commodity prices have arguably played an important role in accounting for historical macroeconomic fluctuations. The two oil price shocks in the 1970s remain the most common explanation for the Great Inflation of the 1970s and the stagflationary patterns observed after these episodes.¹ Hamilton (2009) argues that the oil price run-up of 2007-2008 can account for much of the early stages of the Great Recession. Hamilton (1983) and Bernanke, Gertler and Watson (1997) note the broader point that most US recessions have been preceded by large oil price increases. The evidence linking commodity price shocks to macroeconomic fluctuations is not limited to oil prices, however. For example, the twin oil price shocks of the 1970s were accompanied by twin food price shocks of similar magnitude, a point emphasized early on by Bosworth and Lawrence (1982) and more recently by Blinder and Rudd (2008). In addition, small developing economies have often been dependent on a primary commodity for much of their exports (e.g. Chile and copper) and have, as a result, experienced dramatic boom-bust patterns as a result of commodity price changes.

Given this historical relationship between commodity prices and macroeconomic fluctuations, forward-looking policy-makers and researchers have long been interested in predicting commodity price movements.² This paper studies one source of information about expected spot prices: futures prices. In particular, we examine whether futures prices are (1) an unbiased and/or (2) accurate predictors of subsequent spot prices, in the markets for energy, precious metals, base metals, and agricultural commodities. While there is a long literature

¹ See Blinder and Rudd (2008) for a recent exposition of this viewpoint and Barksy and Kilian (2002) for a contrarian view.

² See Wu and McCallum (2005) and Chen et al (2009) for examples.

studying futures prices for energy markets, we build on this literature by extending the analysis to other commodity markets and by emphasizing recent changes in the properties of futures prices. In our view, a re-examination is warranted both because of recent public policy concerns about sharp movements in a broad range of commodity prices as well as the fact that the use of futures for non-energy markets has grown rapidly in recent years. For example, the volume of trade in the base metals and agricultural goods 3-month and 6-month futures markets considered in this study grew by an average rate of 30 percent per year from 2005 to 2009, compared to less than 20 percent per year for energy markets.

Our first contribution is to show, using commodity futures data of multiple horizons, that while energy futures prices are overall not inconsistent with the unbiasedness hypothesis, we can reject the null of unbiased predictions of future spot prices for most other commodity markets, particularly once one explicitly takes into account time-varying heteroskedasticity in residuals via GARCH. One possible explanation for this result is the difference in the depth of these respective markets. Consistent with this, we find that commodity futures markets which had heavier average trading volumes were more likely to have futures prices for which we could not reject the null of unbiasedness.

Our second contribution is to study the unbiasedness hypothesis over time via rolling regressions. For energy markets, we are consistently unable to reject the null of unbiased futures prices, consistent with the notion that these markets have long been heavily in use.³ Other commodities, on the other hand, have experienced sometimes dramatic periodic swings in coefficient estimates. Yet, a common pattern across non-energy commodities is the fact that, with few exceptions, we cannot reject the null of unbiased futures prices over the last five years.

³ One exception is heating oil at the end of the sample. The unusual behavior of the basis during the first quarter of 2009 leads to a dramatic departure from the null over this time period.

Thus, there appears to have been a common increase in the efficiency of these markets over recent years. This result is particularly striking given the fact that futures prices were heavily criticized in 2008 for having missed the reversal of the 2007-08 price run-up, as illustrated in the quote from Chairman Bernanke in the summer of 2008.

Our third contribution is to study the magnitude of forecast errors from futures prices relative to random walk predictions as well as forecasts from a reduced form empirical model. Focusing on the period since 2003, we find that futures prices have fared no worse than random walk forecasts, and sometimes much better. In addition, futures prices have vastly outperformed the out-of-sample forecasts from our empirical model, with all of these results being broadly similar across commodity types.

Section 1 discusses the theory of storage and its implications for properties of futures prices as predictors of future spot prices, as well as some of the previous empirical evidence of futures prices. Section 2 contains a data description and our baseline empirical results. Section 3 extends our results to a GARCH framework while section 4 considers rolling estimates of our baseline specification. Section 5 presents comparisons of the forecasting accuracy of futures prices relative to a random walk and an empirical model. Section 6 concludes.

1. Theory and Previous Literature

The notion that the futures price is the optimal forecast of the spot price is an implication of the efficient market hypothesis. In an efficient market, new information is reflected instantly in commodity prices. If this is true, then price patterns are random, and no system based on past market behavior can do other than break even. The link between efficiency and forecastability arises from realizing the difference between the current futures price and the future spot price

represents both the forecasting error and the opportunity gain or loss realized from taking certain positions. The requirement that the forecasting error is zero on average is consistent with both market efficiency (the absence of profitable arbitrage opportunities) and the unbiasedness property of the forecaster (zero forecasting error on average).

The futures price of a storable commodity such as crude oil is determined by the spot price and the cost incurred while the commodity is stored awaiting delivery some time in the future. The cost associated with holding the commodity until the delivery date is known as the cost-of-carry. The cost-of-carry consists of the cost of storing oil in a tank (and perhaps insurance) and the financial cost in the form of the opportunity cost of holding oil, or the cost of funding, and perhaps a risk premium.⁴

The spot/futures pricing relationship is based on the assumption that market participants are able to trade in the spot and futures markets simultaneously, i.e. they can utilize spot/futures arbitrage. The relationship between the futures rate and the current spot rate for petroleum is given by:

$$f_{t,t-k} - s_{t-k} = d_{t,t-k} + Q_{t,t-k} \quad (1)$$

where $f_{t,t-k}$ is the observed (log) time $t-k$ futures contract price that matures at time t , and s_{t-k} is the time $t-k$ spot rate, $d_{t,t-k}$ the log cost-of-carry (the sum of storage costs minus convenience yield, plus interest costs and a risk premium), and $Q_{t,t-k}$ is a term accounting for the marking-to-market feature of futures. The object on the left hand side of (1) is called the “basis” in the commodity futures literature.⁵

⁴ Williams and Wright (1991) provide an excellent overview to the behavior of commodity prices and futures. See also Fama and French (1987) and Pindyck (2001).

