

The Fundamentals of Commodity Futures Returns

Gary B. Gorton
The Wharton School, University of Pennsylvania
and National Bureau of Economic Research
gorton@wharton.upenn.edu

Fumio Hayashi
University of Tokyo
and National Bureau of Economic Research
hayashi@e.u-tokyo.ac.jp

and

K. Geert Rouwenhorst
School of Management, Yale University
k.rouwenhorst@yale.edu

First Draft: March 20, 2007

This version: December 19, 2007

Abstract

Commodity futures risk premiums vary across commodities and over time depending on the level of physical inventories, as predicted by the Theory of Storage. Using a comprehensive dataset on 31 commodity futures and physical inventories between 1969 and 2006, we show that the convenience yield is a decreasing, non-linear relationship of inventories. Price measures, such as the futures basis (“backwardation”), prior futures returns, and prior spot returns reflect the state of inventories and are informative about commodity futures risk premiums. We reject the Keynesian “hedging pressure” hypothesis that these positions are an important determinant of risk premiums.

The Fundamentals of Commodity Futures Returns

Abstract

Commodity futures risk premiums vary across commodities and over time depending on the level of physical inventories, as predicted by the Theory of Storage. Using a comprehensive dataset on 31 commodity futures and physical inventories between 1969 and 2006, we show that the convenience yield is a decreasing, non-linear relationship of inventories. Price measures, such as the futures basis (“backwardation”), prior futures returns, and prior spot returns reflect the state of inventories and are informative about commodity futures risk premiums. We reject the Keynesian “hedging pressure” hypothesis that these positions are an important determinant of risk premiums.

1. Introduction

The relationship between storage and commodity futures' risk premiums is a classic question in the history of financial economics.¹ In this paper we analyze the fundamentals of commodity futures risk premiums and show that time-series variation and cross-sectional variation in commodity futures risk premiums are determined by the level of inventories of the commodity in the economy. The starting point of our analysis is the traditional Theory of Storage. Originally proposed by Kaldor (1939), the theory provides a link between the term structure of futures prices and the level of inventories of commodities. This link, also known as “cost of carry arbitrage,” predicts that in order to induce storage, futures prices and expected spot prices of commodities have to rise sufficiently over time to compensate inventory holders for the costs associated with storage.

In addition to market expectations of future spot prices, futures prices potentially embed a risk premium that is a compensation for insurance against future spot price risk. Whether futures prices also embed risk premiums has been more controversial in the literature. In part, this controversy stems from the difficulty in detecting risk premiums in volatile markets using small samples and short time series, and the lack of correlation of commodity futures returns with conventional measures of systematic risk suggested in the asset pricing literature.

To formalize the link between futures prices and risk premiums, Gorton, Hayashi, and Rouwenhorst (GHR 2007) present a simple theoretical extension of the theory of inventory behavior developed by Deaton and Laroque (DL 1992), and Routledge, Seppi and Spatt (RSS, 2000). Their models predict a link between the level of inventories and future spot price volatility. Inventories act as buffer stocks which can be used to absorb shocks to demand and supply dampening the impact on spot prices. DL show that at low inventory levels, the risk of a “stock-out” (exhaustion of inventories) increases and expected future spot price volatility rises. In an extension of the DL model which includes a futures market, RSS show how the shape of the futures curve reflects the state of inventories and signals expectations about future spot price volatility. Deaton and Laroque (1992) and Routledge, Seppi and Spatt (2000) have explained the existence of a convenience yield as arising from the probability of a stock-out of inventories. Because DL and RSS study storage in a risk-neutral world, risk premiums are zero by construction, and futures prices simply reflect expectations about future spot prices.

¹ See, for example, Keynes (1927), Kaldor (1939), Hicks (1939), Working (1949), Telser (1958), and Cootner (1960, 1967).

To allow for a link between inventories and futures risk premiums, GHR (2007) extend the DL model to include risk-averse agents and a hedging motive on behalf of producers. The GHR model predicts a link between the state of inventories, the shape of the futures curve, and expected futures risk premiums. Given that futures contracts provide insurance against price volatility, the level of inventories is negatively related to the required risk premium of commodity futures. The main contribution of our paper is to provide empirical tests of these predictions.

Despite the long history of the traditional Theory of Storage, surprisingly few researchers have attempted to directly test the theory using inventory data.² Often cited reasons include problems related to the availability and the poor quality of inventory data, and issues regarding the appropriate definition of relevant inventories. Most tests of the Theory of Storage have focused instead on testing predictions about the (relative) volatility of spot and futures prices.

The first contribution of this paper is to present measures of inventories for a large cross-section of 31 commodities between 1969 and 2006, and show that these measures of inventories are reflected in the shape of the futures curve as predicted by the Theory of Storage. As with much of the previous literature our initial focus is on the *basis*, the difference between the *current* spot commodity price and the current (nearest to maturity) futures price (expressed as a percentage of the spot price). We link the basis to the level of inventories, and empirically document the nonlinear relationship predicted by the existence of the non-negativity constraint on inventories. In particular, low inventory levels for a commodity are associated with an inverted (“backwardated”) term structure of futures prices, while high levels of inventories are associated with an upward sloping futures curve (“contango”). In addition we show that the relationship between inventories and the shape of the futures curve is non-linear: the slope of the futures curve becomes steeper as inventories decline.

The second contribution of the paper is to document an empirical link between inventories and risk premiums. We present two sets of tests to examine whether inventory levels are negatively associated with risk premiums on commodity futures. The first set of tests uses inventories directly as explanatory variables for risk premiums. In addition to simple regression based evidence, we show that sorting commodity futures into portfolios based on inventory measures is correlated with future average returns. While a direct test of the theory, the interpretation of these findings is complicated by an unknown timing lag in the information release of inventories data, and subsequent data revisions. The second set of tests uses price-

² Exceptions include Dincerler, Khokher and Titman (2003), and Dinceler, Khokher and Simin (2004). The former paper examines the effect of storage on Natural Gas futures returns between 1994 and 2001; the latter paper examines the role of inventories and hedging pressure for risk premiums in futures of Gold, Copper, Crude Oil, and Natural Gas between 1995 and 2004.

based signals to proxy for inventories. We first show that the futures basis, prior futures returns, and prior spot price changes are correlated with current inventory levels. Next, we show that these price-based measures of inventories are informative about the expected returns of portfolios sorted on these measures. Inspection of the inventory characteristics of these sorted portfolios confirms that the risk premiums carry a common component, earned in part by investing in commodities in low inventory states. The returns earned on “momentum” and “backwardation” strategies can therefore be interpreted as compensation earned for bearing risk during times when inventories are low.

Finally, we characterize the behavior of market participants in futures markets in response to changes in inventories. This is of interest because much of the literature on commodity futures has assigned an important role to the behavior of market participants in setting risk premiums. For example, in the Theory of Normal Backwardation, Keynes (1930) conjectured that the long side of a commodity futures contract would receive a risk premium because of hedging demand by producers. And in empirical implementations of the Theory of Normal Backwardation, researchers have linked “hedging pressure” to variation in futures risk premiums (e.g., Carter et al (1983), Bessembinder (1992), DeRoos et al (2000)). Using data obtained from the Report of Traders released by the Commodity Futures Trading Commission, we show that the positions of traders are contemporaneously correlated with inventories and futures prices. However, we find no evidence that these positions are correlated with ex-ante risk premiums of commodity futures. We therefore reject the hedging pressure hypothesis as an alternative explanation for the variation of risk premiums documented in our empirical work.

Our research builds on two strands of literature. The first starts with the traditional Theory of Storage developed by Kaldor (1939), Working (1949), and Brennan (1958), who explain futures prices in terms of the cost of storage, interest rates, and a convenience yield. The convenience yield was posited to explain why inventory holders would hold inventories during periods of expected decline of spot prices. Tests of the Theory of Storage include Fama and French (1988) and Ng and Pirrong (1994), among others. Both papers use the interest-adjusted basis as a proxy for inventories and examine the relation between the futures basis and price volatility.³ Fama and French (1988) analyze daily futures prices of metals over the period 1972 to 1983. Without inventory data, they use two proxies for determining when inventories are low. One proxy is the sign of the interest-adjusted basis. The second proxy is the phase of the business cycle. Fama and French (1988) argue that inventories are relatively low during recessions. With these proxies for inventory levels, they test their hypothesis that futures prices are less variable

³ Equation (1) below defines the basis.

than spot prices when inventory is low, an implication of the Theory of Storage, according to French (1986). They find evidence to support the predictions of Theory of Storage. Ng and Pirrong (1994) study four industrial metals. Like Fama and French (1988) they use the interest adjusted basis as a summary of supply and demand conditions and do not use inventory data. They examine the marginal impact of the basis (the “spread”) on variances, correlations, and elasticities of spot and futures. Their evidence is consistent with a concave, increasing relation between the adjusted spreads and inventories for spot and future return volatilities. Our contribution to this literature is that we directly examine the relationship between the basis and inventories using a large cross-section of commodities. In addition, our sample covers a longer span of time than previous research.

The second strand of literature primarily focuses on variation of risk premiums. Fama and French (1987) study 21 commodity futures using monthly data, over various periods, all ending in July 1984 and starting as early as March 1966. They examine both the variation in the futures basis and the information content in the basis about futures risk premiums. They find evidence that the basis varies with interest rates and seasonals (a proxy for convenience yields, since inventories are higher just after the harvest for agricultural commodities). They also decompose changes in the basis into the change in the expected spot price and the risk premium and conclude that most of the information in the basis concerns expected future spot price movements. Nash (2001), Erb and Harvey (2006), and Gorton and Rouwenhorst (2006) provide recent evidence of a relationship between the futures basis and futures risk premiums. Momentum in commodity futures has been documented by Pirrong (2005), Erb and Harvey (2006), Miffre and Rallis (2007), and Shen, Szakmary, and Sharma (2007). Chang (1985), Bessembinder (1992) and De Roon, Nijman and Veld (2000), Dincerler, Khokher and Titman (2003) and Dincerler, Khokher and Simin (2004) provide empirical evidence that traders’ positions are correlated with expected futures returns. Our contribution relative to these papers is to explain the relation between the returns and commodity characteristics as arising from fundamental variation in inventories as predicted by the Theory of Storage. And we show that expected futures returns are driven by inventories, instead of positions of traders.

In addition to these papers, there is a large literature about unconditional risk premiums in commodity futures markets. Attempts to empirically measure the risk premium on individual commodity futures have yielded mixed results (see, for example, Bessembinder (1992), Kolb (1992), and Erb and Harvey (2006)). Most of these studies use small samples in both the time series and cross sectional dimensions. Looking at portfolios of commodity futures returns has produced different results. Bodie and Rosansky (1980), and Gorton and Rouwenhorst (2005,

2006) provide empirical evidence that, consistent with Keynes' and Hicks' prediction, long investors in commodity futures have historically earned a positive risk premium. The issue of reconciling commodity risk premiums with received asset pricing theory has generally been met with limited success (see, for example, Dusak (1973), Jagannathan (1985)). The current paper sheds little light on this debate, other than to suggest that one avenue to look for a unified explanation of risk premiums is to consider systematic components of risk that are correlated with variation of inventories.

The remainder of the paper is organized as follows. In Section 2 we examine the relationship between inventories and futures prices in more detail. We summarize the theoretical results of GHR (2007) in this section. Section 3 documents our data and some stylized facts. Section 4 presents the empirical evidence on the link between futures prices and inventories, and provides evidence that the state of inventories is correlated with expected commodity futures risk premiums. In Section 5 we analyze the returns to price-based commodity selection strategies, linking these price-based signals to time-series and cross-sectional variation in commodity risk premiums. Section 6 looks in detail at the risk and return relationship between futures risk premia and the volatility of returns. In Section 7 we characterize the behavior of futures markets participants depending on the state of inventories. The final section summarizes our results and suggests some possible avenues for future research.

