Fractional Cointegration and Price Discovery in Canadian Commodities

Ke Xu* Kenneth G. Stewart Zeyang Cao

Department of Economics
University of Victoria
Victoria, British Columbia
Canada V8W 2Y2

September 2020

Abstract

The fundamental value of commodities, characterized as the common stochastic trend shared by cointegrated spot and futures prices, is studied using recently developed tools: the fractionally cointegrated vector autoregression model of Johansen and Neilsen, the permanent-transitory decomposition of Gonzalo and Granger, and a distributional result for forecast evaluation due to Clark and West. For ten Canadian agricultural and mining commodities, we find that the generalization to fractional cointegration is statistically significant. However the economic significance of this generalization—in terms of forecast accuracy and the profitability of mean-variance dynamic trading strategies—is more fragile than may have been appreciated.

Keywords: Canadian commodities, price discovery, fractional cointegration, futures markets, vector error correction model

JEL classifications: C32, G11

Funding

Ke Xu acknowledges financial support from the Social Sciences and Humanities Research Council of Canada (SSHRC) Insight Development Grants (430-2018-00557).

Declarations of Interest: None

* Corresponding author: Ke Xu, kexu@uvic.ca

Among the most intuitive and longstanding notions about financial markets is that assets have, in some sense, a fundamental value associated with long run equilibrium. Observed prices may depart temporarily from this underlying value owing to short run transitory forces.

The role of markets in revealing or "discovering" this fundamental value is called *price discovery* in the asset pricing literature, and has been studied for assets as diverse as equities and foreign exchange. Of particular interest is assets having both a spot or cash market and a derivatives market, so that the fundamental value is jointly determined by both markets. A natural question then becomes: what are the relative contributions of the spot and derivatives markets to price discovery?

1 Introduction

This paper uses recent econometric developments to contribute to the literature studying this question for commodity markets, the original home of derivative instruments—namely futures contracts, which have long traded simultaneous with commodity spot markets. We compare the standard modelling framework—the cointegrated vector autoregressive (CVAR) model—with its generalization, the fractionally cointegrated vector autoregressive (FCVAR) model of Johansen and Nielsen (2012, 2018), in studying price discovery using the permanent-transitory decomposition of Gonzalo and Granger (1995). The methodology follows Figuerola-Ferretti and Gonzalo (2010) and Dolatabadi, Narayan, Nielsen, and Xu (2018). Whereas many studies have investigated price discovery using the CVAR model, few have generalized this to permit the long memory that fractional cointegration allows, a generalization that our results tend to endorse.

Several related analyses follow as natural adjuncts. First, as one dimension on which the statistical success of the CVAR and FCVAR can be assessed, their ability to forecast spot and futures returns is compared. Second, their economic significance is evaluated by comparing the profitability of trading strategies based on the two models. Portfolio returns, optimal holdings, and Sharpe ratios are derived from a mean-variance utility framework. Being a generalization of the CVAR, in principle the FCVAR is superior in the behavior it permits. Nevertheless, we find that the CVAR outperforms the FCVAR in some cases.

To some extent our results serve as robustness checks on earlier findings, both with respect to the sample period and the choice of commodities. Whereas the literature tends to focus on the most heavily traded commodities in European and U.S. markets, we choose as our unifying theme agricultural and mining commodities that are major products in Canada, in terms of both production capacity and trading volumes. Canada is a global leader in mineral production and ranks among the top five global producers of gold, nickel, potash, uranium, diamonds, and platinum. The mineral sector contributes significantly to the Canadian economy, accounting for 19% of exports and 5% of GDP (Mines Canada, 2019). Similarly, agricultural commodities also play a crucial role in the Canadian economy. According to Farm Credit Canada (2018), oilseeds (e.g. soybeans and canola), cereals (e.g. wheat and oats), and meats (e.g. pork and beef) constitute 41% of the total value of Canadian exports of agricultural commodities and food products. In 2018 the agricultural sector was Canada's third-largest export category, accounting for 11% of exports and roughly 6.8% of GDP.

Ten commodities are included in our empirical analysis. We find that either the spot or futures market can dominate price discovery, depending on the commodity. Across the commodities, the CVAR and FCVAR models yield fairly consistent stories about the relative roles of the two markets. In terms of the profitability of trading strategies based on the forecasts of these models, a couple of possibly surprising findings emerge. First, portfolio returns from the CVAR model are slightly higher on average than those from the FCVAR. And second, this result is insensitive to the level of investor risk aversion.

2 Econometric Methodology

Much empirical evidence supports modelling spot and futures prices as cointegrated I(d) processes, where the univariate order of integration of the individual series is $d = 1$. The standard framework for studying their joint evolution has therefore been the CVAR model; Figuerola-Ferretti and Gonzalo (2010; henceforth FFG) is a prominent example. But although, in their univariate behavior, a unitary order of integration $d = 1$ for spot and futures prices is the accepted specification, the nature of their joint evolution is less obvious. It may be characterized by the long memory that can be modelled with fractional cointegration, as in Dolatabadi, Nielsen, and Xu (2015,2016) and Dolatabadi, Nielsen, Narayan and Xu (2018). That is, the order of cointegration b, where b can be fractional, may differ from the shared univariate order of integration $d = 1$.

2.1 The FCVAR and CVAR Models

These distinctions can be captured with the FCVAR model of Johansen and Nielsen (2012, 2018), a generalization of the CVAR:

$$\Delta^d X_t = \alpha\beta'\Delta^{d-b}L_b X_t + \sum_{i=1}^{k}\Gamma_i\Delta^d L_b^i X_t + \varepsilon_t. \tag{1}$$

Here $X_t = [s_t, f_t]$ is the vector of (the logs of) spot and futures prices, L is the lag (or back-shift) operator, and $\Delta \equiv 1 - L$ is the first-difference operator. The symbol L_b denotes $1 - \Delta^b$ which, for fractional b, is defined in the technical literature on fractional integration (Johansen and Nielsen; 2012, 2018). In terms of the mechanics of working with this notation, if $b = 1$ then $L_1 = 1 - \Delta^1 = 1 - (1 - L) = L$. In the augmenting lags, the symbol L_b^i means $(L_b)^i = (1 - \Delta^b)^i$.

The FCVAR nests within it the CVAR, the special case in which $b = d = 1$. For in this case

$$\Delta^{d-b}L_b = \Delta^{1-1}(1 - \Delta^1) = 1 - (1 - L) = L$$

and

$$\Delta^d L_b^i = \Delta^1(L_1)^i = \Delta L^i.$$

The FCVAR model (1) therefore reduces to the standard expression for the CVAR (for example, equ. (25) of FFG):

$$\Delta X_t = \alpha\beta' X_{t-1} + \sum_{i=1}^{k}\Gamma_i\Delta X_{t-i} + \varepsilon_t. \tag{2}$$

Hence, with the exception of permitting distinct integration orders b and d, much of the notation of the FCVAR model is as in the CVAR. In particular, k is the number of lags needed to treat short term dynamics so that the disturbance ε_t is white noise.