⁵ The discussion and notation is based upon the exposition in Brenner and Kroner (1995).

If we assume the log spot rate follows a time random walk with drift, and expectations are rational, then the time $t-k$ *expectation* of the change in the spot rate will equal the basis and the marking-to-market term. Hence, in the regression of change in the spot rate on the basis,

$$s_t - s_{t-k} = \alpha + \beta(f_{t,t-k} - s_{t-k}) + \varepsilon_t \quad (2)$$

α subsumes the terms in the right hand side of (1), as well as the parameters defining the time series process governing the spot rate, while $\alpha = 0$ and $\beta = 1$ if the basis is the optimal predictor of the change in the spot rate. It is important to recall that rejection of the null hypothesis is then a rejection of a composite hypothesis, including both market efficiency and unbiased expectations.

The literature examining the behavior of futures markets is fairly extensive. A number of studies have examined the efficiency of futures markets and have investigated the related issue of the forecastability of spot energy prices. Unsurprisingly, the conclusions are quite diverse. A number of studies provide evidence for efficient markets and an equally large number provide evidence that contradicts an efficient market (unbiased futures price prediction) interpretation. For energy markets, Serletis (1991) found evidence consistent with efficient crude petroleum markets. Bopp and Lady (1991), however, found that either the spot or the futures price can be the superior forecasting variable depending on market conditions, and the information content of the two price series is essentially the same.

The more recent literature has focused on the long-run properties of the spot and forward prices, in the context of cointegration.⁶ We focus on the change in the spot rate, and its relation to the basis, reserving the analysis of long run dynamics to future study.

⁶ See for instance, Crowder and Hamed (1993), Moosa and Al-Loughani (1994), Herbert (1993) and Walls (1995).

2. Baseline Empirical Analysis

We obtained data for the following commodities – petroleum (West Texas Intermediate, or WTI), natural gas (Henry Hub), gasoline (Gulf Coast)⁷, heating oil (No. 2, Gulf Coast), gold, silver, aluminum, copper, lead, nickel, tin, corn, soybean and wheat. All the futures and spot prices are from Bloomberg, and are traded on either the New York Mercantile Exchange (NYMEX), the London Mercantile Exchange (LME) or the Chicago Mercantile Exchange (CME). Appendix 1 provides details on the specific series used for each commodity type.

Figures 1 through 14 plot the spot prices for each series as well as the log of the 3-month and 6-month basis. The most pronounced changes in the basis tend to be associated with large changes in spot prices, indicating that many of these price fluctuations are expected to be transitory. This can be clearly seen in the case of natural gas. During each of the dramatic price spikes in the 4th quarter of 2000, the 1st quarter of 2003, and the 3rd quarter of 2005, the basis turned very negative in expectation of rapid price reversals, each of which promptly materialized. There are, however, some notable exceptions. For example, when oil prices peaked in 2008 after a rapid increase over the previous year, both the 3-month and 6-month basis were close to zero, indicating that market participants expected this price increase to be permanent. A similar pattern is apparent for most other commodity markets during this time period. The subsequent decline in oil prices during the second half of 2008, on the other hand, was associated with a pronounced increase in the basis and therefore a market expectation of at least a partial reversal of the price decline, a pattern which materialized in subsequent months.

⁷ Starting in 2005, the NYMEX phased in a new set of futures for gasoline to take the place of the previous series. The reason has to do with the changing content of ethanol in gasoline. In order to avoid complications involving changing variable definitions, we restrict our analysis to the older gasoline futures, which end at various points in 2006.

These individual episodes seem to indicate that, in many cases, futures markets have been able to correctly predict subsequent spot prices. In order to evaluate more formally the properties of the relationship between the spot and futures prices, we now move to a statistical analysis.

Equation (2) is estimated using OLS using spot and futures prices sampled at a monthly frequency. These data are all sampled at the end of month, and hence allow proper synchronization of prices. Note however that because horizons of 3, 6 and 12 months are used, and the data is of monthly frequency, the regression residuals incorporate overlapping information. Under the null hypothesis of efficient markets (risk neutrality and rational expectations), the regression residuals will exhibit serial correlation. In order to obtain consistent estimates of the standard errors, necessary to conduct proper statistical inference, we calculate heteroskedasticity and serial correlation robust standard errors.⁸ For each series and futures horizon, we use the maximum sample available between 1990:01 and 2009:10. We present estimates of β , the coefficient on the basis, and test statistics for the null hypothesis that $\alpha = 0$ and $\beta = 1$ in Table 1.

Energy Markets: For the crude oil market, the estimates for β at the 3, 6 and 12-month horizons are not statistically distinguishable from unity, as documented in Chinn, Leblanc and Coibion (2005) and Alquist and Kilian (2010). In addition, one cannot reject the joint hypothesis of efficient markets ($\alpha = 0$ and $\beta = 1$) at any horizon.

It is also true that in none of the cases are futures good predictors of subsequent spot prices. At the 3-month horizon, the basis accounts for only 4 percent of total variation in changes

⁸ Using monthly data, if the futures mature 12 months (4 quarters) in the future, then a moving average process of 11 (=12-1) is induced. The Newey-West standard errors are calculated using a Bartlett window and lag order set equal to $k-1$.

in spot rates; at the 12-month horizon, this proportion is only 6 percent. One can also not reject the null of market efficiency for gasoline at any horizon, but the basis can account for a larger fraction of subsequent spot price changes: from 16 percent at the 3-month horizon to 26 percent at the 12-month horizon. Natural gas futures also appear to be more correlated with future spot prices, especially at the longer horizon. While the coefficient is quantitatively and statistically different from the posited value of unity using 3 month futures,⁹ at longer horizons, the slope coefficient is not statistically distinguishable from unity at the 5 percent statistical significance level, and the proportion of variance explained is nearly 25 percent at the long horizon.