2. The Theory of Storage and Commodity Futures

In this section we briefly review some of the existing theories and outline the theoretical model of GHR (2007) and its testable hypotheses.

An upward sloping futures curve is consistent with an expected future spot price that rewards inventory holders for the cost of carrying inventories, including marginal warehousing costs, insurance, and the interest foregone on the capital invested in the inventories. This link between the futures price and the expected future spot price is known as "cost-of-carry" arbitrage. The cost-of-carry argument has difficulty explaining downward sloping futures curves. That is, researchers recognized early on that this argument cannot rationally explain why inventory is held when there is a predictable *decline* in spot prices, when futures prices fall below spot prices – i.e. agricultural products are held over the harvest period when prices predictably fall. To reconcile spot prices at levels above futures prices Kaldor (1937) postulated the existence of a "convenience yield" that holders of physical commodities earn but which does not accrue to holders of futures. This became known as the Theory of Storage.

This Theory of Storage (see, Kaldor (1939), Working (1949), and Brennan (1956)) can be stated in terms of the *basis*, the difference between the contemporaneous spot price in period t , S_t , and the futures price (as of date t) for delivery at date T , $F_{t,T}$.⁴ It views the (negative of) the basis as consisting of the cost-of-carry: interest foregone to borrow to buy the commodity, $S_t r_t$, (where r_t is the interest charge on a dollar from t to T), plus the marginal storage costs w_t , minus a “convenience yield,” c_t :

$$F_{t,T} - S_t = S_t r_t + w_t - c_t. \quad (1)$$

Equation (1) is often rationalized as following from the absence of arbitrage. Because the convenience yield is unobservable, an alternative view of equation (1) is merely that of a definition of the convenience yield. Economic content for equation (1) is provided by the assertion that the convenience yield, which is the basis adjusted for interest charges and storage costs, falls at a decreasing rate as aggregate inventory rises.

The Theory of Storage derives a relationship between contemporaneous spot and futures prices. Another view of commodity futures is the Theory of Normal Backwardation, which compares futures prices to expected future spot prices. As pointed out by Fama and French (1988), these views are not mutually exclusive. The Theory of Normal Backwardation views futures markets as a risk transfer mechanism whereby long (risk-averse) investors earn a risk premium for bearing future spot risk that commodity producers want to hedge. This theory builds on the view that the basis consists of two components: a risk premium, $\pi_{t,T}$, and the expected appreciation or depreciation of the future spot price:

$$F_{t,T} - S_t = [E_t(S_T) - S_t] - \pi_{t,T}, \quad (2)$$

where $\pi_{t,T} \equiv E(S_{t,T}) - F_{t,T}$. Equation (2) merely defines the risk premium. According to Keynes $\pi_{t,T} > 0$, which implies that the futures price is set at a discount (i.e., is “backwardated”) to the expected future spot price at date T , the date the futures contract expires. Keynes and Hicks (1939) view the risk premium as the outcome of the supply and demand for long and short positions in the futures markets (“hedging pressure”). If hedging demand exceeds the supply of long investors, the risk premium will be positive. The content of the Theory of Normal

⁴ The basis is also sometimes referred to as “backwardation”. In empirical applications, the basis is often measured as the difference between the nearest futures contract (i.e., the contract that is closest to maturity), and the next contract. This is due to difficulties observing the spot price.

Backwardation therefore comes from the assertion that hedgers are on net short and offer a risk premium to long investors, who are risk averse.

Since the Theory of Storage and the Theory of Normal Backwardation were first articulated, a large theoretical literature has developed.⁵ Our starting point is the modern version of the Theory of Storage due to Deaton and Laroque (1992, henceforth DL). Their goal is to explain the behavior of observed spot commodity prices, which display high volatility, high positive skewness, and significant kurtosis. Commodity prices show infrequent upward spikes, but no downward spikes. In their model commodity prices, in the absence of any inventories, would be i.i.d. because “harvests” of commodities are i.i.d. These price dynamics are changed fundamentally when inventories are present. Inventories cannot be negative (goods cannot be transferred from the future to the past), so there is a non-negativity constraint on inventories which “introduces an essential non-linearity which carries through into non-linearity of the predicted commodity price series” (DL, p. 1).

DL (1992) do not model futures markets. Routledge, Seppi, and Spatt (RSS, 2000) introduce a futures market into the DL model and show how the “convenience yield” arises endogenously as a function of the inventory level and the shock (“harvests”) affecting supply and demand of the commodity. The convenience yield – the benefit accruing to the physical owners of a commodity – arises from the non-negativity constraint on inventories, which creates an option for the inventory holder of selling commodities in the spot market when inventories are low.

In the DL and RSS models agents are risk-neutral. Hence, the commodity futures risk premium, which is central to the Theory of Normal Backwardation of Keynes and Hicks, is zero by assumption. In the GHR (2007) model of commodity futures, *both* the convenience yield *and* the risk premium emerge endogenously as functions of inventory. In this sense, equations (1) and (2) are both consistent with our equilibrium model. To link the equilibrium spot prices emanating from inter-temporal inventory decisions to commodity futures, GHR extend the DL model by adding futures markets and risk-averse investors to their model. They also assume that inventory holders face a bankruptcy cost, which provides them with a hedging motive. The existence of the futures market provides the inventory holders with an opportunity to hedge bankruptcy costs. They can use the futures market to transfer future spot price risk to risk averse investors, at a price. The model determines the risk premium paid by the inventory holders to the risk-averse

⁵ The literature on commodity futures is vast, and we make no attempt at a comprehensive survey. Reviews of the literature are provided by Carter (1999), Kamara (1982), and Gray and Rutledge (1971), among others.

investors, as a function of the extent of the size of the expected bankruptcy costs, the degree of risk aversion of the investors, and the level of inventories. The level of inventories matters for the risk premium because, as in DL, future spot price variance is negatively related to the level of inventories. That is, when inventories are low, the variance of the future spot price is higher due to an increased likelihood of a stock-out, resulting in the risk-averse long investors demanding a higher risk premium. The actual amount of hedging may either increase or decrease, depending on the relative sensitivities of the inventory holders and the investors to risk. We can summarize the relevant comparative statistics of the GHR model, as follows.

An inverse and nonlinear basis-inventory relation: Positive demand shocks and negative supply shocks lead to a drop in inventories, and result in an increase in spot prices, signalling the scarcity of the commodity in the spot market. Futures prices will also increase, but not by as much as spot prices. First, futures prices reflect expectations about future spot prices, and embed expectations that inventories will be restored over time and spot prices will return to “normal” levels. Second, the risk premium may increase. Both effects act to widen the difference between spot and futures prices. This inverse relation between the basis and inventory should become more pronounced as the inventory level is near stock-out if the demand for the commodity remains positive for very high prices, which is the case during occasional price spikes. We will be looking for evidence of this nonlinearity. This can be viewed as a test of the DL model of storage dynamics.

An inverse risk premium-inventory relation: When inventories are low and spot prices high, the buffer function of inventories to absorb shocks is diminished. In these circumstances the risk of a stock-out increases, which raises the conditional variance (volatility) of the future spot price. Because commodity futures are used to insure price risk, inventory theory predicts an increase in the risk premium.

Momentum in commodity futures excess returns: Although not formally modelled by GHR, inventories can only be restored through new production, a process which can take a considerable amount of time depending on the commodity. Therefore deviations of inventories from normal levels are expected to be persistent, as are the probability of stock-outs and associated changes in the conditional volatility of spot prices. Because past unexpected increases in spot and futures prices are signals of past shocks to inventories, they are expected to be correlated with expected futures risk premiums. This will induce a form of “momentum” in futures excess returns: the

initial unexpected spot price spike due to a negative shock to inventories will be followed by a temporary period of high expected futures returns for that commodity.

We now turn to testing these predictions.

3. Data and Summary Statistics

3.1 Commodity Futures Prices

Monthly data on futures prices of individual commodities were obtained from the Commodities Research Bureau (CRB) and the London Metals Exchange (LME). The details of these data are described in Gorton and Rouwenhorst (2006), who studied all 36 commodities futures that were traded at the four North American exchanges (NYMEX, NYBOT, CBOT, and CME) and the LME in 2004. For the present study, we drop electricity (because no inventory exists by its very nature), and gold and silver (because these are essentially financial futures). This leaves us with 33 commodities. We constructed rolling commodity futures excess returns by selecting at the end of each month the nearest to maturity contract that would not expire during the next month. That is, the excess return from the end of month t to the next month end is calculated as:

$$\frac{F_{t+1,T} - F_{t,T}}{F_{t,T}}$$

where $F_{t,T}$ is the futures price at the end of month t on the nearest contract whose expiration date T is after the end of month $t+1$, and $F_{t+1,T}$ is the price of the same contract at the end of month $t+1$.

Table 1 contains simple summary statistics for the 33 commodities for periods ending in December 2006. In addition to the 33 commodity futures, the first row of the table (labeled “index”) shows the statistics for an equally-weighted, monthly rebalanced, index of the commodity futures returns. It is therefore the simple average for each month of the excess returns for those commodity futures that were traded in that month. The period of calculation, which ends in December 2006, differs across commodities because the starting month varies. We take the starting month to be the latest of: the first month of the inventory series, the 12th month since the futures contract for the commodity started to trade, and December 1969. We require a 12-month trading history because later in the paper we will examine the role of prior 12-month returns. We require the starting month to be December 1969 at the earliest because before 1970 we have only two commodities (Cocoa and Soybeans) for which both futures price data and inventory data are available. The third column indicates the first month of the sample for the commodity. The fourth column of the table lists the number of monthly observations in our sample.

Columns 5-9 of the table have summary statistics of the distribution of excess returns. Although the sample period is slightly different than in Gorton and Rouwenhorst (2006), these summary statistics are qualitatively similar to their study. Of the 33 sample commodities 26 (21) earned a positive risk premium over the sample as measured by the sample arithmetic (geometric) average excess return. An equally-weighted index earned an excess return of 5.48% per annum. The next columns show that the return distributions of commodity futures typically are skewed to the right and have fat tails. DL (1992) make similar observations concerning the distribution of commodity spot prices. Columns 10 and 11 indicate that commodity futures excess returns are positively correlated (on average) with the returns on other commodity futures, but the correlations are on average low (0.12). The average correlation of individual returns with the return on the equally-weighted index is 0.40.

Finally, the last column of the table shows that the sample average (percentage) basis has been negative for two-thirds of the commodities.⁶ An equally-weighted portfolio of the sample commodities had an average basis of -2.10%, indicating that on average across commodities and time periods futures prices have exceeded contemporaneous spot prices. Otherwise stated, on average, commodity futures markets have been in “contango.” At the same time, the average excess return on the equally-weighted index has been positive (5.48% per annum), indicating a historical risk premium to the long side of a commodity futures position.

These observations are of interest, because the futures basis is often referred to by practitioners as the “roll-yield” of a commodity futures position, and a positive roll yield (also referred to as “backwardation”) is sometimes viewed as a requirement for the existence of a positive risk premium to a long position in commodity futures markets. This view is typically based on arguments such as that portrayed in Figure 1. Figure 1 plots the average basis against the average return on individual collateralized futures during the 1991–2006 period. Figure 1 suggests a connection between the risk premium and commodity characteristics, as measured by the basis. A simple linear regression has an R-squared of 52%.