With respect to the forces of long run equilibrium, in general β would be a matrix of cointegrating vectors. In the special case of just two variables—in the present context the spot and futures prices $X_t' = [s_t, f_t]$—and a cointegrating rank of $r = 1$ long run relationship, β is a vector. Allowing for an intercept ρ, let the normalized cointegrating relationship be

$$s_t = \beta_2 f_t + \rho,$$

Then the cointegrating vector is $\beta' = [1, -\beta_2]$ and the error correction term in the CVAR model (2) is (generalized trivially to permit the intercept ρ)

$$s_t - \beta_2 f_t - \rho = \beta' X_t - \rho. \tag{3}$$

The CVAR model assumes that a long run equilibrium relationship between the I(1) variables s_t and f_t takes the form of this linear combination being I(0). In contrast, the FCVAR model (1) with $d = 1$ permits this equilibrium error to be $I(1-b)$. For b in the range $(0, 1/2]$ we have $1/2 \leq 1 - b < 1$, and the nature of this fractional cointegration is that the equilibrium error is nonstationary although mean reverting. For b in the range $(1/2, 1]$ we have $0 \leq 1-b < 1/2$ and the equilibrium error is stationary but with the long memory permitted by fractional integration. As $b \to 1$ so that $1 - b \to 0$ this long memory disappears, in the limit yielding an equilibrium error that is I(0)—the special case of the CVAR model.

Just as in the CVAR model, α is a vector of speed-of-adjustment coefficients. For a two-equation model describing $X_t' = [s_t, f_t]$ this is $\alpha' = [\alpha_1, \alpha_2]$, and these coefficients have the usual interpretations. If the equilibrium error (3) is positive then s_t should adjust downward (the first equation of the system should yield $\Delta s_t < 0$) while f_t should adjust upward (the second equation of the system should yield $\Delta f_t > 0$). These directions of adjustment for the dependent variable ΔX_t in response to previous-period disequilibrium imply that, in the FCVAR model just as in the CVAR, we should find $\alpha_1 < 0$ and $\alpha_2 > 0$.

At its most elementary level, the basis for the hypothesized cointegrating relationship (3) is the common sense observation that spot and futures prices track one another, as is revealed by any plot of the series such as those of Figure 1. The more rigorous formulation of this intuition is modern statements of the classic Kaldor-Working theory of storage, such as that of Garbade and Silber (1983). FFG's contribution was to formulate this theory in the CVAR framework, within which the Gonzalo-Granger decomposition can be brought to bear; the cointegrating relation (3) corresponds to FFG's equation (7), the implication of spot-futures parity. In this model the intercept ρ captures factors such as storage costs, convenience yield, and cost to carry. Because convenience yield can be positive or negative, there is no a priori testable restriction on ρ. The restriction $\beta_2 > 1$ implies a long run equilibrium relationship characterized by backwardation ($s_t > f_t$), while $\beta_2 < 1$ implies long run contango ($s_t < f_t$). Of course, because the futures price is predominantly a forecast of the future spot price, either of these can hold at various points in time, depending on market conditions and expectations about the future. That is, although the parameters β_2 and ρ have interesting economic interpretations, these do not take the form of testable restrictions the rejection of which would falsify the theory.

2.2 Price Discovery

Cointegration between spot and futures prices means that their long run evolution is driven by a single stochastic trend, a stochastic trend that is the joint outcome of trading in the spot and futures markets. What are the relative contributions of these two interconnected markets to price discovery, the fundamental value represented this common stochastic trend?

Because the speed-of-adjustment vector $\alpha' = [\alpha_1, \alpha_2]$ plays the same role in the FCVAR as it does in the CVAR, the permanent-transitory decomposition of Gonzalo and Granger (1995) can be used to investigate this question. To illustrate with the simplest case, consider a two-equation CVAR with one cointegrating relationship. Gonzalo and Granger showed that each variable can be decomposed into permanent and transitory components. The permanent component is the common stochastic trend while, remarkably, the transitory variation of each variable is its individual response to the equilibrium error. These individual responses are governed by the speed-of-adjustment vector α.

The permanent component, the common stochastic trend, is interpreted as the commodity's underlying fundamental value that markets seek to discover. For the price vector specified as $X_t' = [s_t, f_t]$, the first equation of the system describes the evolution of the spot price s_t while the second describes the futures price f_t. Consider the extreme case in which the speed-of-adjustment vector is $\alpha' = [1, 0]$: $\alpha_2 = 0$ means that the futures price does not respond to transitory departures from equilibrium, and instead is driven entirely by its permanent component, the common stochastic trend. Hence the long run evolution of prices is determined in the futures market. It is the spot market that, with non-zero α_1, responds to the short run disequilibrium captured by the equilibrium error, contributing transitory variation to the long run evolution of prices. This transitory variation could be due to factors such as, in the words of FFG (p. 158), "…bid-ask bounce, temporary order imbalances or inventory adjustments."

Notice that, in addition to α having its usual speed-of-adjustment interpretation, according to the Gonzalo-Granger decomposition α also captures the importance of transitory factors in determining each price. It follows that its orthogonal complement $\alpha_\perp$, defined by $\alpha_\perp'\alpha = \alpha'\alpha_\perp = 0$, describes the opposite—the extent to which the permanent component influences each price; equivalently, the extent to which each price reveals, or discovers, the fundamental value. Because $\alpha_\perp$ is unique only up to a multiplicative constant, it is natural to normalize its components to sum to one, so that they measure the relative shares in price discovery. In the extreme example where $\alpha' = [1, 0]$, this is $\alpha_\perp' = [0, 1]$, indicating that all price discovery takes place in the futures market, none in the spot market.

All this was recognized and applied by FFG in using a CVAR to model commodity markets. Dolatabadi, Narayan, Nielsen, and Xu (2018) recognized that, because the role of α is the same in the FCVAR model, the approach extends naturally to fractional cointegration. It is this set of tools that we apply.

2.3 Relationship to Earlier Literature

The recognition that the common stochastic trend is a natural measure of fundamental value, and that the Gonzalo-Granger decomposition can be applied to obtain the relative contributions of the markets to the discovery of that fundamental value, is a seminal insight. It provides a remarkable example of how new methodological tools can not only advance our ability to study empirical questions rigorously, but even alter our conception of what those questions are and how they should be posed.