Finally, in the case of heating oil, we find that we cannot reject the market efficiency hypothesis for our three futures contracts, except at the 6-month horizon, although the explanatory power of the basis is quite low, accounting for approximately 10% of the variation in future spot prices. In short, the evidence from energy markets is mostly consistent with the efficient markets hypothesis, but the ability of futures prices to quantitatively predict subsequent spot prices appears limited, most notably for oil prices.

Precious Metals: Replicating the same analysis for gold and silver yields strikingly different results: we can reject the null of unbiasedness and the efficient markets hypothesis at any standard level of statistical significance for all horizons. Indeed, point estimates of the coefficient on the basis are consistently negative, indicating that the basis tends to predict changes in future spot prices in the wrong direction. We return to the source of this result in our discussion of subsample-estimates in section 4.

⁹ A similar finding is obtained by Chernenko et al. (2004).

Base Metals: The results for base metals appear to be much more mixed. For example, we can reject the null that $\beta = 1$ at all horizons for copper and lead, but cannot reject the null for the others (with the sole exception of aluminum at the 1-year horizon). However, this appears to be driven primarily by the very large standard errors associated with the parameter estimates. For example, the point estimates of β for aluminum at the 3 and 6 month horizons are actually negative, but the estimates are so imprecise that we cannot reject the null at standard levels of statistical significance. With nickel and tin, the point estimates of β are all positive, but again the standard errors do not allow for any meaningful inference about the null. We return to this issue in section 3.

Agricultural Products: For corn and soybean, we cannot reject the null that $\beta = 1$ at either the 3 or 6 month horizon.¹⁰ Unlike the case of base metals, the standard errors associated with the estimates of β are in the same neighborhood as those found for energy and precious metals, so these results do not appear to be driven by excessively large standard errors. We can, however, reject the null of market efficiency for corn at both horizons at the 5 percent level. For wheat, there is strong evidence against the null of unbiasedness at both horizons considered.

In summary, we find a remarkable contrast between energy futures markets and futures prices for other commodities. For energy markets, futures prices are consistent with unbiasedness and the more general predictions of market efficiency (with few exceptions). The other markets, on the other hand, fare much worse. The nulls of unbiasedness and market efficiency are almost always rejected, with most of the exceptions occurring with base metals for

¹⁰ We restrict our analysis to 3 and 6-month futures in the case of agricultural products because 12-month futures are not available over sufficiently long samples for these goods.

which standard errors are too large to reject any meaningful economic null, a point reinforced in section 3.

3. Time-Variation in Shock Volatility

In the previous analysis, we allowed for serial correlation and heteroskedasticity of a general form, using robust standard errors to make inferences regarding statistical significance. However, we know that asset prices, including derivatives based on underlying commodities, often evidence systematic conditional heteroskedasticity. This understanding motivates a formal GARCH approach to modeling the heteroskedasticity.

First, we test for the presence of conditional heteroskedasticity. Formal tests of the null of no ARCH effects in the simple basis regressions are rejected the 1% level for all commodity markets at all horizons. Thus, modeling the heteroskedasticity in errors is likely to increase the efficiency of the estimates. Table 2 then presents GARCH(p,q) estimates of the basis specifications. The p and q terms are chosen via the AIC criterion for each commodity at each horizon.

The use of GARCH reduces the standard errors of our point estimates by a substantial amount, approximately 50% on average across commodities and horizons. The results reinforce the predominant finding from the OLS estimates: namely, that it is difficult to reject the null of unbiasedness for energy futures but not for other commodities. For example, while the standard errors of the point estimates of β are consistently reduced across commodity types, we can only reject the null of unbiasedness in energy futures for natural gas at the 3-month horizon and heating oil at the 1-year horizon.¹¹ This reflects the fact that the point estimates of β are almost

¹¹ For gasoline 1-year futures, multiple missing observations throughout the sample prevent us from estimating equation (2) by GARCH.

all much closer to one than when using OLS. For non-energy markets, we can reject the null of unbiasedness for every commodity for at least one forecasting horizon. In particular, for aluminum, nickel and tin, the three base metals for which we could not consistently reject the null of unbiasedness under OLS, we now find much stronger rejections of the null, despite the fact that standard errors remain large relative to those of other commodity markets. Similarly, while we could not reject the null of unbiasedness for corn and soybeans under OLS at any horizon, the GARCH results yield rejections of the null at one of the two horizons for each product.

In short, the GARCH estimates provide additional evidence consistent with the null that energy futures are unbiased predictors of future spot prices, but that this is not the case for other commodity futures markets. The one qualitative difference worth noting between the OLS and GARCH estimates is that under the latter, we can reject the null of efficient markets (i.e. the joint hypothesis of $\alpha = 0$ and $\beta = 1$) for nearly every commodity at nearly every horizon, with the sole exceptions being gasoline and aluminum at the 3-month horizon and natural gas at the 6-month horizon.

An immediate question that comes to mind is why there might be such a difference in the behavior of futures markets across commodities. A possible explanation is the relative depth of these markets, i.e. how much are these markets actually used? Sparsely used markets are more likely to be dominated by a few insiders so that futures prices might not be very good predictors of future prices, whereas deeper markets, i.e. those markets in which there are many buyers and sellers, are more likely to satisfy no-arbitrage conditions. Energy futures markets have long been used by financial and real market participants, whereas other commodities have historically been traded primarily on spot markets, with futures markets being relatively new additions.

To assess whether there is a link between the depth of each market and how closely the no-arbitrage condition holds. Figure 15 plots the (log) average monthly volume of trades for each type of commodity (averaged across 3-month and 6-month futures) from 2006-2009 versus the average absolute value of the t -statistic for the null of unbiasedness for 3-month and 6-month futures for each commodity type.¹² There is a clear negative correlation between quantity of trades in each commodity futures market and how well the unbiasedness assumption matches the data. Despite how few observations there are, this negative relationship is statistically significant at the 5 percent level. Of course, one should be wary of imposing a causal relationship: an alternative interpretation could be that people are more willing to trade in those markets for which futures prices tend to be more unbiased predictors of subsequent spot prices. However, given the previous lack of evidence on how futures markets fared in non-energy markets, there is ground to view this alternative interpretation with some suspicion.