In our discussion of equations (1) and (2) in Section 2, we already observed that these are not mutually exclusive: the futures basis compares futures prices to contemporaneous spot prices, while the risk premium in equation (2) is the difference between futures prices and expected future spot prices. Equation (1) shows that for commodities to be stored, futures prices have to

⁶ The basis is calculated for each commodity as $(F1/F2 - 1) * 365/(D2 - D1)$, where F1 is the nearest futures contract and F2 is the next nearest futures contract; D1 and D2 are the number of days until the last trading date of the respective contracts. The period over which the sample is calculated for the basis is from the month indicated in third column of the table to November 2006, so the sample size is the same as that for the excess return.

exceed contemporaneous spot prices to compensate inventory holders for the full cost of storage. Only when inventories are sufficiently low can the spot price exceed the futures price corrected for the cost of carry, i.e. when the convenience yield is sufficiently high. The sample average basis of -2.1% simply indicates that inventories have been sufficiently high on average for the convenience yield not to exceed the full cost of storage. At the same time futures prices have been set at a discount to average future spot prices, rewarding the long side of the futures position for providing price insurance.⁷

Figure 1 suggests a link between the presence of risk premiums and the basis. In this paper we explore this link in detail. We will show that the cross-section dependence arises from the fact that some commodities are harder to store than others. The relationship between the basis and ex-ante risk premiums is the subject of Section 5, in which we examine the predictive power of the basis for risk premiums, and the extent to which this predictability stems from variation in inventory levels. In the next sub-section we will present our inventory data.

3.2 Inventory Data

There are many issues involved in compiling a dataset on inventories, the least of which is the absence of a common data source. In addition to data availability, there is the important conceptual question of how to define the relevant inventories. Because most commodity futures contracts call for physical delivery at a particular location, futures prices should reflect the perceived relative scarcity of the amount of the commodity which is available for immediate and future delivery at that location. For example, data on warehouse stocks of industrial metals held at the exchange are available from the LME, but no data is available on stocks that are held off-exchange but that could be economically delivered at the warehouse on short notice. Similarly, relevant Crude Oil inventories would include not only physical stocks held at the delivery point in Cushing, Oklahoma, but also oil which is held at international locations but that could be economically shipped there, or perhaps even government stocks. Aside from the definition of relevant inventories there is a timing issue. Information about inventories is often published with a lag and subsequently revised. This creates a timing issue in matching variation of prices to variation of inventories. Despite these potential caveats, the behavior of inventories is central to the Theory of Storage and for this reason it is important to attempt to document the empirical relationship between measured inventories and futures prices.

⁷ A reference to financial futures may be instructive in this context, as financial futures do not have a convenience yield. When the dividend yield on equities is below the interest rate, equity futures price will exceed spot prices, and the markets will be in “contango” This is not incompatible with the presence of a positive equity risk premium.

We collected a sample of inventory data for the 33 individual commodities of Table 1 from a variety of sources. With the exception of Sugar, Feeder Cattle, and Rough Rice, we were able to find monthly data for all commodities. For Feeder Cattle, we do not use the available inventory series which is quarterly. Instead we use 3-month-ahead values of the Live Cattle inventory for the current monthly level of Feeder Cattle, under the assumption that it takes three months to feed calves to create what are called Feeder Cattle. A detailed description of these data is in the Appendix. In the rest of the paper, we will drop Sugar and Rough Rice and focus on the 31 commodities with monthly inventory data.

Examination of the data reveals that the inventory time-series of most commodities contains a time-trend and exhibits strong seasonal variation. We estimated individual inventory trends by applying a Hodrick-Prescott filter to the log of inventories for individual commodities. We will sometimes refer to the Hodrick-Prescott (HP) filtered inventory data as the “normal” inventory level and denote it by I^* .⁸

To illustrate the seasonal variation of commodity inventories around these trends we ran a regression of the deviations of the log of inventories from their HP-fitted trends on monthly dummy variables. Table 2 reports the regression results along with the autocorrelation of the residuals (which are de-trended and de-seasonalized inventories). The table helps to illustrate two stylized facts about inventories. First, inventory levels are persistent. At 0.71 inventories of Soybean Meal have the lowest sample first-order autocorrelation, and the median first-order autocorrelation exceeds 0.90. Second, there are large cross-sectional differences in the seasonal behavior of inventories. This is illustrated in Figure 2, which shows the seasonal variation of inventories of Natural Gas, Wheat, and Corn. The seasonal variation of inventories stems from both demand and supply. Many agricultural commodities are harvested once a year and inventories are held to meet demand throughout the year. Inventories therefore are lowest just prior to the harvest season and peak at the end of the harvest season. For example, Corn is harvested in late summer to fall in North America. Wheat is harvested in the early summer in the Southern states and late summer in the Northern states. Wheat inventories therefore are lowest

⁸ The smoothness parameter we use when applying the Hodrick-Prescott filter to monthly series is determined as follows. Ravin and Uhlig (2002) recommend adjusting the smoothness parameter in proportion to the fourth power of the relative frequency. So if x is the smoothness parameter for a quarterly series, the monthly equivalent is x times $3^4 (=81)$. In business cycle analysis, it is customary to use 1,600 for quarterly series. As shown in Ravin and Uhlig (2002), this amounts to retaining peak-to-peak cyclical movements of roughly 10 years or longer, so the difference between the raw series and the filtered series consists of movements of relatively short durations. One would think that determinants of a normal inventory, such as storability and production flexibility, change only gradually. If so, the smoothness parameter should be larger. From visual inspection, we chose a smoothness parameter of 160,000 (whose monthly equivalent is this times 81). This amounts to retaining peak-to-peak cyclical movements of about 30 years or longer.

just prior to the harvest season and peak at the end of the harvest season. Contrary to Corn and Wheat, Natural Gas is produced throughout the year, but heating demand has a strong seasonal component which peaks during the winter months. During months of low demand, Natural Gas is stored in underground salt domes. Industrial Metals inventories exhibit little seasonal variation as exhibited by the low regression R-squared given in Table 2. Crude Oil is demanded and produced during the year, but demand for its derivatives --- Heating Oil and Unleaded Gas --- is more seasonal. Because Soybean Oil and Soy Meal are derived commodities and can be produced throughout the year, they exhibit less seasonality than the inventories of Soybeans themselves.

4. Inventories and Futures Prices

This section provides empirical evidence about the relationship between (1) inventory levels and risk premiums of commodity futures and (2) between inventories and the basis. In Section 4.1 we test the central prediction of the Theory of Storage that the marginal convenience yield as proxied for by the basis is a declining function of inventories. This motivates the use of the basis as a measure of the state of inventories. Section 4.2 examines the link between inventories and risk premiums.

4.1. Basis and Inventories

As a preliminary test, we examine whether the futures basis varies between high and low inventory months. Let I and I^* indicate the actual and normal inventory level at the end of the month.⁹ For each commodity we calculate the average basis for months when the normalized inventory I/I^* (the ratio of inventory level to the HP-filtered inventory) is below 1 and above 1. The results are summarized in Panel A of Figure 3. The figure illustrates that for all commodities low inventory months are associated with above average basis for that commodity and that the basis is below average during high inventory months. As indicated by the red line, the difference is statistically significant at the conventional 5% level for most commodities. (The calculation of the t -values is explained in GHR (2007) Appendix C.)

To further explore the non-linear relationship between the basis and inventories we estimate the following regression:

$$Basis = \text{linear function of seasonal dummies} + h(x) + \text{error},$$

⁹ For simplicity we have omitted time subscripts, but keep in mind that the normal inventory level changes through time.

where x is the normalized inventory level I/I^* . The hypothesis is that as the inventory levels fall below “normal,” as measured by I^* , the basis increases at an increasing rate. To allow for this nonlinearity we applied the “cubic spline regression” technique (see. e.g., Green and Silverman (1994) for a textbook treatment). This is a technique for estimating potentially nonlinear functions. Splines are piece-wise polynomial functions that fit together at “knots.” In the case of cubic splines, the first and second derivatives are continuous at the knots.¹⁰

To test whether the basis is negatively related to inventories and whether the relationship is, in fact, nonlinear, we will estimate the slope, implied by the spline function $h(x)$ at the average level of inventories ($I = I^*$) as well as in situations when inventories fall 25% below average ($I/I^* = 0.75$). For each commodity, the sample period is the same as shown in Table 1. The results of these tests are summarized in Table 3, and illustrated in Panel A of Figure 4 for Copper and Panel B for Crude Oil.

The second and third columns of Table 3 show that at the average level of inventories (i.e., at $I=I^*$), the estimated slope of the basis-inventory regression is negative for all commodities except one, and statistically significant for more than half of the commodities. For each commodity group, using pooled OLS we estimate the coefficients under the constraint that they are the same within groups. Inspection of the size of the coefficients shows that the relationship is particularly strong for commodities in the Energy group (the pooled OLS estimate for Energy is -1.546), while many Industrial Metals tend to have slope coefficients that are relatively small in magnitude (the pooled OLS estimate is -0.051). Industrial Metals are relatively easy and cheap to store, and equilibrium inventories of Industrial Metals are expected to be large on average relative to demand. By comparison, Energy which is more bulky and expensive to store, should have lower inventories relative to demand. Cross-sectional differences in storability should therefore be reflected in the sensitivity of the basis to inventory shocks. Perishability also helps to explain why the slope coefficients for Meats are on average larger than for commodities in the Softs and

¹⁰ The internal breakpoints that define the piecewise segments are called “knots.” Let x_j ($j = 1, 2, \dots, J$, $0 < x_1 < x_2 < \dots < x_j$) be so-called “knots”. The cubic spline technique approximates $h(x)$ by:

$$h(x) \approx \beta_1 x + \beta_2 x^2 + \beta_3 x^3 + \sum_{j=1}^J \beta_{3+j} (x - x_j)^3 \mathbf{1}\{x > x_j\},$$

where $\mathbf{1}\{\}$ is the indicator function. By construction, the second derivative of $h(x)$ is continuous at each knot. The attraction of a cubic spline is that the approximating function is linear in powers of x . We experimented with J on our data, and decided to set $J = 1$ and set x_1 to be 1 (i.e., $I = I^*$). For larger values of J , there were too many peaks and troughs in the estimated cubic spline.

Grains groups. Because storage costs provide an incentive to economize on inventories, it is also expected that the variation of inventories is lower for commodities that are difficult to store, relative to commodities that are easy to store: this is illustrated in the two panels of Figure 4 which shows much larger variation in the inventories of copper than in the inventories of crude oil.

To examine the non-linearity of the basis-inventory relationship, the fourth column of Table 3 reports the slope when inventories fall by 25% from their average value. In the case of Copper, for example, the estimated slope measured at the average level of inventories equals -0.032 ($t = -0.61$) and steepens to -0.153 ($t = -2.76$) when inventories drop by 25%. This difference of 0.121, given in column 6, is significant at the 5% level ($t = 5.64$). Inspection of columns 6 and 7 shows a pattern of steepening slopes for many commodities in the Metals, Grains, and Softs group. The results are weaker for Meats and Energies. Inspection of the inventory data for energy commodities shows that historical inventories often fluctuate within a narrow range, and in some cases do not fall to the test level of 0.75. Consequently, the slope coefficients at 0.75 are merely polynomial extrapolations of a relationship constructed to fit a different portion of the sample and should be taken with caution. This point is clearly seen from Panel B of Figure 4 for Crude Oil.

Overall our results are not inconsistent with the Theory of Storage.¹¹ We find that there is a clear negative relationship between normalized inventories and the basis and that for many commodities the slope of the basis-inventory curve becomes more negative at lower inventories levels. And we find steeper slopes at normal inventory levels for commodities that are difficult to store. We turn to the relationship between inventories and risk premiums next.