The historic approach to analyzing price discovery was cast in terms of lead-lag relationships between spot and futures prices. If, for example, futures prices were found to lead spot prices, the inference was that new information is embodied first in futures prices and that they therefore contribute most to price discovery. For example, citing literature going as far back as 1932, Garbade and Silber (1983, p. 289) remark that "The essence of the price discovery function of futures markets hinges on whether new information is reflected first in changed futures prices or in changed cash prices."

Because futures prices embody expectations of future spot prices, in many markets futures prices do indeed tend to lead spot prices. This is particularly so in commodity markets, where the cost of futures trading is lower than for trading in the spot market. The conventional wisdom has therefore been that, especially for commodities, it is futures markets that dominate price discovery.

This conception of price discovery gained traction with the widespread application of Granger causality testing within CVAR systems which, of course, is really about timing relationships. One example is Silvapulle and Moosa (1999), who studied the crude oil market and found that linear causality testing reveals that futures prices lead spot prices, but nonlinear causality testing reveals a bidirectional effect. They conclude that spot and futures crude oil prices react simultaneously to new information, but that the futures market plays the dominant role in price discovery.

This tradition continues today as one stream of the empirical literature. An example is Peri, Baldi, and Vandone (2013), who study the U.S. corn and soybeans markets and conclude (p. 402) that "...futures prices play a major role in price discovery." The basis for this conclusion is (p. 398) "The study of the causal relationship between spot and futures prices ..." and "...the analysis of the 'price discovery' role of spot and futures markets, defined as the lead-lag relationship and information flows between spot and futures markets."

However there are a couple of weaknesses with this approach. First, it yields no precise quantitative measure of the relative roles in price discovery comparable to Gonzalo and Granger's $\alpha_\perp$. Second, and more fundamentally, it is a commonplace observation about Granger causality that, because it is defined in terms of timing relationships, these can be misleading as indicators of true causality—an example of the fallacy of *post hoc ergo propter hoc.* The belief that timing relationships reveal price discovery may amount to essentially the same fallacy.

The Gonzalo-Granger decomposition makes it possible to circumvent these ambiguities by recasting price discovery, not in terms of timing relationships, but in terms of the common stochastic trend shared by cointegrated spot and futures prices.

3 Data and Estimation

We study ten commodities that are important to the Canadian economy and heavily traded enough that daily spot and futures data are available for substantial sample periods. These include five agricultural commodities—soybeans, wheat, oats, live cattle, and lean hogs—and five extracted commodities—platinum, iron ore, nickel, crude oil, and gold. We use daily nearest contracts from the Commodity Research Bureau (accessed via www.barchart.com). For eight commodities the sample begins 1 January 2009; the exceptions are iron ore (19 February 2013) and nickel (16 December 2011). In all cases the sample period ends 1 September 2019. In addition to these spot and futures price series, we follow the standard practice of using the three-month Government of Canada treasury bill rate, downloaded from Statistics Canada, as the risk-free rate of return that would be available to a Canadian investor.

3.1 Data Description

Table 1 summarizes some key features of our futures data, including the Barchart identifier symbols, the exchanges on which the contracts are traded, and the sample start dates. Average daily volume and open interest are also shown, first for each commodity's full sample and then for a common subsample, the end-of-sample months of 2019. *Volume* is the average daily number of contract trades, and varies dramatically across the commodities: futures volume for the most heavily traded commodities, such as crude oil, gold, and soybeans, is many times that of iron ore, oats, or nickel. This variation is paralleled in *open interest,* the average number of contracts outstanding. A large open interest indicates active trading, and so is a useful measure of market liquidity.

Turning to a comparison of spot and futures prices, Figure 1 provides time plots by commodity and Table 2 reports descriptive statistics. At the level of casual inspection, the plots reveal two compelling conjectures. The first is that, typical of prices in financial markets, many of these spot and futures prices do not obviously exhibit trend-reverting behavior, and so are likely driven by stochastic trends. That is, the natural null hypothesis is that they are integrated processes, although in some cases possibly zero-drift ones. This intuition is supported by the Dickey-Fuller tests in the final column of Table 2 which, for both spot and futures prices and across all commodities, do not reject the unit root null at anything approaching conventional significance levels.

Given the acceptance of univariate I(1) specifications, the second compelling conjecture suggested by Figure 1 is that spot and futures prices are cointegrated: they clearly move together, to the point of being indistinguishable in some of the graphs. (Spot prices are plotted in blue, futures in red.) This cointegration is, of course, the basis for the CVAR model (2) as the standard framework for analysis. We therefore turn to the application of this model and its generalization, the FCVAR.

3.2 Estimation Results

Table 3 reports estimation results for the CVAR and FCVAR models, which are obtained with the Matlab package of Nielsen and Popiel (2018). For each commodity and model, an appropriate lag length k is selected by a combination of testing and Schwarz's Bayesian information criterion, using Ljung-Box Q tests as a diagnostic to verify the absence of residual autocorrelation. Whereas Dolatabadi et al. (2018) tended to find that the FCVAR required fewer lags, and attributed this to the serial dependence permitted by a non-unitary b, we did not particularly find this to be the case.

In the FCVAR model, the fractional parameter b is estimated to range between 0.591 (lean hogs) and 0.931 (iron ore). With the exception of iron ore, these estimates differ from unity to an extent that is statistically significant, and so on this basis the CVAR model is rejected. With respect to the point estimates, the conventional intuition is that a higher value of b implies less memory in the error correction term and hence better market efficiency. Yet some of our estimates are not entirely consistent with this intuition. Iron ore, the least-traded commodity, yields the highest estimate of b at 0.931, while crude oil, the most heavily traded, has one of the lower estimates at 0.669. Even so, the latter is more consistent with the other estimates than the extraordinarily low value of 0.194 found by Dolatabadi et al. (2018, Table 3) for crude oil, which was for the period March 1983–October 2012.

Although the CVAR model is rejected relative to the FCVAR, note that the estimates of the cointegration coefficient β_2 are, for any given commodity, usually very similar across the two models, oats being the main exception. This is a manifestation of the asymptotic (but not finite-sample) equivalence noted by Lasak (2008, p. 9) of the maximum likelihood estimator of the cointegrating vector β in the two models. The cointegration coefficient β_2 indicates whether the futures market is systematically in a state of backwardation ($-\beta_2 > 1$) or contango ($-\beta_2 < 1$). For both the CVAR and FCVAR models, many of the estimates differ little from the razor's edge case of $-\beta_2 = 1$ in which neither backwardation nor contango dominates in the long run. The main exception is wheat, where both models yield $-\hat{\beta}_2 \approx 1.5$ and so suggest backwardation. The one commodity where the two models yield quite different point estimates is oats: the CVAR model suggests backwardation ($-\hat{\beta}_2 =$

1.230) while the FCVAR suggests contango ($-\hat{\beta}_2 = 0.779$).