4. Time Variation in Coefficients

Our baseline approach is to make use of the longest time sample available for each commodity. However, an important feature of these futures markets is that they are relatively new. For example, futures prices for base metals are only listed starting in the late 1990s. In addition, even though some futures markets may have been in existence over the whole sample, this does not mean that there was enough trading volume for anything like the efficient markets hypothesis

¹² We consider only trading volumes between 2006 and 2009 because earlier quantity data is not available for many commodities, particularly base metals. We omit gasoline from the figure because of the change in gasoline types in 2006 as well as lead because we do not have volume measures for 3-month and 6-month figures for this commodity. We should note, however, that we do have 1-month volume measures, which yield an average log volume of 10.5 and the average t -stat for lead is 3.5, so if we were to add it to Figure 15 it would be somewhat below the trend line but would not qualitatively change the results. Note also that we omit 12-month futures because these are not available for all commodity types.

to be expected to hold. In this section, we therefore consider 5-year rolling estimates of equation (2) for each commodity and forecasting horizon.

Figure 16 presents the time-varying estimates of β and their associated two-standard deviation confidence intervals. Turning first to energy markets, which appear to be the least biased predictors of future spot prices over the whole sample, the Figures reveal some notable caveats to our earlier findings. For example, in the case of 6-month oil futures, the null of unbiasedness can be rejected during the very early part of our sample, approximately 1991:1-1996:1, as well as for a narrow period toward the end of the sample (the 5-year period ending in very early 2009). Similarly, using the whole sample, we found that we could reject the null of unbiasedness for natural gas at the 3-month horizon but not the one-year horizon. The sub-sample estimates thus present a more nuanced picture in which, for both futures horizons, there were 5-year periods during which one could reject the null and others during which one could not reject the null. Note that this sub-sample instability can help explain why previous studies have reached such contrasting results when testing the null over different time samples.

Particularly striking are the results for heating oil, which exhibit a pronounced decline in the basis coefficient (to approximately -1 for 3-month and 6-month futures) in the last part of the sample, a feature which is absent for oil and natural gas futures. The contrast with respect to oil futures is notable because the spot prices for the two series are very highly correlated (0.992 for logs of the prices over the whole sample). The decline in the coefficient on the basis for heating oil can be attributed to the fact that, when heating oil prices were near their minimum levels in the first quarter of 2009, the 3-month and 6-month basis turned negative and failed to predict the subsequent partial recovery in heating oil prices (see Figure 3). The downward movement in the basis followed by an increase in prices explains why the coefficient on the basis turns negative

for heating oil when this time period is included in the rolling estimates. This is particularly striking given that the 3-month and 6-month basis for oil jumped *up* during this same time period, thereby correctly anticipating the recovery in oil prices, as can be seen in Figure 1. Given the very high correlation between the spot prices of these two commodities, the divergent behavior of the basis over this time period is puzzling.

The time-varying estimates are also revealing for non-energy commodities. In particular, we find that, with almost no exceptions, we cannot reject the null of unbiasedness for the last five years of our sample, despite numerous departures from the null in the early and middle periods of the sample. For example, our baseline OLS estimates for gold and silver strongly rejected the null of unbiasedness with estimated coefficients on the basis that were negative at all horizons. For gold, this result appears to be driven by the early 2000's, during which the basis fell significantly while spot prices began to rise persistently (see Figure 5), thereby generating a negative coefficient on the basis in estimates of equation (2). However, when focusing on the last five to six years, the estimated coefficient on the basis is close to and not statistically different from one. Similar results hold for silver.

The pattern for base metals is also consistent with this. For each base metal, there are periods at the beginning or the middle of the sample during which unbiasedness can be strongly rejected, and yet, at the end of the sample, we can only reject the null of unbiasedness for one out of the fifteen commodity/horizon combinations, namely aluminum at the 12-month horizon. Thus, as with gold and silver, the evidence against the null of unbiasedness uncovered in the baseline estimates of equation (2) appear to be driven almost exclusively by the early periods of the sample. Over the last five years, there is little evidence against the null of unbiasedness for any of the base metals in our dataset. Similar qualitative results obtain for agricultural products.

We should note that this feature of the data, the increasing difficulty of rejecting the null of unbiasedness across commodities over the last five years, is also consistent with the earlier interpretation of our results as being driven by the relative depth of these markets. For example, energy futures markets have long been used by financial and real market participants and we cannot reject the null of unbiasedness in these markets for much of the sample. On the other hand, base metals and agricultural futures have both seen much more rapid rates of growth in trading volumes than energy futures. For example, trading volume in base metals has on average grown by nearly 30 percent a year between 2005 and 2009 while trading in energy futures grew at less than 20 percent a year over the same period. Similarly, trading volume of agricultural futures has grown by 17 percent a year since 1999 (and 31 percent a year since 2005) compared to 10 percent a year for energy futures over the same time period.

This rapid increase in trading volume means that by 2009, average trading volume in agricultural futures and base metals had reached 57 percent and 23 percent of the average trading volume in energy products (again using only 3-month and 6-month futures, and excluding lead and gasoline from all calculations). In short, trading in agricultural and base metal futures has grown disproportionately rapidly and this increase in the depth of these markets could potentially explain why we are observing a greater difficulty in rejecting the null of unbiasedness in futures in recent years relative to earlier in the sample. We should note, however, that precious metals are an outlier here: trading in these futures markets has grown much more slowly (around 6 percent a year since 1999) and in 2009 experienced trading volumes around 10 percent of volumes in energy markets, yet the pattern of failing to reject the unbiasedness hypothesis in recent years also holds for precious metals.

5. Forecast Comparison

We also evaluate the relative forecasting abilities of futures, by treating the futures price as the expectation of the future spot rate. We then compare these implied forecasts against a simple naïve forecast and forecasts generated from a simple univariate time series model. The naïve forecast is merely that the current spot rate is the best guess of the future rate (i.e., a random walk without drift). The time series model takes each log price as an ARIMA(1,1,1) process.