4.2. Inventories and Futures Risk Premiums

As mentioned previously, the Theory of Storage due to DL (1992) does not make direct predictions about futures risk premiums, but instead makes predictions about the future volatility of spot prices. This prediction stems from the fact that when inventories are low, the ability of inventories to absorb shocks to demand and supply is diminished, raising the conditional volatility of future spot prices. In our model, to the extent that the risk premium on long futures positions is compensation paid by hedgers to obtain insurance against price risk, the mean excess return from commodity futures should increase when future spot price risk increases. Therefore, the Theory of Storage implies that the state of inventory at the end of the month is a key predictor

¹¹ The results of Table 3 are not significantly altered if the dependent variable is the interested-adjusted basis; see Equation (1).

of the excess return from the end of the month to the next and that the mean excess return and inventory are inversely related.

As a first test of this prediction, we perform a linear regression of the monthly excess return on I/I^* measured at the end of the previous month as well as monthly dummies. The Theory predicts that I/I^* , our measure of the state of inventories, should have a negative effect on the subsequent excess returns. The results are reported in Table 4. Unlike in the basis-on-inventory regression of Table 3, we only consider the linear specification because the excess return is a hard variable to predict, as evidenced in the low R-squared in Table 4. As is apparent from the low t -values, the I/I^* coefficients are not sharply estimated. However, most of them have the expected negative sign. If we impose the restriction of a common slope coefficient within groups, we find marginally significant negative slope coefficients for Meats and Energy. These groups also exhibit a larger sensitivity of returns to inventories, which is consistent with our findings in Table 3 that futures prices of commodities that are difficult to store are more sensitive to inventory shocks than commodities that are relatively easy to store.

In a second test, we examine the results of a simple sorting strategy, whereby at the end of each month we cross-sectionally rank the commodities based on their level of normalized inventories. We compare the average return of a portfolio of commodities in the top half in terms of normalized inventories to the average return of a portfolio comprised of the commodities in the bottom half of this ranking. We measure the total futures returns of these portfolios during the month until the last day of the month when we re-sort and rebalance. The portfolios are equally-weighted. This test is nonparametric; it allows for a non-linear relationship between inventories and the risk premium. And comparing the returns of characteristic-sorted portfolios has the additional attractive feature that it controls for the cross-sectional dependence.

The results are given in Table 5. The returns of the inventory-sorted portfolios are consistent with the predictions of the theory that low inventories are associated with high future risk premiums. Panel A summarizes the returns to these portfolios in deviation from the equally-weighted index. The first columns show that the Low Inventory portfolio has outperformed the High Inventory portfolios in 56% of the months between 1969 and 2006. The annualized average out-performance was 8.06 % ($t = 3.19$). The next columns show that the performance difference between the inventory-sorted portfolios has been relatively stable during the most recent period.

In Panel B of Table 5, we summarize various characteristics of the commodities in the inventory sorted portfolios: for reasons we will discuss in greater detail in the next section, we report the average prior 12-month futures return prior to portfolio formation, the average percentage 12-month change in spot prices, the average futures basis and the average commodity

volatility (measured as the standard deviation of daily excess returns during the month) during the holding period. The Low Inventory portfolio selects commodities with a high basis: the difference between the basis of the Low and High Inventory portfolios exceeds 14% ($t = 14.51$). This is, of course, a direct implication of the Theory of Storage, and consistent with our earlier findings in Table 3, and Figure 3. In addition to having a higher basis, Low Inventory commodities also have higher prior spot and prior futures returns than High Inventory commodities. Over the full sample, the 12-month futures return difference prior to inclusion in the portfolio is about 14.9% per annum ($t = -6.45$). The high prior futures return of the Low Inventory portfolio suggests that our portfolio sorts capture more than variation of inventories that is predictable. High prior futures returns are an indication of past negative shocks to supply and/or positive shocks to demand. Because inventories cannot be replenished instantaneously, the prior futures return history carries information about the current state of inventories. We will return to this issue in the next section when we investigate the extent to which inventory dynamics can be responsible for the presence of momentum in commodity futures markets.

Finally, the right hand of Panel B summarizes the positions of traders in futures markets. It shows that Commercial traders are net short in commodity futures markets and as a percentage of open interest, that their positions are larger for High Inventory commodities. Data on positions of large traders is published by the Commodity Futures Trading Commission (CFTC). In the CFTC's Commitment of Traders Reports large traders are classified as "commercials" or "non-commercials." This is discussed further in Section 7.

Two caveats are in order about our trading rule test. First, the tests do not control for (unknown) publication delays in the release of inventory data. If news about inventories is negatively correlated with contemporaneous spot prices, and inventory data is released with a lag, this will create a negative correlation between innovations to inventories and subsequent spot price innovations. Because futures prices will inherit spot price innovations, the delay of news about inventories will create a correlation between inventories and subsequent futures returns that is unrelated to futures risk premiums. Second, our test does not exploit cross-sectional differences between commodities. Because commodities differ in terms of storability (perishability, bulkiness, and capacity constraints of storage) the Theory of Storage predicts that equilibrium inventory policies will differ across commodities. Furthermore, uncertainty about future demand and supply is also likely to vary across commodities, leading to cross-sectional differences in optimal inventory policies that are positively associated with future risk premiums.

Absent a structural equilibrium model that includes multiple commodities we have no guide as to how to compare the state of inventories across commodities. Theoretically, the

important state variable is the “likelihood of stock-out,” which we have proxied for by using I/I^* , the inventory level relative to normal inventories, but this measure does not permit comparisons across commodities. In the next section we will examine three predictions of the Theory of Storage that use price-based measures of the state of inventories that circumvent these difficulties.

5. Price-Based Tests of the Cross-Sectional Variation of Futures Risk Premiums

In the previous section we provided evidence that the shape of a futures curve, i.e., the basis, reflects information about the state of that commodity’s inventory, and that inventory levels are negatively related to subsequent excess returns to commodity futures. In this section we discuss three additional and related predictions of the Theory of Storage about spot and futures prices. First, when inventories fall spot prices will increase, signalling the scarcity of the commodity for immediate delivery. High spot prices are therefore a signal of the state of inventories. Second, shocks to current inventories also raise futures prices although not by as much as spot prices reflecting expectations that inventories will be restored over time and spot prices will return to “normal” (and perhaps because the risk premium rises). Hence, the futures basis widens. Third, to the extent that inventories are slow to adjust, past demand and supply shocks will persist in current inventory levels. Because unanticipated shocks to demand and supply affect futures prices, the futures return history of a commodity carries information about past demand and supply shocks that may not be fully resolved due to the slow adjustment of inventories. In sum, the level of spot prices, the futures basis, and prior futures returns can be expected to carry information about the current state of inventories, and hence will be correlated with risk premiums.

Panel B of Figure 3 illustrates the relation between inventories and 12-month prior futures returns for individual commodities. Similar to Panel A, for the basis, we calculate average prior 12-month futures returns for each commodity for months when I/I^* is above unity and when I/I^* is below 1. The Figure illustrates that for most commodities, high normalized inventories are associated with low futures returns over the prior year, while low inventory states are associated with high prior 12-month futures returns. Taken together, Figure 3 shows that prior futures returns and the basis are informative price-based signals of the level of inventories. To the extent that the level of inventories is relevant for futures risk premiums, as suggested in Table 5, it can be expected that prior futures and spot returns and the basis predict risk premiums on commodity futures. In the remainder of this section we will examine the extent to which these price signals carry information about expected futures returns.

There are two advantages to using observable prices as indicators of the state of inventories. First, price information does not suffer from revisions and publication delays associated with inventory data. Second, using price information opens the potential to exploit cross-sectional differences between expected commodity futures returns. For example, if a particular commodity is difficult or costly to store, then all else equal, the Theory of Storage predicts a lower level of equilibrium inventories. Lower average inventories will make a commodity more susceptible to the risk of stock-outs, and the associated futures contract is expected to have a higher equilibrium risk premium. To the extent that these cross-sectional differences are embedded in the shape of the futures curve such as the basis, we expect our price signals to capture this information about cross-sectional differences in expected futures returns

To quantify the information in price signals about both the cross-sectional and time-series variation in risk premiums, we divide the sample of commodities into halves at the end of each month based on their prior performance and the futures basis. We measure the total futures returns of these portfolios during the month until the last day of the month when we re-sort and rebalance. The portfolios are equally-weighted. The performance and characteristics of the portfolios are given in Tables 6, 7 and 8.

Panel A of Table 6 summarizes the returns on the portfolios formed by sorting based on the basis. Over the full sample period since 1969, the High Basis portfolio outperformed the equally-weighted index by 5.42% annualized ($t = 3.98$) while the Low Basis portfolio underperformed the average commodity by 4.82% ($t = -3.44$). The difference between the High and Low Basis portfolio was positive in 58% of the months and averaged 10.23% annualized ($t = 3.73$).

Panel B of Table 6 reports several characteristics of the basis-sorted portfolios. To the extent that the futures basis carries information about the state of inventories, it can be expected that the High Basis portfolio selects commodities that have below average inventories, high spot prices (measured relative to the same time last year), and high prior 12-month futures returns. And as predicted by DL, High Basis commodities are expected to have relatively high future price volatility. These predictions are indeed borne out by the data: the High Basis portfolio selects commodities with low inventories ($t = -17.08$), high futures returns during the 12-month period prior to portfolio formation ($t = 12.93$), and high spot prices relative to the same time a year prior ($t = 10.45$). Somewhat surprisingly, the difference between the volatility (measured as the standard deviation of daily excess returns during the month) of the commodities is both economically as well as statistically relatively small ($t = 2.13$). We return to this issue below in Section 6.

The right two-thirds of Table 6 examines two more recent sub-periods. These panels show that these returns and portfolio characteristics have been relatively stable during the first and second halves of our sample. The last three rows of Panel B summarize the positions of Traders in the basis-sorted portfolios, as reported by the Commodities Futures Trading Commission (CFTC). These will be discussed in more detail in Section 7 of the paper, but for now note that Commercials are on average net short in both the High and Low Basis portfolios, and Non-Commercials and small (Non-Reportable) traders are net long. Non-Commercials are over-weighted in the High Basis commodities, and the reverse holds for the Non-Reportable positions. There is no significant difference between the positions of Commercials between the two portfolios.

Inspection of the portfolio characteristics suggests that the basis-sorted portfolios capture time-series variation of risk premiums by selecting commodities when inventories are low. However, as pointed out before, differences in the basis can also reflect cross-sectional differences in storability of commodities that are correlated with (unconditional) risk premiums. To examine whether the returns to the basis strategies capture time-series variation of risk premiums or simply select commodities that are difficult to store, we repeat the portfolio sorts after subtracting the full sample mean from the basis for each commodity. This isolates the returns that can be attributed to time-series variation of the basis from return variation attributed to cross-sectional variation in the average basis. Unreported results show that the sample average return difference between High and Low (de-measured) Basis portfolios is 9.72% ($t = 3.75$) which is not significantly different from the returns associated with sorting on the raw basis. This suggests that the returns of sorting commodities on the raw basis primarily captures time-variation of futures returns that is associated with time variation of inventories.

Table 7 summarizes the returns to sorting commodities on Futures Momentum, measured as the prior 12-month futures return. Although momentum has been documented at horizons ranging from one month to one year, we chose to report results for a relatively long prior return interval (e.g., see Pirrong (2005) and Shen, Szakmary, and Sharma (2007)). Our choice is driven by our goal of constructing a price-based measure of inventories. Based on the empirical evidence of Table 2 that inventories are slow to adjust, we expect relatively distant prior shocks to inventories to carry information about current inventories. Because some commodities have distinct annual seasonal variation in production, we include a history of up to one year in our price-based measure of past positive demand shocks or negative supply shocks. Unreported results show that sorting on longer term measures of past futures returns increases the dispersion between the inventory characteristics of the momentum portfolios.