Instead of being manifested in the cointegration coefficient, the statistically significant difference between the CVAR and FCVAR models mainly shows up in the adjustment coefficients and the associated price discovery shares. Even so, the estimates are generally consistent with short run adjustment of prices to bring their markets into long run equilibrium. Both the CVAR and FCVAR models yield estimates of the speed-of-adjustment coefficients that are predominantly negative for the spot price ($\hat{\alpha}_1 < 0$) and positive for the futures price ($\hat{\alpha}_2 > 0$). Most of the coefficients violating this pattern are very close to zero, and so their signs can easily be the outcome of sampling error. To the extent that any importance is attached to unintuitive signs, however, the FCVAR model has more than the CVAR.

Finally, estimates of the price discovery shares $\alpha_{\perp,1}$ and $\alpha_{\perp,2}$ are shown in the final two columns of Table 3. These are only loosely "shares" in that, although their values are normalized by requiring that they sum to unity, they are not constrained to be positive fractions, and indeed some lie outside the range [0,1]. Nevertheless their values indicate the relative importance of the two markets in price discovery. In this respect CVAR and FCVAR models are fairly consistent, finding the spot market to be dominant for wheat, platinum, iron ore, and gold, while the futures market is dominant for soybeans, live cattle, lean hogs, and nickel. The two models differ for only two commodities: for oats the CVAR model indicates that the spot market dominates, while the FCVAR suggests it is the futures market. The reverse is the case for crude oil.

A comparison with other results in the literature shows that which market dominates price discovery can change over time. A good example is soybeans which, according to both the CVAR and FCVAR models, are dominated by the futures market, whereas Dolatabadi et al. (2018) find it to be the spot market.

Of course, nothing requires that price discovery be dominated by one market; both the spot and futures markets can contribute. But when, in our results, one market is dominant according to its coefficient point estimate, that coefficient is generally strongly statistically significant.

4 Forecasting

Although our CVAR and FCVAR models yield sensible within-sample inferences, a rather different criterion by which they can be judged is their forecasting ability. Here we consider two dimensions of this: out-of-sample forecast accuracy, and the ability of our models to yield portfolio profits in a mean-variance utility framework. The latter serves to gauge forecasting success in terms of economic significance rather than purely statistical significance.

4.1 Out-of-Sample Forecasting

To compare out-of-sample forecasting ability we re-estimate both the CVAR and FCVAR models using, for each commodity, the first 75 percent of the sample. The remaining 25 percent of the sample is then forecasted recursively, each day updating the sample on which the forecast is based with the most recent day's observation. These forecasts are then compared with the observed sample.

In asset pricing a common statistic for evaluating predicted returns is the out-of-sample R^2 (see, for example, equ. (1) of Campbell and Thompson (2008)) which, unlike a conventional regression R^2, can be negative. A positive (negative) out-of-sample R^2 indicates that the model forecasts have a smaller (larger) average mean square prediction error than the historical average return.

The first four columns of Table 4 report these out-of-sample R^2's and show that the models sometimes underperform historical returns: the out-of-sample R^2 is negative. For both the CVAR and FCVAR models, the spot prices for lean hogs and nickel are by far the best forecasted. Overall, the forecasting performance of the two models is remarkably comparable, the out-of-sample R^2's of the two being similar. And spot and futures prices seem to be about equally difficult to forecast, in that there is little systematic tendency for either model to consistently predict spot prices better than futures prices, or the reverse.

Although the out-of-sample R^2's suggest that the forecasting performance of the CVAR and FCVAR models is broadly similar, a more rigorous comparison is possible. For this purpose Table 4 reports two statistics, the relative root mean square error (RMSE) and the CW statistic of Clark and West (2007). The relative RMSE is defined so that a negative value favors the more general model—in our case, the FCVAR. And indeed this is the pattern for the majority of commodities in both the spot and futures markets.

But to what extent is this statistically significant? The advantage of the CW statistic is that, unlike the other statistics of Table 4, it has a determinable sampling distribution, which Clark and West (2007) showed to be asymptotically standard normal, an asymptotic result that should easily hold for our sample sizes. The CW statistic takes the null hypothesis to be that the more general model (in this case the FCVAR) yields forecasts that are no better than the special-case model (the CVAR). The alternative hypothesis is that the more general model is better. So a large CW statistic favors the more general model. Consistent with our observations about the R^2's, Table 4 shows that this null is rejected at conventional significance levels in only a few instances: the spot prices for soybeans, oats, lean hogs, and nickel, and the futures prices for lean hogs and iron ore.

On balance, to the extent that the evidence favors the forecasting performance of one of our models, it is the FCVAR. But it is hardly surprising that the more general model does somewhat better; perhaps the greater surprise is how modest the improvement seems to be.

4.2 Profits in a Mean-Variance Utility Framework

Can the forecasts of our CVAR or FCVAR models be used to trade profitably? We now consider portfolios constructed according to mean-variance utility, allowing investors to trade dynamically in the sense that portfolios can be rebalanced daily. Daily rebalancing is most likely to reveal profitable trading opportunities, because forecast accuracy deteriorates at longer forecast horizons.

Applying the standard framework in the literature (for example, Marquering and Verbeek (2004) or Dolatabadi et al. (2018)) investors are treated as holding two assets, one risky and the other risk-free. The log returns between periods t and $t+1$ are denoted r_{t+1} and $r_{f,t+1}$, respectively. In general the risky return r_{t+1} could be for either a spot or futures position. However commodity spot markets tend to be dominated by producers and users of the physical commodity. It is futures markets that provide the forum for a broader range of traders to participate, enabling risk transfer. Hence futures markets have lower transactions costs and are the principal vehicle by which financial investors trade. This explains why it is futures markets that are often found to be the dominant contributor to price discovery.

For these reasons, we assume that investors who are trading dynamically with daily rebalancing are using the futures market, and the risky return r_{t+1} is on a futures contract. Let w_t denote the proportion of the portfolio invested in this futures position, and let θ denote the trading cost of altering this portfolio weight. The transactions cost of daily portfolio rebalancing is therefore $\theta|w_{t+1} - w_t|$ and the return on the portfolio is

$$r_{p,t+1} = w_{t+1}r_{t+1} + (1 - w_{t+1})r_{f,t+1} - \theta|w_{t+1} - w_t|. \tag{4}$$

Investors make their portfolio decisions w_{t+1} by maximizing mean-variance utility. The standard utility function used in this literature is

$$U(r_{p,t+1}) = \mathrm{E}_t(r_{p,t+1}) - \frac{1}{2}\gamma \mathrm{Var}_t(r_{p,t+1}). \tag{5}$$

The parameter γ captures risk aversion, which is increasing in γ. One appeal of this utility specification is that utility maximization yields a solution for the optimal portfolio weight that is elegant in its simplicity:

$$w^*_{t+1} = \frac{\mathrm{E}_t(r_{t+1}) - r_{f,t+1}}{\gamma \mathrm{Var}_t(r_{t+1})}. \tag{6}$$

Intuitively, the portfolio share invested in the risky asset, a futures contract, is directly related to its expected return but inversely related to both the return variance and the degree of risk aversion. Conveniently, w^*_{t+1} is unaffected by the transactions cost θ, which enters the calculations only via the portfolio return (4).