We use a rolling regressions methodology in conjunction with out-of-sample forecasting to assess the estimated specifications. For instance, we estimate the 3 month horizon regression for 1990m01 to 2002m10, forecast out 3 months, predicting the 2003m01 spot rate. Then “roll” the sample up to 1990m02 to 2002m11, re-estimate and then re-forecast, predicting the 2003m02 spot rate. This process is repeated until 2009m10 is predicted. This yields a set of forecasts over the 2003m01-2009m10 period.

As indicated in Table 3, futures prices do a quite good job, in terms of both unbiasedness and smallest forecast errors, for oil, gasoline and heating oil. The ratios of the Root Mean Squared Errors (RMSE) and Mean Absolute Errors (MAE) for futures prices and ARIMA models relative to that of the no-drift random walk are presented. *Numbers less than one indicate smaller forecast errors than random walk.* Significance levels are also indicated.¹³

For the oil market, futures prices marginally outperform the random walk at the 3-month horizon, but not at the 6-month or 12-month horizons, with none of the differences being statistically significant, thereby largely confirming the results of Alquist and Kilian (2010). In addition, the random walk consistently outperforms the ARIMA model. Very similar results

¹³ Significance levels are computed via simulations of the random walk model, in which we generate time series under the random walk by drawing from the actual sequence of errors with replacement. These simulated series are then used to generate a distribution of RMSE's and MAE's from the random walk specification.

obtain for heating oil, but futures perform noticeably better than the random walk for both natural gas and gasoline, although again the differences are not statistically significant.

Futures prices also outperform the random walk at most horizons for precious metals and agricultural commodities, although the differences are again not statistically significant. For base metals, the random walk modestly outperforms futures prices at all horizons. By far, the ARIMA model fares worse, and is often statistically rejected against the null of the random walk. In short, futures prices fare well in forecast comparisons, frequently outperforming the random walk prediction, a notable result given the well-known difficulty of outperforming random walk predictions for other financial markets.¹⁴

6. Conclusion

Commodity prices have long played an important role in accounting for economic fluctuations. Forecasting changes in commodity prices is therefore an important component for forward-looking policy-makers. The growing use of futures markets has raised the question of how much information these prices incorporate about future movements in spot prices. We show that while energy futures can adequately be characterized as unbiased predictors of future spot prices, there is much stronger evidence against the null of unbiasedness for other commodities, especially once one explicitly takes into account time-varying heterogeneity in shocks variances. In part, this failure of futures markets for many commodities to satisfy the unbiasedness hypothesis likely reflects the fact that these markets suffered from only light trading volumes. In recent years, as the depth of these markets has increased, we find much weaker evidence against the null of unbiasedness. In addition, futures prices have frequently outperformed random walk

¹⁴ The most well-known of these cases is likely exchange rates, as famously illustrated in Meese and Rogoff (1983), and more recently by Cheung, Chinn and Garcia Pascual (2005).

predictions since 2003, despite substantial and protracted price changes, and vastly outperform reduced form statistical models of price changes. This result leads us to be cautiously optimistic about the broader use of futures prices as predictors of subsequent spot price movements, particularly for those markets which continue to be actively used by a wide range of financial and real market participants.

Appendix 1

All spot and futures prices are downloaded from Bloomberg. The specific series used and associated mnemonics for spot and futures prices are:

Commodity	Market	Spot Ticker	Futures Ticker
<i>Energy</i>			
Oil	NYMEX	USCRWTIC	CL
Natural Gas	NYMEX	NGUSHHUB	NG
Gasoline	NYMEX	MOINY87P	HU
Heating Oil	NYMEX	NO2INYPR	HO
<i>Base Metals</i>			
Aluminum	LME	LMAHDY	LA
Copper	LME	LMCADY	LP
Lead	LME	LMPBDY	LL
Nickel	LME	LMNIDY	LN
Tin	LME	LMSNDY	LT
<i>Precious Metals</i>			
Gold	NYMEX	GOLDS	GC
Silver	NYMEX	SILV	SIA
<i>Agricultural</i>			
Corn	CME	CORNCH2Y	C_
Soybeans	CME	SOYBCH1Y	S_
Wheat	CME	WEATCH2S	W_

Appendix 2

To generate the forecasts using time series models, the following algorithm was used.

1. An ARIMA(1,1,1) is estimated.
2. Estimate using the in-sample period (up to 2002m12); then roll recursively the estimation until all the out-of-sample observations are exhausted.
3. Forecasts are compared on a 2003m01-2009m10 sample.

References

- Alquist, Ron and Lutz Kilian, 2010, "What Do We Learn from the Price of Crude Oil Futures?" forthcoming in *Journal of Applied Econometrics* 25(4).
- Barsky, Robert B. and Lutz Kilian, 2002, "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in B. Bernanke and K. Rogoff (eds) *NBER Macroeconomics Annual* 2001, 137-183.
- Bernanke, Ben S., Mark Gertler, and Mark Watson, 1997. "Systematic Monetary Policy and the Effects of Oil price Shocks," *Brookings Papers on Economic Activity* 1, 91–142.
- Blinder, Alan S. and Jeremy B. Rudd, 2008, "The Supply-Shock Explanation of the Great Stagflation Revisited," forthcoming in Michael Bordo and Athanasios Orphanides, eds. The Great Inflation, University of Chicago Press.
- Bopp, Anthony E. and George M. Lady, 1991, "A Comparison of Petroleum Futures versus Spot Prices as Predictors of Prices in the Future," *Energy Economics* 13(4):274-282.
- Bosworth, Barry P. and Robert Z. Lawrence, 1982, Commodity Prices and the New Inflation, The Brookings Institution.
- Brenner, Robin and Kenneth Kroner, 1995, "Arbitrage, Cointegration, and Testing the Unbiasedness Hypothesis in Financial Markets," *Journal of Financial and Quantitative Analysis* 30(1) (March), pp. 23-42.
- Chen, Yu-Chin, Kenneth Rogoff, and Barbara Rossi, 2009, "Can Exchange Rates Forecast Commodity Prices?" forthcoming in the *Quarterly Journal of Economics*.
- 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," *International Finance and Discussion Papers* No. 808 (Washington, DC: Board of Governors of the Federal Reserve System, June).
- Cheung, Yin-Wong, Menzie D. Chinn, and Antonio Garcia Pascual. 2005. "Empirical Exchange Rate Models of the Nineties: Are Any Fit to Survive?" *Journal of International Money and Finance* 24 (November): 1150-1175.
- Chinn, Menzie D., Olivier Coibion, and Michael LeBlanc, 2005, "The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline Markets," NBER Working Paper No. 11033 (January).
- Crowder, William and Anas Hamed, 1993, "A Cointegration Test for Oil Futures Market Efficiency," *Journal of Futures Markets* 13(8), pp. 933-41.