Panel A shows that High Momentum commodities have outperformed a portfolio of Low Momentum commodity futures by 13.36% per annum ($t = 4.93$), earning positive excess returns in 58% of the months. The point estimates for the excess returns are slightly higher for the second half of the sample, as well as the fraction of the months the excess return is positive (65% since 1990, versus 58% over the full sample). Panel B shows that Momentum portfolios take positions in similar commodities as the Basis-sorted portfolios. In particular, the High Momentum portfolio selects commodities with High Basis and below average inventories, while the Low Momentum portfolio does the opposite. The t -statistics associated with these characteristics differences are large and clearly indicate that portfolios sorted on inventories, the basis, and prior performance take correlated positions in ways that is predicted by the Theory of Storage. This is reflected in the correlation between the returns to High Basis and High Momentum portfolios, which is 0.87 over the full sample period. Inspection of the Positions of Traders reveals that Commercials increase their short positions in commodities that experience price increases, while Non-Commercials take larger long positions following a price run-up.

Finally, Table 8 reports the results from sorting commodities based on the change in the year-on-year percentage change of the commodity spot price. In light of the seasonality of spot prices of many commodities the 12-month prior spot return captures the change in the relative scarcity of each commodity compared to the same time a year ago. Panel A of the Table shows that the results for portfolios sorted on Spot Momentum are quantitatively similar to those sorted on Futures Momentum. The High Spot Momentum portfolio has outperformed the Low Momentum portfolio by 13.85% annualized ($t = 4.95$) over the full sample, and by 16.03% during the last 16 years ($t = 4.47$). And High Spot Momentum commodities have relatively low inventories, high futures momentum, and a high basis. Inspection of the Positions of Traders shows that Commercials hedge more after spot prices have increased, and that much of the liquidity to them is provided by the Non-Commercials.

The main conclusion from Tables 5-8 is that, consistent with the predictions of the Theory of Storage developed in GHR (2007), risk premiums of commodity futures vary with the state of inventories. Portfolios that take positions based on prior futures return, prior spot returns and the futures basis select commodity futures with below average inventories which the Theory predicts are expected to earn higher risk premiums. Moreover, these risk premiums are highly significant, both in a statistical sense as well as in an economic sense. We also presented some evidence that the Position of Traders varies with the return of the price-based portfolio strategies – especially Momentum and Inventories, although the interpretation of the positions evidence is somewhat ambiguous. Commercials increase their short positions after price run-up, but also

when inventories are high. Non-commercials take larger long positions in commodities with high momentum, and to a lesser extent high basis.

In the GHR model, the correlation between inventories and the amount of open interest in the futures market is ambiguous and depends on the relative sensitivities of the risk-averse investors and the inventory holders, seeking to hedge bankruptcy costs. However, the co-movement between the basis, inventories, momentum and traders' positions raises the question of a causal relationship; in particular we are interested in whether the positions of market participants can provide an alternative explanation for our results. We explore this issue in more depth in Section 7.

6. The Risk Premium and Price Volatility

In the GHR extension of the Theory of Storage, the futures risk premium is linked to the state of inventories through expected future spot price volatility. At low inventory levels risk-averse agents require a higher risk premium, because lower buffer stocks increase the volatility of expected future spot prices (DL (1992)). In our empirical work thus far, we find evidence of a link between the level of inventories (and price measures of that level) and risk premiums, but the evidence of the relationship with volatility is weak. Our measures of monthly futures return volatility, calculated from the volatility of daily returns, are generally higher for low inventory portfolios than for high inventory portfolios, but Tables 5-8 show that a test of statistical significance generally fails to reject the hypothesis of equal volatilities.

At first glance, these weak volatility results are of some concern as they are central to the insurance argument of the theory. However, our ability to detect volatility effects may be weakened by the migration of commodities between portfolios in the presence of cross-sectional differences in volatility. For example, the arrival of spring marks the end of a period of peak demand for Natural Gas as well as the start of the growing season for Wheat. Uncertainty about Wheat prices is likely to rise relative to uncertainty about Natural Gas prices. However, if Wheat were to replace Natural Gas in the High Basis portfolio during the spring, the volatility of the average commodity in the High Basis portfolio is likely to fall relative to the average volatility of commodities in the Low Basis portfolio. This is because Natural Gas has much higher unconditional volatility than Wheat (see Table 1), despite the fact that Wheat Prices become more volatile during the growing season, and Natural Gas prices become less volatile after the end of winter.

We perform two tests of the inventory-volatility channel in which we control for cross-sectional differences among commodity price volatilities. In the first test, we recalculate Tables

5-8 using the demeaned time series of returns and characteristics of each commodity. Demeaning the data prior to sorting removes the effects of cross-sectional differences between commodities, and isolates the time series variation of the returns and characteristics in the data. It therefore addresses the question of whether the volatility of a commodity is high or low relative to its own average when it is included in a particular portfolio. Table 9 summarizes the t -statistics for the difference between the average volatility of commodities in the High and Low portfolios sorted by the various characteristics: panel A applies to the demeaned data and, for comparison, panel B collects the t -statistics from the original data in Tables 5-8.¹²

Table 9 shows that controlling for the cross-sectional differences between commodities significantly sharpens our previous findings for volatility. The top panel shows that High Basis and High Momentum are associated with above average subsequent return volatility, while low values of these characteristics are associated with below average volatility. The size of the t -statistics indicates that these differences are significant at the 5% level for all sub-periods considered. The results for the inventory sorted portfolios remain qualitatively unchanged.

In our second test we directly examine the distribution of monthly returns conditional on a proxy for the state of inventories. The hypothesis is that both the mean and standard deviation will be higher when inventories are lower. To control for cross-sectional differences of average returns and price volatility, we examine the conditional distribution of normalized commodity returns, obtained by rescaling the time-series of excess returns for each commodity to have zero mean and unit standard deviation. Next, we rank all available observations by commodity on the Basis observed at the end of the prior month. The returns corresponding to the lowest (highest) Basis quintile are assigned to the Low (High) Basis normalized distribution. These two conditional distributions therefore each contain 20% of the (normalized) return observations of each of the sample commodities. The normalization of the returns facilitates the comparison across commodities.

The two distributions are displayed in Figure 5, which visualizes the properties of the conditional (normalized) distributions of returns as summarized in the embedded table. The High Basis distribution has higher mean and standard deviation than the Low Basis distribution. Testing for the difference in the High Basis and Low Basis distributions is complicated by the fact that the (normalized) returns used to calculate the two distributions are not independent draws from the same distribution, so the usual tests (e.g., the Kolmogorov-Smirnov test) for the

¹² For brevity, the full tables are not included in the paper but are available from the authors on request. With the exception of the volatility results reported in Table 9, the results for the demeaned data are quantitatively similar those of Tables 5-8.

difference in two distributions cannot be used. The returns conditional on the Basis signal for each commodity may be serially uncorrelated, but they are contemporaneously correlated across commodities. In order to address this issue, we construct a bootstrap of the return distributions which preserves the cross section correlation of commodity Basis measures, but delinks these signals from the timing of the returns.¹³ The delinking of the signals from the returns will allow us to produce bootstrapped distributions which can be used to test whether the differences in moments, displayed in Figure 5, are statistically significant.

The results are shown in Table 10. The Table shows four bootstrap experiments, one corresponds to i.i.d. draws of different months of signals. The other three are block bootstraps, which draw three, six, and twelve month blocks, preserving the potential autocorrelation of excess returns.¹⁴ The Table provides *p*-values and confidence intervals for difference in means, difference in standard deviations (Std), difference in skewness (Skew), and difference in kurtosis (Kurt). As expected, the Table shows that the average excess return of the High and Low Basis portfolios are significantly different when evaluated against the bootstrapped distribution (*p*-value < 0.002). The average difference of 14.3% is slightly higher than the difference between the average returns of sorted portfolios in reported Table 6 for the same period (11.57%). More interestingly the difference between the standard deviations of the two distributions is 29% and highly significant (*p*-value < 0.001). The significant difference in the standard error of returns of the two distributions is consistent with the prediction of the theory. It is this higher risk which commands the higher mean, also significantly higher than the Low Basis portfolio. The table also shows that skewness and kurtosis are not significantly different. These results are robust to the time period considered.¹⁵

7. Risk Premiums and the Positions of Traders

It is difficult to reconcile commodity futures risk premiums with traditional asset pricing models, because historical excess returns to commodity futures have low correlations with equities and

¹³ First, construct a matrix where the rows are signals for each commodity (1, 0, -1; where 1 (-1) indicates a month in the highest (lowest) 20% Basis months for that commodity; 0 indicates the remaining months). Columns of the matrix are monthly dates. To construct a bootstrapped distribution of returns, randomly rearrange the columns (or blocks of columns) so that the time dimension of basis signal is scrambled. High and Low Basis return distributions are constructed by applying the scrambled signals.

¹⁴ The sample correspond to the sixteen year period of 1990/12 – 2006/12 (28 commodities). As the bootstrap procedure described above requires no missing data three commodities milk, butter and coal are dropped.

¹⁵ Other sample periods considered were 1959/7 – 2006/12 (9 commodities), 1970/1 – 2006/12 (15 commodities), and 1973/1 – 2006/12 (17 commodities)

aggregate consumption, which are important measures of risk in traditional asset pricing models [e.g., Jagannathan (1985), and Gorton and Rouwenhorst (2006)]. In part for this reason, the prevailing explanation for commodity futures risk premiums in the empirical literature has been hedging-pressure, which is based on the Keynesian Theory of Normal Backwardation. This section re-examines the evidence for the hedging-pressure hypothesis, and tests whether hedging pressure can provide an alternative explanation for the variation of the risk premiums documented in this paper.

In the Keynesian view, the function of commodity futures markets is to enable a risk transfer between hedgers and investors/speculators. The Theory of Normal Backwardation postulates that hedgers are on net short, and offer speculators a risk premium by setting futures prices at a discount relative to expected future spot prices. Academic researchers have tested this prediction by examining the relation between futures returns and “hedging pressure” – defined as the relative size of the short positions taken by hedgers. Empirically, hedging pressure is measured using data on positions of large traders published by the Commodity Futures Trading Commission (CFTC). In the CFTC’s Commitment of Traders Reports large traders are classified as “commercials” or “non-commercials.” The CFTC omits information about the specific identities of traders, but it has become customary in the academic literature to view commercials as hedgers and non-commercial as investors.¹⁶ Several papers, including papers by Carter, Rausser, and Schmitz (1983), Chang (1985), Bessembinder (1992), and DeRoos, Nijman, and Veld (2000), Dincerler, Khokher and Titman (2003) and Dincerler, Khokher and Simin (2004) show that the relative size of the commercial positions is correlated with futures risk premiums.¹⁷

The interpretation of the empirical evidence on hedging pressure is complicated by two issues. First, most papers document a *contemporaneous* correlation between futures prices and traders’ positions. The contemporaneous correlation may simply reflect the response of traders to changes in futures prices, and does not speak to a causal relationship.¹⁸ The first question we ask therefore is whether hedging pressure is correlated with expected *future* commodity risk premiums. Second, these papers treat hedging pressure as exogenous, but it seems reasonable to

¹⁶ In addition the CFTC has a category of “non-reportable positions,” which includes either commercial or non-commercial positions that are below the reporting limits set by the CFTC. These would include either small hedgers or speculators therefore. For the exact definitions see <http://www.cftc.gov/opa/backgrounder/opacot596.htm>. See also Ederington and Lee (2002) for a discussion about the accuracy of the classifications.