The calculation of trading profits proceeds by first using the estimated CVAR or FCVAR models to forecast one-day-ahead futures returns r_{t+1}. Using this forecasted return, the optimal portfolio weight (6) is calculated. This requires several additional inputs: a risk-free rate of return $r_{f,t+1}$, measured by the treasury bill yield; an estimate of the time-varying variance $\mathrm{Var}_t(r_{t+1})$, for which we follow the standard practice of using a GARCH(1,1) model; and the risk aversion parameter γ. For the latter we follow Dolatabadi et al. (2018) in using γ values of 3, 6, and 12 to gauge the sensitivity of the results to different degrees of risk aversion. In addition to these informational inputs, the optimal weights are constrained to lie in the range $-0.5 \leq w^*_{t+1} \leq 1.5$, which limits short selling and borrowing/leverage to 50% of the position.

How profitable ex post is the ex ante trading strategy embodied in w^*_{t+1}? Applying w^*_{t+1} to the *realized* futures return r_{t+1}, the actual portfolio return (4) is calculated. Only at this step is the transactions cost θ needed. We adopt the value used by Dolatabadi et al. (2018) based on their reading of the literature, $\theta = 0.000167$ (that is, 0.0167% of the nominal value of the re-weighting), which reflects the comparatively low cost of futures trading.

The results are shown in Table 5, reported as average annualized percentage excess returns, with standard errors. Several regularities are evident. First, most are statistically insignificant, including all that are negative: for most commodities, neither the CVAR nor the FCVAR model does dependably better than investing at the risk-free rate. But there are exceptions: live cattle, lean hogs, nickel, and, to a lesser extent, iron ore all yield significantly positive returns, regardless of model or level of risk aversion. For only two commodities is there any substantive difference between the CVAR and FCVAR models: for wheat the FCVAR model yields a significant return, while for oats it is the CVAR model. But with this exception, perhaps the most striking aspect of Table 5 is the insensitivity of the results across both models and levels of risk aversion; none of these conclusions is conditional on the risk aversion parameter γ. And there is no evidence to support the conjecture that incorporating fractional cointegration yields more profitable trading strategies. On the contrary, the final row of Table 5 shows that, on average across the commodities, the CVAR yields higher returns than the FCVAR.

What is it about the nature of these commodities that some have profitable trading strategies, others not? A comparison with Table 1 shows that the most profitable strategies tend to be associated with low trading volumes (oats, live cattle, lean hogs, iron ore, nickel), while the absence of profitable strategies is associated with high volumes (soybeans, crude oil, gold). The common-sense inference is that it is the relatively illiquid markets that offer profitable trading opportunities; the more liquid the market, the fewer such opportunities.

These results are to a considerable extent consistent with previous research, but not entirely. Dolatabadi et al. (2018) studied seventeen commodities, some the same as ours, over the period March 1983 through October 2012. Like us, they found little difference between the CVAR and FCVAR models, and dramatic differences in profitability across commodities. They too obtained largely insignificant futures returns for soybeans, wheat, and crude oil; but, unlike us, they found significant returns for the precious metals gold and platinum.

Sharpe Ratios

Are any of these findings about profitability altered if risk is controlled for? A simple way of examining this is to calculate Sharpe ratios, the ratios of the excess returns to their standard errors. Table 6 reports these, calculated from the information underlying Table 5. Given that Table 5 revealed little variation across degrees of risk aversion, Table 6 focuses on the intermediate case of $\gamma = 6$. It shows that, for the most part, the essential conclusions from

Table 5 are unchanged. Half the Sharpe ratios are statistically insignificant, including all that are negative. The results for the CVAR and FCVAR models are remarkably similar, both yielding insignificant Sharpe ratios for soybeans, platinum, crude oil, and gold, yet significant ones for live cattle, lean hogs, nickel, and, to a lesser extent, iron ore. The FCVAR model yields modestly significant Sharpe ratios for wheat, the CVAR model for oats. This is essentially the same pattern of statistical significance as in Table 5. And as before, incorporating fractional cointegration does not yield a more profitable risk-adjusted trading strategy: the final line of Table 6 shows that, on average across the commodities, the CVAR model actually yields a slightly higher Sharpe ratio than the FCVAR.

5 Conclusions

The traditional approach to price discovery has been cast in terms of timing relationships. In commodity markets, spot markets tend to be dominated by producers and consumers of the physical commodity. Instead it is futures markets that are the vehicle by which the much larger class of financial investors participate. Futures prices embody expectations of future spot prices. Hence the conventional wisdom about commodity markets is that new information is reflected most quickly in futures prices, which therefore tend to lead spot prices. The typical finding has therefore been that the fundamental value of commodities is primarily discovered in futures rather than spot markets.

This paper has explored the sensitivity of this conventional wisdom to recently developed tools for studying price discovery, in particular the Gonzalo-Granger decomposition as implemented by Figuerola-Ferretti and Gonzalo (2010). Instead of casting price discovery in terms of timing relationships, it draws on the structure of cointegrated processes to characterize the fundamental value of a commodity in terms of the common stochastic trend shared by spot and futures prices. We have done this within the traditional CVAR framework used by FFG, and within the fractionally cointegrated VAR of Johansen and Neilsen (2012, 2018) as applied by Dolatabadi et al. (2018). As a variation on the tradition of studying commodities important to the U.S. or European economies, we have concentrated on a subset of agricultural and extracted commodities significant to Canada.

Our empirical results strongly reject a unitary order of cointegration for almost all these commodities, supporting fractional cointegration. Even so, in other respects the two models yield results that are, for the most part, not dramatically different. For most commodities the cointegrating vector relating spot and futures prices is essentially $[1, -1]$, indicating neither systematic backwardation nor contango. Dynamic adjustment to disequilibrium largely conforms with standard intuition. And out-of-sample forecasting performance is similar, as is the potential that the two models offer for profitable trading strategies.