Executive Office of the President, 2001, *Report of the National Energy Policy Development Group* (Washington, DC: U.S. GPO).

Fama, Eugene F. and Kenneth R. French, 1987, "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage," *The Journal of Business*, 60(1) 55-73.

Hamilton, James D., 1983, "Oil and the Macroeconomy Since World War II," *Journal of Political Economy* April 228-248.

Hamilton, James D., 2009, "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity*, Spring: 215-259.

Herbert, John 1993, "The Relation of Monthly Spot to Futures Prices for Natural Gas," *Energy* 18(1), pp. 1119-1124.

Meese, Richard A. and Kenneth Rogoff, 1983, "Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?" *Journal of International Economics* 14, 345-373.

Moosa, Imad and Nabeel Al-Loughani, 1994, "Unbiasedness and time varying risk premia in the crude oil futures market," *Energy Economics* 16(2), pp. 99-105.

Pindyck, Robert S., 2001, "The Dynamics of Commodity Spot and Futures Markets: A Primer," *Energy Journal* 22(3): 1-29.

Serletis, Apostolos, 1991, "Rational Expectations, Risk and Efficiency in Energy Futures Markets," *Energy Economics* 13(2):111-115.

Walls, David W., 1995, "An Econometric Analysis of the Market for Natural Gas Futures," *Energy Journal* 16(1): 71-83.

Williams, Jeffrey and Brian Wright, 1991, *Storage and Commodity Markets* (Cambridge: Cambridge University Press).

Wu, Tao and Andrew McCallum, 2005, "Do Oil Futures Prices Help Predict Future Oil Prices?" Federal Reserve Bank of San Francisco *Economic Letter* 2005-38.

Table 1: Regressions of Price Change on the Basis

	3 months						6 months						1 year					
	β	se(β)	Wald	adj. R ²	SER	N	β	se(β)	Wald	adj. R ²	SER	N	β	se(β)	Wald	adj. R ²	SER	N
<i>Energy Products</i>																		
Oil	1.06	(0.45)	0.65	0.04	0.18	238	0.84	(0.43)	1.35	0.05	0.25	238	0.76	(0.37)	2.05	0.06	0.31	238
Natural Gas	0.63***	(0.09)	10.95***	0.13	0.27	223	0.80*	(0.12)	1.53	0.18	0.35	220	1.18	(0.22)	0.37	0.24	0.42	214
Heating Oil	0.82	(0.16)	1.47	0.10	0.18	238	0.64*	(0.21)	3.06**	0.07	0.25	238	0.87	(0.31)	1.18	0.10	0.33	231
Gasoline ***	0.97	(0.24)	0.02	0.16	0.17	203	0.97	(0.19)	1.01	0.22	0.20	205	1.27	(0.27)	1.56	0.26	0.28	141
<i>Precious Metals</i>																		
Gold	-1.13***	(0.74)	4.15**	0.01	0.07	238	-1.02***	(0.52)	7.60***	0.02	0.09	238	-0.74***	(0.44)	7.95***	0.03	0.13	234
Silver	-2.01***	(0.66)	10.57***	0.03	0.11	238	-1.33***	(0.71)	6.22***	0.03	0.16	238	-0.93***	(0.56)	8.65***	0.05	0.19	232
<i>Base Metals</i>																		
Aluminum	-0.27	(0.87)	1.13	-0.01	0.12	145	-0.08	(1.32)	0.35	-0.01	0.18	142	-1.37**	(1.07)	2.85*	0.05	0.23	136
Copper	-0.71**	(0.71)	2.90*	0.00	0.18	145	-0.66*	(0.96)	1.59	0.00	0.27	142	-1.39***	(0.78)	5.33***	0.09	0.32	136
Lead	-0.29***	(0.46)	4.44**	-0.01	0.18	145	-0.51***	(0.56)	4.31**	0.00	0.27	142	-0.45***	(0.53)	3.80**	0.00	0.39	133
Nickel	0.69	(1.09)	0.44	0.00	0.21	145	0.70	(1.48)	0.51	0.00	0.33	142	1.24	(0.98)	0.55	0.03	0.49	136
Tin	1.10	(0.91)	0.77	0.00	0.14	145	0.65	(1.21)	0.45	0.00	0.22	138	0.38	(1.51)	0.72	-0.01	0.31	133
<i>Agricultural Products</i>																		
Corn	0.96	(0.25)	4.66**	0.14	0.13	238	0.77	(0.30)	4.50**	0.13	0.20	238						
Soybean	1.25	(0.29)	0.45	0.15	0.12	238	1.02	(0.33)	0.06	0.17	0.16	238						
Wheat	0.53***	(0.15)	6.52***	0.07	0.13	238	0.34***	(0.18)	9.40***	0.03	0.19	235						

Note: The table presents estimated results by OLS of equation (2) in the text for different commodities and futures prices horizons. For oil, corn, soybean, and wheat, the spot price is replaced with the 1-month futures price. For gasoline, only the old version (“HU”) of futures are used in the regressions. Statistical significance at the 10%, 5%, and 1% level are denoted by *, **, and *** respectively. For β , the null is that $\beta=1$. Newey-West HAC standard errors, using truncation equal to the futures horizon minus one.