¹⁷ See also Van der Goorbergh (2004) and Szymanowska (2006). Bryant, Bessler and Haigh (2006) question the hedging pressure hypothesis.

¹⁸ DeRoos et al (2000) is the only paper to examine the correlation between returns and ex-ante hedging pressure, but we were unable to qualitatively replicate their results. They appear to be studying the contemporaneous correlation.

assume that traders' positions reflect an equilibrium response to demand and supply shocks to physical commodity markets. For example, when a negative supply shock drives down inventories and increases current spot and futures prices, hedgers might find it advantageous to hedge more in equilibrium, despite the fact that the compensation they have to offer to speculators has increased due to increased uncertainty about future spot prices. Therefore, the second question of interest is: if hedging pressure predicts ex-ante risk premiums, to what extent does this reflect an optimal response to fundamental shocks?

Table 11 provides a summary of the net positions of traders. For each commodity we report the average net position by trader category, the standard deviation of the position, the percentage of the months the position is long, as well as the persistence of the position as measured by the first-order autocorrelation coefficient (ρ). All positions are measured as a fraction of the total open interest in that commodity. The first observation about Table 11 is that commercials are on average net short in most markets, while non-commercials and non-reportable positions are on average net long. This is broadly consistent with the Keynesian hypothesis. Exceptions include Corn, Feeder Cattle, Lean Hogs and Milk, where the average position of the commercials is net long. If all short positions were taken up by commercials, their average position would be 100% of the open interest. Instead, the average net short position of commercials across commodities is about 10%, which indicates that commercials are both long and short in a given month. In addition, the table shows that there are large cross-sectional differences in net positions over time: the average standard deviation of the net position of commercials is 15% per month. Also, there are large cross-sectional differences across commodities. For example, commercials in Oats and Platinum are short more than 90% of the months, while the Crude Oil and Corn commercials are almost equally likely to be long or short. Non-Reportable positions in Coffee are always net long, while non-reportable positions in Corn and Feeder Cattle are almost always short. Positions are uniformly persistent for all commodities: the first-order autocorrelations of the positions of commercials range from 0.59 for Coffee to 0.92 for Palladium. It is notable that the non-reportable positions are on average net long in most contracts, and most of the time. A detailed explanation of these differences is beyond the scope of this paper, as our main focus is on the question whether these positions predict risk premiums.

Table 12 summarizes the results of regressions of futures excess returns on hedging pressure, defined as the net long position of commercials scaled by the open interest as in Table 11. Hedging pressure enters this regression either contemporaneously or predictively: in the left columns of each panel the monthly futures return between $t-1$ and t are regressed on the hedging pressure measured at time t , in the right columns hedging pressure is measured at the at time $t-1$.

A negative slope coefficient in the table means that an increase in hedging (decrease of long position) by commercials is associated with a higher futures return. The results in Table 12 show that the slope coefficients are generally significantly negative when hedging pressure is measured at the end of the return interval (i.e., contemporaneously), but insignificantly different from zero when hedging pressure is measured at the beginning of the return interval (i.e., lagged). The R-squared of the predictive regressions is on average below 1%, compared to 10% on average in the coincident regressions. These results are therefore inconsistent with the hypothesis that hedging pressure is an important determinant of ex-ante risk premiums, and consistent with a story that traders adjust their positions as futures prices change.¹⁹ In particular, the significantly negative slope coefficients in the coincident regressions indicate that commercials increase their short positions as prices go up, while non-commercials increase their long positions in a rising market. This would make non-commercials appear to be momentum investors. Indeed, the results in Tables 7 and 8 which summarize the characteristics of portfolios sorted on prior futures or prior spot price returns indicate that non-commercials take larger long positions in high momentum commodities than in commodities with poor prior performance.

The main conclusion of this section is that contrary to the existing literature, we find no evidence that supports a hedging pressure explanation for risk premiums in commodity futures markets. Instead, we have shown that risk premiums systematically vary with the state of inventories, as predicted by the Theory of Storage. Two questions remain. First, does our single factor explanation capture most of the predictable variation of risk premiums? And second: can we reconcile these risk premiums with modern asset pricing theories of risk?

Answering these questions is beyond the scope of the current paper, but we leave the reader with partial answers to these questions. We have shown that portfolios sorted on basis and momentum take positions in low inventories. If we regress the excess returns of the High Basis portfolio on the excess return of the Low Inventories portfolio, we find a significant intercept of 2.4% p.a. ($t = 1.58$). A similar regression of the excess returns of the High Momentum portfolio on the excess returns of the Low Inventories portfolio has an intercept of 4.1% p.a. ($t = 2.27$). This suggests that there is an orthogonal component to the returns to basis and momentum portfolios that is not captured by variation of inventories. This may be due to a combination of noise in the measurement of inventories, or that our inventory sorts do not capture cross-sectional differences in inventory dynamics across commodities that are correlated with risk premiums.

¹⁹ We also conducted a sort of commodities into portfolios based on beginning of period hedging pressure, along the same lines as the portfolio sorts in tables 5-8. Unreported results show that we find no evidence that these sorts were informative about spreading futures risk premiums.

Alternatively it may be due to an omitted risk factor that drives both the returns to basis and prior-return sorted portfolios. If we run a regression of the components of the returns to basis and momentum portfolios that are orthogonal to inventories on each other, we find that the resulting intercept of that regression to be insignificantly different from zero, both economically as well as statistically. This suggests that basis and momentum portfolios contain a common source of risk that is orthogonal to variation in inventories, and is compensated for in average returns. We leave a full exploration of these issues to future research.

8. Summary and Conclusions

This paper examines the relationship between the state of inventories and risk premiums of individual commodity futures, as predicted by the Modern Theory of Storage. For this purpose, we collect a comprehensive historical dataset of inventories for 31 individual commodities over a 37-year period between 1969 and 2006. Our major findings can be summarized as follows. First, consistent with the predictions of the Theory of Storage, we empirically document a negative, non-linear relationship between the futures basis (convenience yield) and the level of inventories: at low inventory levels the basis increases at an increasing rate. Second, we show that the state of inventories is informative about futures risk premiums. Although inventory data suffer from measurement error, we show that commodity futures and spot prices carry relevant information about the state of inventories that can be used to provide additional evidence about the role of inventories for futures risk premiums. In particular we show that prior futures returns, prior spot price changes and the futures basis are correlated with futures risk premiums as predicted by the Theory. We also looked specifically at the causal link between futures risk premiums and the volatility of expected future spot prices. Volatility is higher when inventories are low, resulting in higher futures risk premiums. We confirmed this risk-return tradeoff. Finally, we distinguish our explanation of risk premiums from the Theory of Normal Backwardation, which – in its empirical implementation – attributes risk premiums to “hedging pressure” by Commercial hedgers. While the positions of participants in futures markets vary with both returns and the state of inventories, we find no evidence that the positions of futures traders predict risk premiums on commodity futures.

Appendix: Inventory Data

1. **Aluminum:** Source: London Metals Exchange: “Warehouse stocks.” Start date: 12/29/1978. Periodicity: daily.
2. **Butter:** Source: U.S. Department of Agriculture: “Commercial stocks of butter in the U.S. on first of the month in thousands of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data to 12/31/2004 is compiled in an Excel table by Economic Research Services (ERS-USDA) from National Agricultural Statistics Services (NASS-USDA) data. Then, the data is taken directly from NASS-USDA monthly Cold Storage reports. Data as of the first of the month is shifted to the end of previous month.
3. **Coal:** Department of Energy: Monthly Energy Review: “U.S. Coal stocks.” Start date: 03/31/1999. Periodicity: monthly
4. **Cocoa:** Source: New York Board of Trade, sum of three series: “Visible stocks of cocoa in New York warehouses” (thousands of bags), same for Philadelphia warehouses and for Port of Hampton Road warehouses. Data to 04/30/1999 is 1000 times the monthly series of the same data compiled by Commodity Research Bureau (CRB Yearbooks CD) in millions of bags and rounded to one decimal place. Start date: 1/31/1931. Periodicity: monthly.
5. **Coffee:** Source: New York Board of Trade: “Exchange warehouse stocks, 60 kg bags.” Start date: 1/31/1983. Periodicity: monthly.
6. **Copper:** Source: London Metals Exchange. Start date: 1/2/1970. Periodicity: daily.
7. **Corn:** Source: USDA Livestock and Seed Division, Portland OR. “Stocks of Grain at Selected Terminals and Elevator Sites, thousands of bushels.” Start date: 06/25/1974. Periodicity: weekly.
8. **Cotton:** Source: New York Board of Trade: “Certificated Warehouse Stocks, 480 lb bales.” Start date: 12/31/1989. Periodicity: weekly.
9. **Crude Oil:** Source: Department of Energy: Monthly Energy Review “U.S. crude oil ending stocks non-SPR, thousands of barrels.” Start date: 1/31/1945. Periodicity: monthly.
10. **Feeder Cattle:** For Feeder Cattle, we do not use the available inventory series which is quarterly. Instead we use 3-month-ahead values of the Live Cattle inventory for the current monthly level of Feeder Cattle, under the assumption that it takes three months to feed calves to create what are called Feeder Cattle. Source: U.S. Department of Agriculture: “Cold storage holdings of frozen beef in the U.S. on first of the month in thousands of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data to 12/31/2004 is compiled in an Excel table by Economic Research Services (ERS-USDA) from National Agricultural Statistics Services (NASS-USDA) data. Then, the data is taken directly from NASS-USDA monthly Cold Storage reports. Data as of the first of the month is shifted to the end of previous month and then shifted 3 months forward to account for the average time feeder cattle spends in feedlots.
11. **Heating Oil:** Department of Energy: Monthly Energy Review “U.S. total distillate stocks.” Start date: 1/31/1945. Periodicity: monthly.

12. **Lead:** Source: London Metals Exchange: “Warehouse stocks.” Start date: 1/2/1970. Periodicity: daily.
13. **Lean Hogs:** Source: U.S. Department of Agriculture: “Cold storage holdings of frozen pork in the U.S. on first of the month in thousands of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data to 12/31/2004 is compiled in an Excel table by Economic Research Services (ERS-USDA) from National Agricultural Statistics Services (NASS-USDA) data. Then, the data is taken directly from NASS-USDA monthly Cold Storage reports. Data as of the first of the month is shifted to the end of previous month.
14. **Live Cattle:** Source: U.S. Department of Agriculture: “Cold storage holdings of frozen beef in the U.S. on first of the month in thousands of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data to 12/31/2004 is compiled in an Excel table by Economic Research Services (ERS-USDA) from National Agricultural Statistics Services (NASS-USDA) data. Then, the data is taken directly from NASS-USDA monthly Cold Storage reports. Data as of the first of the month is shifted to the end of previous month.
15. **Lumber:** American Forest & Paper Association: “Stocks (gross) of softwood in the United States, on the first of the month, in millions of board feet.” Data compiled by Commodity Research Bureau (CRB Yearbooks CD) and rounded to one decimal place. Start date: 12/31/1969. Periodicity: monthly. Data as of the first of the month is shifted to the end of previous month.
16. **Milk:** Source: U.S. Department of Agriculture: “Commercial stocks of milk in the U.S. on first of the month in thousands of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data compiled in an Excel table by Economic Research Services (ERS-USDA) from National Agricultural Statistics Services (NASS-USDA) data. Data as of the first of the month is shifted to the end of previous month.
17. **Natural Gas:** Department of Energy: Monthly Energy Review: “U.S. total natural gas in underground storage (working gas), millions of cubic feet.” Start date: 9/30/1975. Periodicity: monthly.
18. **Nickel:** Source: London Metals Exchange: “Warehouse stocks.” Start date: 7/13/1979. Periodicity: daily.
19. **Oats:** Source: USDA Livestock and Seed Division, Portland OR. “Stocks of Grain at Selected Terminals and Elevator Sites, thousands of bushels.” Start date: 06/25/1974. Periodicity: weekly.
20. **Orange Juice:** Source: U.S. Department of Agriculture: “Cold storage stocks of orange juice concentrate in the U.S., millions of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data to 12/31/2004 is compiled by Commodity Research Bureau (CRB Yearbooks CD) from National Agricultural Statistics Services (NASS-USDA) data and rounded to 1 decimal place. Then, the data is taken directly from NASS-USDA monthly Cold Storage reports. Data as of the first of the month is shifted to the end of previous month.
21. **Palladium:** Source: New York Mercantile Exchange: “Warehouse stocks.” Start date: 10/31/1995. Periodicity: daily.