Hence, although the difference between the CVAR model and its FCVAR generalization is statistically significant, economic significance is more questionable; previous findings of economically significant differences may not be robust to variations in the commodities and sample periods. The generalization to fractional cointegration does not even reliably yield a more parsimonious modeling of dynamics, as is revealed by the lag lengths of Table 3 that we found necessary to adequately treat temporal dependence. In these respects, then, the FCVAR model may not provide significant value added over the CVAR.

Another respect in which the two models are similar is their implications for price discovery. Consistent with the conventional wisdom, for many commodities the futures market is the dominant contributor to price discovery, or at least contributes substantially. But for about half the commodities the spot market dominates, and in this there is considerable consistency between the CVAR and FCVAR models.

References

Campbell, J.Y., and Thompson, S.B. (2008) "Predicting the equity premium out of sample: can anything beat the historical average?" *Review of Financial Studies* 21, 1509–1531.

Clark, T.E., and West, K.D. (2007) "Approximately normal tests for equal predictive accuracy in nested models," *Journal of Econometrics* 138, 291–311.

Dolatabadi, S., Nielsen, M.Ø., and Xu, K. (2015) "A fractionally cointegrated VAR analysis of price discovery in commodity futures markets," *Journal of Futures Markets* 35, 339–356.

Dolatabadi, S., Nielsen, M.Ø., and Xu, K. (2016) "A fractionally cointegrated VAR model with deterministic trends and application to commodity futures markets," *Journal of Empirical Finance* 38, 623–639

Dolatabadi, S., Nielsen, M.Ø., Narayan, P.K., and Xu, K. (2018) "Economic significance of commodity return forecasts from the fractionally cointegrated VAR model," *Journal of Futures Markets* 38, 219–242.

Farm Credit Canada (2018) "Trade rankings report: Navigating trade disruptions and volatility," *FCC Ag Economics*, 1–12.

Figuerola-Ferretti, I., and Gonzalo, J. (2010) "Modelling and measuring price discovery in commodity markets," *Journal of Econometrics* 158, 95–107.

Garbade, K.D., and Silber, W.L. (1983) "Price movements and price discovery in futures and cash markets," *Review of Economics and Statistics* 65, 289–297.

Gonzalo, J., and Granger, C.W.J. (1995) "Estimation of common long-memory components in cointegrated systems," *Journal of Business and Economic Statistics* 13, 27–35.

Jobson, J.D., and Korkie, B.M. (1981) "Performance hypothesis testing with the Sharpe and Treynor measures," *Journal of Finance* 36, 889–908.

Johansen, S., and Nielsen, M.Ø. (2012) "Likelihood inference for a fractionally cointegrated vector autoregressive model," *Econometrica* 80, 2667–2732.

Johansen, S. and Nielsen, M.Ø. (2018) "Testing the CVAR in the fractional CVAR model," *Journal of Time Series Analysis* 39, 836–849.

Lasak, K.A. (2008) "Maximum likelihood estimation of fractionally cointegrated systems," *CREATES Research Papers,* 2008-53.

Marquering, W., and Verbeek, M. (2004) "The economic value of predicting stock index returns and volatility," *Journal of Financial and Quantitative Analysis* 39, 407–429.

Mines Canada (2019) *The Canadian Minerals and Metals Plan,* M4-175/2019E.

Nielsen, M.Ø., and Popiel, M.K. (2018) "A Matlab program and user's guide for the fractionally cointegrated VAR model," QED working paper 1330, Queen's University.

Peri, M., Baldi, L., and Vandone, D. (2013) "Price discovery in commodity markets," *Applied Economics Letters* 20:4, 397-403.

Silvapulle, P., and Moosa, I.A. (1999) "The relationship between spot and futures prices: Evidence from the crude oil market," *Journal of Futures Markets* 19, 175–193.