Table 2: GARCH Regressions of Price Change on the Basis

	3 months						6 months						1 year					
	β	se(β)	Wald	adj. R ²	SER	N	β	se(β)	Wald	adj. R ²	SER	N	β	se(β)	Wald	adj. R ²	SER	N
<i>Energy Products</i>																		
Oil	0.83	(0.29)	14.12***	-0.01	0.19	238	1.06	(0.14)	36.43***	0.00	0.26	238	1.01	(0.07)	12.19***	-0.01	0.32	238
Natural Gas	0.66***	(0.11)	8.67***	0.11	0.27	223	0.96	(0.07)	0.15	0.15	0.35	220	1.10	(0.09)	35.70***	0.18	0.44	214
Heating Oil	0.99	(0.13)	4.01**	0.06	0.19	238	1.00	(0.10)	15.77***	0.02	0.26	238	0.74***	(0.09)	6.19***	0.06	0.33	231
Gasoline ***	0.99	(0.18)	0.09	0.11	0.18	203	1.15	(0.11)	6.74***	0.20	0.20	205						
<i>Precious Metals</i>																		
Gold	-1.76***	(0.37)	51.44***	-0.10	0.07	238	-1.63***	(0.21)	91.44***	-0.04	0.09	238						
Silver	-1.83***	(0.31)	79.35***	-0.02	0.12	238	-1.21***	(0.20)	170.73***	-0.06	0.16	238						
<i>Base Metals</i>																		
Aluminum	0.84	(0.53)	1.22	-0.11	0.12	145	1.66**	(0.31)	3.12**	-0.09	0.19	142	0.49**	(0.20)	10.64***	-0.16	0.25	136
Copper	-0.26***	(0.32)	15.75***	-0.06	0.18	145	-1.19***	(0.26)	44.79***	-0.04	0.28	142	-1.11***	(0.16)	152.27***	0.06	0.32	136
Lead	0.37**	(0.25)	3.21**	-0.10	0.19	145	0.38***	(0.14)	14.95***	-0.10	0.28	142	1.11	(0.21)	59.68***	-0.37	0.45	133
Nickel	1.16	(0.69)	2.97*	-0.05	0.21	145	4.32***	(0.38)	39.11***	-0.23	0.37	142	2.91***	(0.27)	73.21***	-0.11	0.52	136
Tin	1.43	(0.69)	7115***	-0.05	0.14	145	2.21***	(0.45)	4.55**	-0.09	0.23	138	1.12	(0.32)	11.71***	-0.18	0.34	133
<i>Agricultural Products</i>																		
Corn	0.87	(0.14)	30.61***	0.11	0.14	238	0.71***	(0.09)	67.90***	0.10	0.20	238						
Soybean	1.27***	(0.09)	6.63***	0.10	0.12	238	0.91	(0.12)	5.47***	0.12	0.17	238						
Wheat	0.63**	(0.18)	60.70***	-0.06	0.14	155	0.51***	(0.12)	71.98***	-0.09	0.20	155						

Note: The table presents estimated results by GARCH(p,q) of equation (2) in the text for different commodities and futures prices horizons. For oil, corn, soybean, and wheat, the spot price is replaced with the 1-month futures price. For gasoline, only the old version (“HU”) of futures are used in the regressions. Statistical significance at the 10%, 5%, and 1% level are denoted by *, **, and *** respectively. For β , the null is that $\beta=1$. p and q are selected via the AIC for each specification. Newey-West HAC standard errors, using truncation equal to the futures horizon minus one.

Table 3: Out-of-Sample Forecasting Performance of Commodity Futures relative to Random Walk and Econometric Forecasts:

	Relative Root Mean Squared Forecast Error						Relative Mean of Absolute Value of Forecast Errors					
	Futures			ARIMA			Futures			ARIMA		
	3-mo	6-mo	12-mo	3-mo	6-mo	12-mo	3-mo	6-mo	12-mo	3-mo	6-mo	12-mo
<i>Energy Products</i>												
Oil	0.99	1.00	1.06	1.28***	1.24**	1.14	0.98	1.01	1.13	1.22**	1.17	1.11
Natural Gas	1.02	0.95	0.91	1.23	1.18	1.13	1.02	0.97	0.94	1.23	1.21	1.18
Heating Oil	0.99	1.02	1.03	1.26***	1.23*	1.15	0.99	0.99	1.08	1.25**	1.14	1.13
Gasoline ***	0.87	0.86		1.04	1.08		0.88	0.91		1.07	1.06	
<i>Precious Metals</i>												
Gold	0.94	0.87	0.75	1.21*	1.12	0.96	0.95	0.88	0.73	1.18	1.13	0.97
Silver	0.99	0.97	0.95	1.20**	1.15	0.99	0.98	0.94	0.94	1.17	1.12	0.95
<i>Base Metals</i>												
Aluminum	1.01	1.01	1.07	1.29***	1.33***	1.03	1.00	0.98	1.08	1.30***	1.25**	1.07
Copper	1.01	1.03	1.12	1.32***	1.30***	1.07	1.04	1.06	1.14	1.36***	1.25**	1.06
Lead	1.01	1.02	1.04	1.24***	1.21*	1.31**	1.02	1.04	1.09	1.30***	1.15	1.22
Nickel	1.01	1.02	1.03	1.23**	1.21*	1.18	1.02	1.01	1.03	1.22**	1.19	1.18
Tin	1.00	1.00	1.00	1.31***	1.29***	1.29**	0.99	0.98	0.98	1.28***	1.22**	1.27**
<i>Agricultural Products</i>												
Corn	0.98	0.95		1.34***	1.19		0.97	0.99		1.37***	1.23*	
Soybean	0.91	0.89		1.23**	1.23*		0.90	0.91		1.24**	1.21	
Wheat	0.98	1.04		1.23	1.18		0.99	1.01		1.23	1.21	

Note: Table displays the root mean squared forecast error and mean of absolute value of forecast errors at each forecast horizon and commodities product for two forecasting approaches – using futures prices and using an ARIMA model– *relative* to the relevant forecast error measure from a random walk (without drift) prediction. Out-of-sample forecasts are evaluated over 2003M1 to 2008M7. The *, **, and *** denote whether the p -value of the two-sided test of the null that the forecast error measure was generated by a random walk process is less than 10%, 5%, and 1% respectively. p -values are calculated by simulating random walk processes with same variance as in each commodity market and generating a distribution of RMSEs and MAEs for each commodity at each forecast horizon.