22. **Platinum:** Source: New York Mercantile Exchange: “Warehouse stocks.” Start date: 10/31/1995. Periodicity: daily.
23. **Pork Bellies:** Source: U.S. Department of Agriculture: “Frozen pork belly storage stocks in the United States, on first of the month, in thousands of pounds.” Start date: 12/31/1969. Periodicity: monthly. Data to 12/31/2004 is compiled in an Excel table by Economic Research Services (ERS-USDA) from National Agricultural Statistics Services (NASS-USDA) data. Then, the data is taken directly from NASS-USDA monthly Cold Storage reports. Data as of the first of the month is shifted to the end of previous month.
24. **Propane:** Source: Department of Energy: Monthly Energy Review: “U.S. ending stocks of Propane/Propylene, thousands of barrels.” Start date: 1/31/1971. Periodicity: monthly.
25. **Soybeans:** Source: USDA Livestock and Seed Division, Portland OR. “Stocks of Grain at Selected Terminals and Elevator Sites, Thousands of Bushels”. Start date: 06/25/1974. Periodicity: weekly. From 12/31/1961 to 05/31/1974 1000 * monthly series of the same data compiled (in millions of bushels) by Commodity Research Bureau (CRB Yearbooks CD) and rounded to 1 decimal place as “Commercial stocks of soybeans in the United States, on the first month, in millions of bushels.” and shifted to the end of previous month.
26. **Soybean Oil:** Source: Economic Research Services, U.S. Department of Agriculture: “Stocks of soybean oil (crude and refined) at factories and warehouses in the United States on the first of month, millions of pounds. Data compiled by Commodity Research Bureau (CRB Yearbooks CD) and rounded to 1 decimal place. Start date: 9/30/1970. Periodicity: monthly. Data as of the first of the month is shifted to the end of previous month.
27. **Soybean Meal:** Source: Economic Research Services, U.S. Department of Agriculture: “Stocks, including mill feed and lecithin of soybean cake and meal at oil mills in the United States on the first of the month in thousands of short tons.” Data compiled by Commodity Research Bureau (CRB Yearbooks CD) and rounded to 1 decimal place. Start date: 9/30/1970. Periodicity: monthly. Data as of the first of the month is shifted to the end of previous month.
28. **Tin:** Source: London Metals Exchange: “Warehouse stocks.” Start date: 1/2/1970. Periodicity: daily. Gap in the data from Jan. 1986 – 6/30/1989 during the suspension of trading due to tin crisis. Contract resumed trading in June 1989, but it took another 12 months or so for warehouse stocks to rise from extremely low levels. We only used data from 6/30/1990.
29. **Unleaded Gas:** Source: Department of Energy: Monthly Energy Review: “U.S. total motor gasoline ending stocks, thousands of barrels.” Start date: 3/31/1981. Periodicity: monthly.
30. **Wheat:** Source: USDA Livestock and Seed Division, Portland OR. “Stocks of Grain at Selected Terminals and Elevator Sites, Thousands of Bushels”. Start date: 06/25/1974. Periodicity: weekly. From 06/30/1970 to 05/31/1974 1000 * monthly series of the same data compiled (in millions of bushels) by Commodity Research Bureau (CRB Yearbooks CD) and rounded to 1 decimal place as “Commercial stocks of domestic wheat in the United States, on the first month, in millions of bushels of domestic wheat in storage in public and private elevators in 39 markets and wheat afloat in vessels or barges at lake and seaboard ports, the first Saturday of the month.” and shifted to the end of previous month.

31. **Zinc:** Source: London Metals Exchange: "Warehouse stocks." Start date: 1/2/1970.
Periodicity: daily.

References

- Bessembinder, Hendrik (1992), "Systematic Risk, Hedging Pressure, and Risk Premiums in Futures Markets," *Review of Financial Studies* 5: 637-667.
- Bodie, Zvi and Victor Rosansky (1980), "Risk and Return in Commodity Futures," *Financial Analysts Journal* 36: 27-39.
- Brennan, Michael (1958), "The Supply of Storage," *American Economic Review* 48: 50-72.
- Bryant, Henry, David Bessler, and Michael Haigh (2006), "Causality in Future Markets," *Journal of Futures Markets* 26: 1039 – 1057.
- Carter, Colin (1999), "Commodity futures markets: A Survey," *The Australian Journal of Agricultural and Resource Economics* 43: 209-247.
- Carter, Colin, Gordon Rausser, and Andrew Schmitz (1983), "Efficient Asset Portfolios and the Theory of Normal Backwardation," *Journal of Political Economy* 91: 319-331.
- Chang, Eric (1985), "Returns to Speculators and the Theory of Normal Backwardation," *Journal of Finance* 40: 193-208.
- Cootner, P. (1960), "Returns to Speculators: Telser versus Keynes," *Journal of Political Economy* 68: 396-404.
- Cootner, P. (1967), "Speculation and Hedging," *Food Research Studies*, supplement to vol. 7: 63-106.
- Deaton, Angus and Guy Laroque (1992), "On the Behavior of Commodity Prices," *Review of Economic Studies* 59: 1-23.
- DeRoos, Frans, Theo Niiman, and Chris Veld (2000), "Hedging Pressure Effects in Futures Markets," *Journal of Finance* 55: 1437-1456.
- Dusak, Katherine (1973), "Futures Trading and Investor Returns: An Investigation of Commodity market Risk Premiums," *Journal of Political Economy* 81: 1387-1406.
- Dinceler, Cantekin, Zeigham Khoker, and Timothy Simin (2004), "The Convenience Yield and Risk Premia of Storage," University of Western Ontario, working paper.
- Dinceler, Cantekin, Zeigham Khoker, and Sheridan Titman (2003), "Futures Premia and Storage," University of Western Ontario, working paper.
- Ederington, Louis and Jae Ha Lee (2001), "Who Trades Futures and How: Evidence From the Heating Oil Futures Market," *Journal of Business* 75: 353 – 373.
- Erb, Claude and Campbell Harvey (2006), "The Strategic and Tactical Value of Commodity Futures," *Financial Analysts Journal* 62: 69-97.
- Fama, Eugene and Kenneth French (1988), "Business Cycles and the Behavior of Metals Prices," *Journal of Finance* 43: 1075-1093.

- Fama, Eugene and Kenneth French (1987), "Commodity Futures Prices: Some Evidence on Forecast Power, Premiums, and the Theory of Storage," *Journal of Business* 60: 55-73.
- Getzy, Chris, Bernadette Minton, and Catherine Schrand (1997), "Why Firms Use Currency Derivatives," *Journal of Finance* 52: 1323-1354.
- Gorton Gary, and K. Geert Rouwenhorst (2005), "A Note on Erb and Harvey (2005)," working paper Yale School of Management
- Gorton, Gary and K. Geert Rouwenhorst (2006), "Facts and Fantasies about Commodity Futures," *Financial Analysts Journal* 62: 47- 68.
- Gorton, Gary, Fumio Hayashi, and K. Geert Rouwenhorst (2007), "The Fundamentals of Commodity Futures Returns," Yale ICF Working Paper No. 07-08.
- Graham, John and Cliff Smith (1999), "Tax Incentives to Hedge," *Journal of Finance* 54: 2241-2262.
- Gray, Roger (1961), "The Search for a Risk Premium," *Journal of Political Economy* 69: 250-260.
- Gray, Roger and D.J.S. Rutledge (1971), "The Economics of Commodity Futures Markets: A Survey," *Review of Marketing and Agricultural Economics* 39: 57-108.
- Green, P.J. and Bernard Silverman (1994), Nonparametric Regression and Generalized Linear Models (Chapman & Hall/CRC; Boca Raton, Florida).
- Hayashi, Fumio (2000), Econometrics (Princeton University Press; Princeton, NJ).
- Hicks, John R. (1939), Value and Capital (Cambridge, England: Oxford University Press).
- Hodrick, Robert and Edward Prescott (1997), "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking* 29: 1-16.
- Jagannathan, Ravi (1985), "An Investigation of Commodity Futures Prices Using the Consumption-Based Intertemporal Capital Asset Pricing Model," *Journal of Finance*, vol. 40: 175-191.
- Kaldor, Nicholas (1939), "Speculation and Economic Stability," *Review of Economic Studies* 7: 1-27.
- Kamara, Avraham (1982), "Issues in Futures Markets: A Survey," *Journal of Futures Markets* 2: 261-294.
- Keynes, John M. (1930), A Treatise on Money, vol. 2 (London: Macmillan).
- Kolb, Robert W. 1992. "Is Normal Backwardation Normal?" *Journal of Futures Markets*, vol. 12, no. 1: 75-91.

- Miffre, Joëlle and Georgios Rallis (2007), "Momentum in Commodity Futures Markets," *Journal of Banking and Finance*, 31 no. 6 (June): 1863- 1886
- Nash, Daniel J. 2001. "Long-term Investing in Commodities." *Global Pensions Quarterly*, Morgan Stanley Global Pensions Group, vol. 4, no.1: 25-31.
- Ng, Victor and Stephen Pirrong (1994), "Fundamentals and Volatility: Storage, Spreads, and the Dynamics of Metals Prices," *Journal of Business*, vol. 67, no. 2 (April): 203-230.
- Pirrong, Craig (2005), "Momentum in Futures Markets," Bauer College of Business, University of Houston, working paper.
- Ravn, Morten and Harald Uhlig (2002), "On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations," *The Review of Economics and Statistics* 84: 371-376.
- Routledge, Bryan, Duane Seppi, and Chester Spatt (2000), "Equilibrium Forward Curves for Commodities," *Journal of Finance* 55: 1297-1338.
- Samuelson, Paul A. (1965), "Proof that Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review* 6 (Spring): 41-49.
- Shen, Qian, Andrew C. Szakmary and Subhash C. Sharma (2007), "An examination of momentum strategies in commodity futures markets," *Journal of Futures Markets* 26 (3): 227-256.
- Szymanowska, Marta (2006), "Essays on Rational Asset Pricing," PhD Dissertation, University of Tilburg.
- Telser, Lester G. (2000), Classic Futures: Lessons from the Past for the Electronic Age (London, U.K.: Risk Books).
- Telser, Lester (1958), "Futures Trading and the Storage of Cotton and Wheat," *Journal of Political Economy* 66: 233-55.
- Van der Goorbergh, Rob (2004), "Essays on Optimal Hedging, and Investment Strategies, and on derivatives Pricing," PhD Dissertation, University of Tilburg.
- Williams, Jeffrey and Brian Wright (1991), Storage and Commodity Markets (Cambridge University Press).
- Working, Holbrook (1949), "The Theory of the Price of Storage," *American Economic Review* 39: 1254-62.