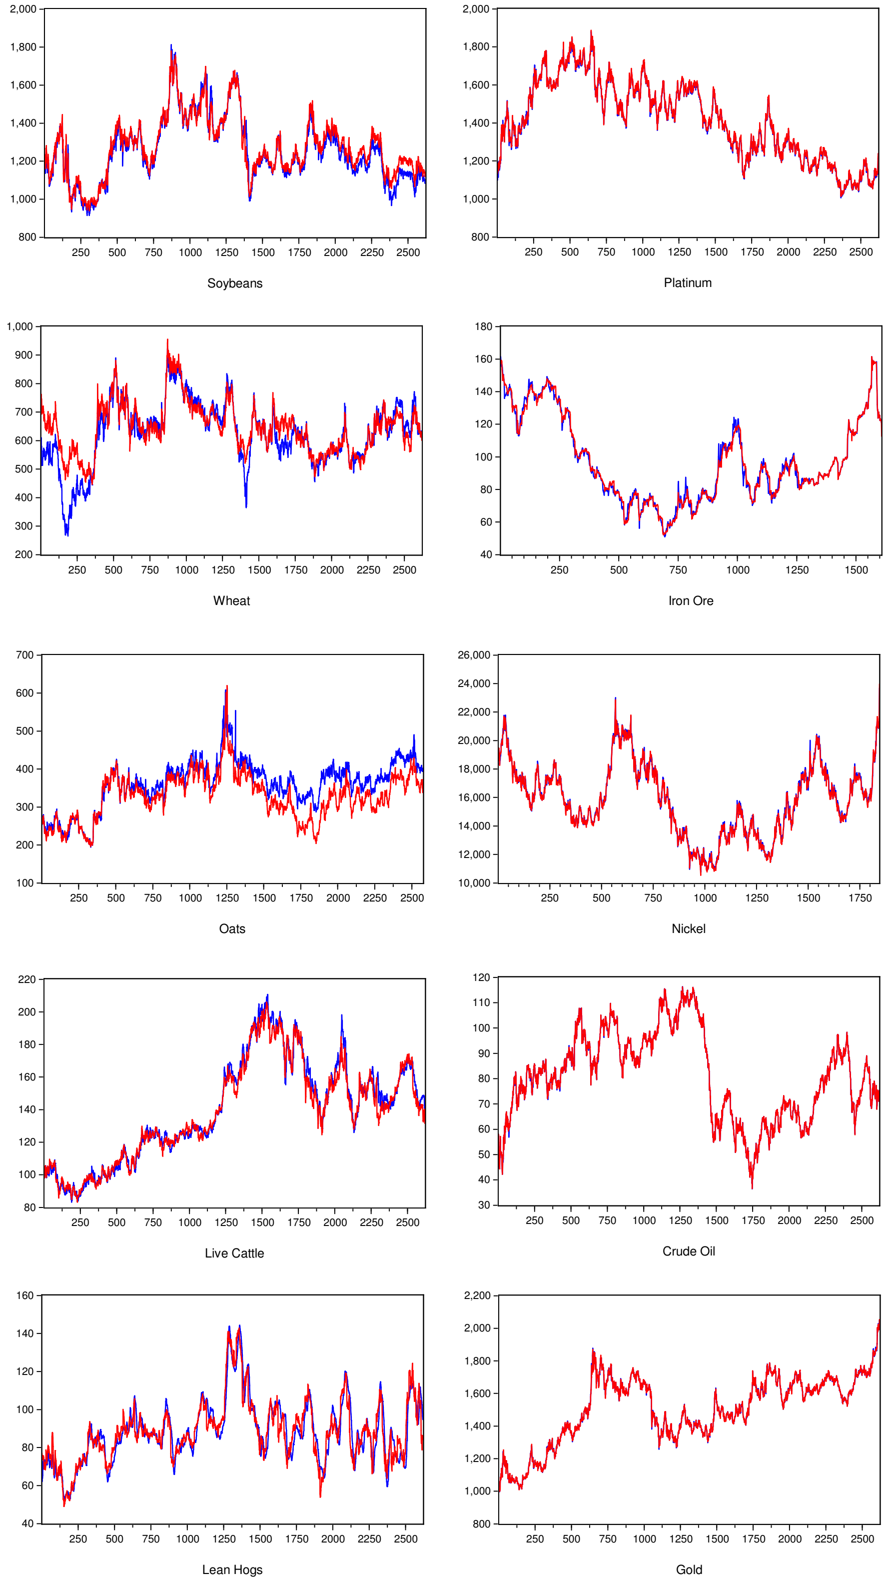

Figure 1: Daily Commodity Spot (blue) and Futures (red) Prices

Table 1: Average Daily Futures Contract Volume and Open Interest

Commodity	Barchart[a] symbol	Exchange	Sample start date	Observations[b]	Full sample[c] volume	Full sample[c] open interest	2019 sample[d] volume	2019 sample[d] open interest
Soybeans	ZS	CBOT	01/02/09	2621	196,956	638,897	205,732	705,033
Wheat	ZW	CBOT	01/02/09	2621	111,038	431,170	134,142	439,992
Oats	ZO	CBOT	01/02/09	2577	973	9959	664	5474
Live cattle	LE	CME	01/02/09	2621	53,759	314,654	68,858	381,402
Lean hogs	HE	CME	01/02/09	2621	41,880	227,601	64,271	273,059
Platinum	PL	NYMEX	01/02/09	2621	13,059	58,512	22,184	79,533
Iron ore	TR	NYMEX	02/19/13	1608	296	9426	316	7517
Nickel	Q0	LME	12/16/11	1845	1445	47,266	1307	205,360
Crude oil	CL	NYMEX	01/02/09	2621	822,011	1,723,293	1,205,514	2,056,436
Gold	GC	COMEX	01/02/09	2621	211,625	462,719	330,880	519,045

[a] www.barchart.com

[b] Number of observations is the number of trading days for which both a spot and futures price are available. Although the oats sample begins 01/02/09, it had fewer trading days with matched prices.

[c] Sample ends 30 August 2019 for all commodities.

[d] 2019 sample is the eight months January 2 through August 30.

Table 2: Descriptive Statistics: Daily Prices

Commodity	Mean	Standard deviation	Skewness	Kurtosis	Maximum	Minimum	ADF test (p-value)
Spot Market:							
Soybeans	1255.60	169.10	0.60	3.10	1812.50	913.90	0.4979
Wheat	625.38	112.33	−0.37	3.41	923.97	264.85	0.4522
Oats	360.93	63.19	−0.41	3.50	608.36	194.39	0.5313
Live cattle	139.75	31.40	0.14	2.06	210.71	83.16	0.6480
Lean hogs	88.33	16.49	0.62	3.75	144.37	52.38	0.7683
Platinum	1412.30	204.10	−0.00	1.90	1880.60	1006.50	0.5811
Iron ore	97.64	25.98	0.54	2.22	161.50	50.89	0.2488
Nickel	15,892.00	2515.00	0.00	2.00	23,689.00	10,758.00	0.7069
Crude oil	80.92	17.27	−0.02	2.18	116.28	36.43	0.5534
Gold	1503.80	205.40	−0.50	2.60	2049.40	998.00	0.9165
Futures Market:							
Soybeans	1275.00	158.40	0.40	3.10	1779.60	938.00	0.5017
Wheat	644.22	87.65	0.43	3.27	955.11	446.26	0.3878
Oats	328.50	59.68	0.17	3.71	620.04	196.27	0.5032
Live cattle	138.37	29.96	0.15	2.10	205.95	83.19	0.6228
Lean hogs	88.40	16.11	0.69	4.13	142.86	49.02	0.4577
Platinum	1416.70	204.10	−0.00	1.90	1887.20	1014.20	0.5709
Iron ore	97.58	26.06	0.53	2.22	161.43	52.08	0.2038
Nickel	15,844.00	2521.00	0.00	2.00	23,962.00	10,538.00	0.7258
Crude oil	80.96	17.26	−0.03	2.18	116.20	36.52	0.5533
Gold	1503.80	205.50	−0.50	2.60	2052.70	998.30	0.9076

Table 3: CVAR and FCVAR Models of Commodity Spot and Futures Prices

Commodity	Lag length k	Fractional parameter[a] b	Cointegration coefficient[a] $-\beta_2$	Adjustment coefficients α α_1	α_2	Price discovery shares $\alpha_\perp$[b,c] $\alpha_{\perp,1}$	$\alpha_{\perp,2}$
CVAR model:							
Soybeans	6	1	1.074**	−0.045	0.011	0.2069***	0.7931***
Wheat	5	1	1.519***	0.005	0.015	1.5611***	−0.5611***
Oats	11	1	1.230*	−0.000	0.017	0.9753***	0.0247
Live cattle	5	1	1.048***	−0.048	0.008	0.0821	0.9179***
Lean hogs	3	1	1.012	−0.035	−0.006	−0.2035	1.2035***
Platinum	4	1	1.000	−0.033	0.550	0.9439***	0.0561
Iron ore	5	1	0.997	−0.034	0.184	0.8447***	0.1553***
Nickel	0	1	0.999	−0.943	0.053	0.0528	0.9472***
Crude oil	2	1	1.000	−0.924	−0.400	−0.7622	1.7622***
Gold	1	1	1.000	0.205	1.172	1.2113***	−0.2113***
FCVAR model:							
Soybeans	8	0.647***	1.080**	−0.321	−0.087	−0.3735	1.3735***
Wheat	6	0.911**	1.462***	0.009	0.026	1.5399***	−0.5399
Oats	14	0.728***	0.779*	−0.158	−0.010	−0.0668	1.0668**
Live cattle	4	0.695***	1.021**	−0.247	0.010	0.0373	0.9627***
Lean hogs	6	0.591***	1.025	−0.202	−0.182	0.0632	0.9368**
Platinum	3	0.745***	0.998**	−0.204	0.974	0.8266***	0.1734
Iron ore	5	0.931	0.997*	−0.032	0.237	0.8804***	0.1196
Nickel	1	0.857***	1.002	−1.242	−0.003	−0.0021	1.0021***
Crude oil	8	0.669***	1.001	−0.899	2.273	0.7167	0.2833
Gold	1	0.909***	1.000	0.232	1.356	1.2064***	−0.2064**