Figure 1: Price of Petroleum (WTI), end of month, and the log of 3 and 6 month basis.




Figure 2: Price of Natural Gas (Henry Hub), end of month, and the log of 3 and 6 month basis.




Figure 3: Price of Heating Oil, end of month, and the log of 3 and 6 month basis.




Figure 4: Price of Gasoline (NY Harbor), end of month, and the log of 3 and 6 month basis.




Figure 5: Price of Gold, end of month, and the log of 3 and 6 month basis.




Figure 6: Price of Silver, end of month, and the log of 3 and 6 month basis.




Figure 7: Price of Aluminum, end of month, and the log of 3 and 6 month basis.




Figure 8: Price of Copper, end of month, and the log of 3 and 6 month basis.




Figure 9: Price of Lead, end of month, and the log of 3 and 6 month basis.




Figure 10: Price of Nickel, end of month, and the log of 3 and 6 month basis.




Figure 11: Price of Tin, end of month, and the log of 3 and 6 month basis.




Figure 12: Price of Corn, end of month, and the log of 3 and 6 month basis.




Figure 13: Price of Soybeans, end of month, and the log of 3 and 6 month basis.




Figure 14: Price of Wheat, end of month, and the log of 3 and 6 month basis.






Figure 15: Relationship between Market Depth and the Unbiasedness Hypothesis




Note: The figure plots the average monthly trading volume, in logs, of 3-month and 6-month futures for each commodity between 2006 and 2009, as well as the average of the absolute value of the t-statistics for the null of unbiasedness for 3-month and 6-month futures for each commodity from the empirical estimates of Table 1.

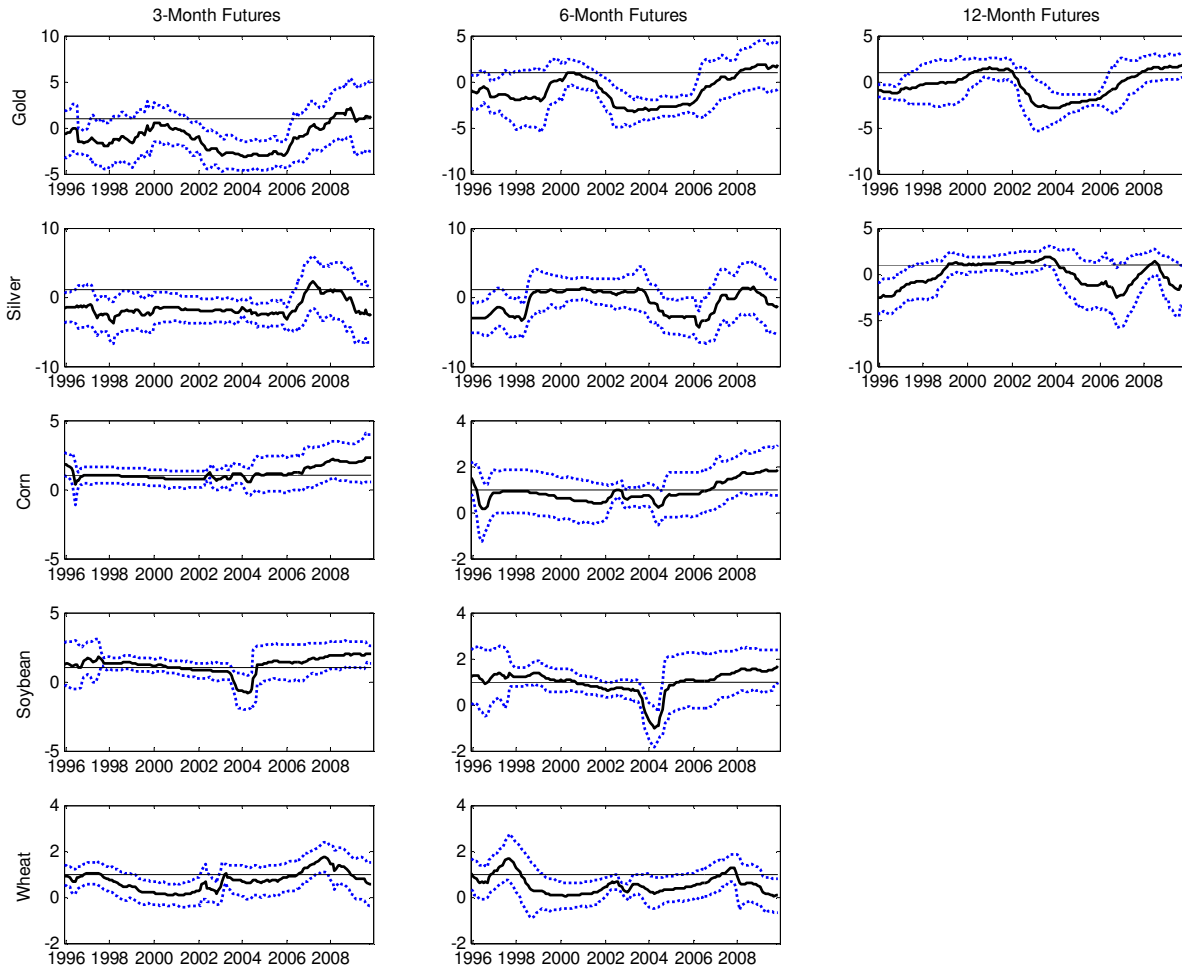
Figure 16: Rolling Estimates of Basis Equation for each Commodity and Futures Horizon



Panel B: Base Metals



Panel C: Precious Metals and Agricultural Commodities



Note: The three panels plot 5-year rolling OLS estimates of β from equation (2) for each commodity and futures horizon, as well as 2-standard deviation confidence intervals (dotted blue lines). The horizontal line is for the null of unbiasedness ($\beta = 1$). The date associated with each point is the last observation of each sample period considered.