[a] Asterisks indicate whether coefficients differ from one at 10(*), 5(**), and 1(***) percent significance.
[b] Asterisks indicate whether coefficients differ from zero at 10(*), 5(**), and 1(***) percent significance.
[c] Shares are defined to satisfy $\alpha'\alpha_\perp = 0$ and to sum to one, which yields $\alpha_{\perp,1} = -\alpha_2/(\alpha_1 - \alpha_2)$ and $\alpha_{\perp,2} = \alpha_1/(\alpha_1 - \alpha_2)$.

Table 4: Out-of-Sample Forecast Comparisons

	CVAR R^2		FCVAR R^2		Relative RMSE		CW statistic[a]	
Commodity	Spot	Futures	Spot	Futures	Spot	Futures	Spot	Futures
Soybeans	−0.0311	−0.0024	−0.0205	−0.0107	−0.5137	0.4136	2.5843***	−0.7003
Wheat	0.0145	−0.0027	0.0145	−0.0022	−0.0023	−0.0246	0.3366	0.6354
Oats	0.0142	0.0176	0.0184	0.0106	−0.2117	−0.3580	1.8037*	1.1352
Live cattle	0.0657	−0.0039	0.0692	−0.0051	−0.1871	−0.0614	1.4694	0.1338
Lean hogs	0.6678	0.0024	0.6580	0.0130	1.4628	−0.5348	−1.6634*	2.5337**
Platinum	−0.0181	0.0809	−0.0185	0.0783	−0.0181	−0.1457	0.3387	0.0236
Iron ore	0.0940	0.0091	0.0895	−0.0005	0.2450	0.4861	−0.8999	−1.8805*
Nickel	0.4262	−0.0056	0.4293	−0.0067	−0.2699	0.0573	1.9567*	−0.2302
Crude oil	0.0099	0.0100	0.0140	0.0132	−0.2074	−0.1643	1.5448	1.2130
Gold	−0.0141	0.1336	−0.0145	0.1318	0.0220	0.1033	−0.7051	−0.5975

[a] Asterisks indicate whether coefficients differ from zero at 10(*), 5(**), and 1(***) percent significance.

Table 5: Percentage Average Annualized Excess Returns: Futures Markets with Daily Rebalancing

	$\gamma = 3$		$\gamma = 6$		$\gamma = 12$	
Commodity	CVAR	FCVAR	CVAR	FCVAR	CVAR	FCVAR
Soybeans	0.0017	−0.0059	0.0008	−0.0030	0.0004	−0.0015
	(0.0041)	(0.0123)	(0.0020)	(0.0061)	(0.0010)	(0.0031)
Wheat	0.0148	0.0378*	0.0074	0.0189*	0.0037	0.0094*
	(0.0126)	(0.0215)	(0.0063)	(0.0107)	(0.0032)	(0.0054)
Oats	0.1451**	−18.4912	0.0726**	−6.9155	0.0363**	−3.8896
	(0.0584)	(29.0787)	(0.0292)	(18.8221)	(0.015)	(9.6521)
Live cattle	4.2353***	4.2312***	2.1176***	2.1156***	1.0588***	1.0578***
	(0.7586)	(0.8068)	(0.3793)	(0.4034)	(0.1896)	(0.2017)
Lean hogs	20.4989***	20.4067***	10.2494***	10.2034***	5.1247***	5.1017***
	(1.2533)	(1.2530)	(0.6266)	(0.6265)	(0.3133)	(0.3132)
Platinum	−0.0023	−0.0022	−0.0011	−0.0011	−0.0006	−0.0005
	(0.0020)	(0.0019)	(0.0010)	(0.0010)	(0.0005)	(0.0005)
Iron ore	4.7256*	4.4948*	2.3628*	2.2474*	1.1814*	1.1237*
	(2.6676)	(2.5589)	(1.3338)	(1.2795)	(0.6669)	(0.6397)
Nickel	0.0014***	0.0014***	0.0007***	0.0007***	0.0004***	0.0004***
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.00004)	(0.00004)
Crude oil	−0.1339	−0.5447	−0.0670	−0.6936	−0.0335	−0.3794
	(0.4762)	(3.2410)	(0.2381)	(1.7284)	(0.1190)	(0.8719)
Gold	−0.0015	−0.0015	−0.0008	−0.0007	−0.0004	−0.0004
	(0.0012)	(0.0011)	(0.0006)	(0.0006)	(0.0003)	(0.0003)
Average	2.9485	1.0126	1.4742	0.6972	0.7371	0.3022

Note: Standard errors are in parentheses. Asterisks indicate whether coefficients differ from zero at 10(*), 5(**), and 1(***) percent significance.

Table 6: Sharpe Ratios Under Intermediate Risk Aversion: $\gamma = 6$

Commodity	Number of forecasts τ	Model[a,b] CVAR	FCVAR
Soybeans	655	0.0156	−0.0192
		(0.0391)	(0.0391)
Wheat	655	0.0459	0.0690*
		(0.0391)	(0.0391)
Oats	644	0.0980**	-0.0145
		(0.0395)	(0.0394)
Live cattle	655	0.2181***	0.2049***
		(0.0395)	(0.0395)
Lean hogs	655	0.6391***	0.6364***
		(0.0429)	(0.0428)
Platinum	655	−0.0430	−0.0430
		(0.0391)	(0.0391)
Iron ore	402	0.0884**	0.0876*
		(0.0450)	(0.0450)
Nickel	461	0.3260***	0.3260***
		(0.0478)	(0.0478)
Crude oil	655	−0.0110	−0.0157
		(0.0391)	(0.0391)
Gold	655	−0.0521	−0.0456
		(0.0391)	(0.0391)
Average		0.1325	0.1186

[a] Sharpe ratios are $\hat{S} = \bar{r}/s$, where $\bar{r}$ is the average excess return and s is the standard deviation of the excess return. (By comparison, the standard errors of Table 5 are $s/\sqrt{\tau}$.)

[b] Values in parentheses are standard errors calculated using the formula $[(1 + \hat{S}^2/2)/\tau]^{1/2}$ from Jobson and Korkie (1981, p. 893). Asterisks indicate whether coefficients differ from zero at 10(*), 5(**), and 1(***) percent significance.