Table 1: Summary of Commodity Futures Returns 1969/12-2006/12

The table summarizes the average excess returns to individual commodity futures, expressed as percent per annum. Column 2 gives the number of monthly observations in the sample, followed by the arithmetic and geometric average returns. The next columns give the annualized volatility (defined as the standard deviation of monthly excess return multiplied by the square root of 12), skewness, and kurtosis of the monthly return distribution, followed by the average pair-wise correlation with the other sample commodities and an equally-weighted index that includes all commodities. The final column gives the average futures basis, measured as the percentage difference between the nearest and next-to-maturity futures contracts and expressed as a percent per annum.

Commodity Group	Commodity	Start	N	Arithm Mean	Geom. Mean	Stdev	Skew	Kurt	Corr w/ others	Corr w/ Index	Avg Basis
Index	Index	1969/12	444	5.48	4.58	13.5	0.91	8.71	0.40	1.00	-2.10
Metals	Copper	1969/12	444	7.77	4.06	27.7	0.77	6.10	0.19	0.40	0.37
	Platinum	1995/10	134	12.82	11.28	17.6	-0.16	2.83	0.12	0.33	2.59
	Palladium	1995/10	134	13.41	6.76	37.2	0.60	4.89	0.09	0.25	-0.11
	Zinc	1989/1	215	3.77	1.17	23.2	0.71	4.21	0.14	0.35	-3.51
	Lead	1989/1	215	4.90	2.18	23.7	0.69	4.05	0.11	0.28	-3.22
	Nickel	1988/2	226	16.65	9.37	41.5	3.06	25.99	0.14	0.40	2.67
	Aluminum	1988/6	222	-2.06	-3.93	19.3	-0.11	4.03	0.16	0.41	-4.09
	Tin	1990/7	197	4.11	2.39	18.9	0.88	5.76	0.14	0.37	-1.54
Softs	Cotton	1989/12	204	-4.10	-7.64	26.7	0.27	3.46	0.08	0.22	-7.30
	Cocoa	1969/12	444	5.94	0.89	32.5	0.83	4.62	0.07	0.27	-1.86
	Sugar	1989/12	204	6.78	2.71	28.6	0.09	3.20	0.07	0.20	3.18
	OJ	1969/12	444	4.58	0.11	31.3	2.01	14.87	0.03	0.16	-3.97
	Lumber	1970/10	434	0.50	-3.95	30.0	0.42	4.34	0.09	0.19	-7.61
	Coffee	1983/1	287	2.17	-4.95	38.9	1.09	5.71	0.07	0.21	-5.83
Grains	Wheat	1970/6	438	-0.80	-3.84	25.0	0.79	5.91	0.15	0.52	-5.02
	Corn	1974/6	390	-5.42	-8.06	23.4	1.12	9.34	0.18	0.57	-9.76
	Soybeans	1969/12	444	3.31	-0.52	28.6	1.70	13.54	0.20	0.70	-2.80
	Soybean Oil	1970/9	435	7.49	2.46	33.1	1.64	9.79	0.16	0.60	0.02
	Soybean Meal	1970/9	435	6.80	1.46	34.6	2.55	21.08	0.18	0.65	-0.72
	Oats	1974/6	390	-2.02	-6.77	32.6	2.83	27.67	0.13	0.49	-7.91
	Rough Rice	1987/8	232	-6.35	-10.48	29.4	1.19	8.40	0.07	0.18	-13.06
Meats	Pork Bellies	1969/12	444	1.77	-5.36	38.3	0.58	4.41	0.13	0.44	4.58
	Live Cattle	1969/12	444	6.37	4.61	18.7	-0.23	4.49	0.14	0.44	1.33
	Lean Hogs	1969/12	444	7.54	3.81	27.4	0.15	4.35	0.15	0.50	-5.86
	Feeder Cattle	1972/11	409	2.87	1.37	17.2	-0.53	5.74	0.08	0.32	1.55
	Milk	1997/1	119	5.14	1.34	28.2	0.87	5.46	0.05	0.12	1.16
	Butter	1997/9	111	11.03	4.78	36.1	0.68	5.40	0.03	0.11	-7.18
Energies	Heating Oil	1979/11	325	9.00	4.44	30.7	0.62	4.66	0.15	0.45	2.24
	Crude Oil	1984/3	273	14.47	8.98	33.7	0.60	5.83	0.13	0.47	5.81
	Gasoline	1985/12	252	18.35	12.24	36.0	0.90	5.71	0.14	0.51	7.85
	Propane	1988/8	220	27.03	17.56	48.7	3.95	37.52	0.14	0.52	9.41
	Natural Gas	1990/12	192	8.67	-5.63	54.5	0.55	3.74	0.10	0.42	-15.66
	Coal	2002/7	53	-2.53	-4.05	17.6	0.33	2.61	0.17	0.49	-8.45

Table 2: Inventories and Seasonality

The table summarizes results from a regression of de-trended inventories on monthly dummies. De-trended inventories are defined as $\log(I/I^*)$ where $\log(I)$ is the log level of inventories and $\log(I^*)$ is the fitted value of applying a Hodrick-Prescott filter to the log of inventories. The second column gives the start of the time series of inventories for each commodity; the end of the sample is December 2006. Subsequent columns give the estimated dummy coefficients in percents, and the R-squared of the regression. The final column gives the first-order autocorrelation of monthly de-trended inventories

Commodity Group	Commodity Name	Start	N	Coefficients of Monthly Dummies												R-sq	rho
				Jan	Feb	March	Apr	May	June	July	Aug	Sept	Oct	Nov	Dec		
Metals	Copper	1969/12	445	4.9	2.4	-6.4	-4.5	-5.3	-10.3	-6.4	4.5	8.9	4.4	3.6	4.2	0.01	0.98
	Platinum	1995/10	135	-6.4	-10.4	-5.6	-3.5	-13.3	-15.7	19.3	18.1	20.8	2.7	-3.1	-2.6	0.05	0.86
	Palladium	1995/10	135	-5.3	9.4	14.0	-9.4	-5.7	-1.8	-16.3	-17.2	2.3	-2.1	-1.9	31.6	0.01	0.94
	Zinc	1969/12	445	0.7	-1.0	-2.1	0.4	1.2	-0.1	-0.7	3.3	4.5	0.9	-2.9	-4.1	0.00	0.97
	Lead	1969/12	445	-0.7	-1.9	-4.0	-3.1	-2.7	2.1	4.4	3.3	4.7	2.3	0.2	-4.5	0.00	0.97
	Nickel	1980/4	321	10.4	2.8	0.2	5.5	8.1	-4.1	-8.7	-9.6	-2.8	-5.4	-3.1	7.0	0.00	0.96
	Aluminum	1979/12	325	5.0	3.3	1.5	0.2	-2.9	-2.3	-4.2	-2.3	-0.5	-0.8	3.4	-0.2	0.00	0.98
	Tin	1990/6	199	7.9	2.0	1.9	-1.3	-4.4	-7.0	-5.0	0.8	-1.0	-3.7	-0.1	10.3	0.01	0.95
Softs	Cotton	1989/12	205	-3.8	20.7	34.7	34.4	28.3	16.3	7.9	-45.6	-57.2	-34.1	-8.8	6.7	0.12	0.81
	Cocoa	1969/12	445	-9.4	-6.2	0.0	5.6	14.3	15.9	19.7	9.8	2.5	-16.1	-23.1	-12.6	0.04	0.94
	OJ	1969/12	445	-1.8	10.3	11.7	18.7	26.8	24.9	15.4	1.9	-14.1	-29.3	-37.8	-25.9	0.49	0.92
	Lumber	1969/12	445	0.0	1.3	1.8	1.4	-0.9	-0.9	-1.4	-1.4	-0.2	0.7	-0.6	0.1	0.01	0.90
	Coffee	1983/1	288	-10.4	-9.9	6.5	-0.7	13.2	11.6	16.2	10.5	7.4	-16.4	-28.3	0.3	0.01	0.96
Grains	Wheat	1970/6	439	-0.5	-9.3	-17.9	-31.2	-41.9	-19.3	11.4	26.0	30.5	26.3	16.4	6.9	0.34	0.96
	Corn	1974/6	391	30.0	27.3	23.3	12.6	-7.0	-24.8	-39.7	-48.9	-34.2	6.5	29.0	28.4	0.42	0.91
	Soybeans	1969/12	445	37.2	31.6	25.2	11.6	-6.7	-24.6	-47.3	-79.6	-84.8	37.8	52.6	45.9	0.61	0.76
	Soybean Oil	1970/9	436	6.3	6.7	7.4	7.9	7.2	3.4	1.7	-5.2	-13.8	-13.3	-8.0	0.7	0.07	0.97
	Soybean Meal	1970/9	436	5.0	3.4	1.9	4.7	8.2	-2.1	3.4	-14.2	-18.4	-2.5	6.1	4.9	0.07	0.71
	Oats	1974/6	391	8.6	2.6	-3.6	-12.6	-26.6	-32.1	-23.7	1.3	21.3	27.0	21.6	15.4	0.12	0.94
Meats	Pork Bellies	1969/12	445	20.7	23.9	43.8	57.2	61.8	48.9	5.4	-64.6	-121.1	-82.5	-15.0	20.9	0.70	0.87
	Live Cattle	1969/12	445	7.8	4.2	3.5	1.2	-3.0	-5.3	-5.2	-5.8	-4.4	-1.2	2.0	6.0	0.09	0.92
	Lean Hogs	1969/12	445	0.7	3.0	7.8	17.3	16.6	7.5	-4.5	-16.8	-16.3	-8.6	-3.0	-3.6	0.30	0.91
	Feeder Cattle	1969/12	445	1.3	-3.0	-5.3	-5.2	-5.8	-4.4	-1.2	2.0	6.1	7.8	4.1	3.5	0.09	0.92
	Milk	1970/1	444	-6.9	-4.4	-2.9	2.5	9.0	12.1	13.3	8.3	1.6	-5.0	-13.7	-13.9	0.28	0.95
	Butter	1970/1	444	-13.3	0.5	5.2	18.5	29.1	26.1	25.2	13.4	-2.6	-14.3	-40.9	-46.7	0.23	0.91
Energies	Heating Oil	1969/12	445	2.0	-9.4	-19.5	-21.0	-14.9	-7.7	2.6	9.1	13.7	14.9	17.1	12.9	0.50	0.92
	Crude Oil	1969/12	445	-1.5	-1.2	1.0	2.6	3.1	1.7	0.4	-1.1	-2.5	-0.1	0.1	-2.3	0.07	0.94
	Gasoline	1969/12	445	5.5	5.9	2.4	0.6	0.2	-0.8	-1.8	-4.4	-2.0	-3.9	-1.5	-0.1	0.26	0.85
	Propane	1971/1	432	-17.8	-34.7	-36.8	-24.8	-7.5	5.9	16.6	23.2	26.3	24.3	19.9	5.5	0.66	0.90
	Natural Gas	1975/9	376	-10.4	-34.3	-47.3	-38.6	-19.3	-2.5	10.7	21.5	31.1	36.2	32.5	16.7	0.79	0.91
	Coal	1999/3	94	-4.9	-4.5	-0.2	4.5	6.7	5.9	1.4	-2.8	-4.4	-2.0	0.1	-0.8	0.17	0.94

Table 3: Futures Basis and Inventories

The table reports the results of a regression of the futures basis on I/I^* (the ratio of actual to normal inventory level) and monthly dummies, using a cubic spline regression (see Section 3.2 for how I^* is calculated). The sample period for each commodity is the same as in Table 1. The basis is defined as the annualized difference between the nearest to maturity futures price and the next futures price. Columns 2 to 5 report the slope and associated t -statistics of the regression at $I/I^* = 1$ and $I/I^* = 0.75$. The next two columns report the difference in the slopes and a t -value for the difference. The standard errors of the coefficient estimates used for calculating the t -values are calculated using the Newey-West method for correcting error serial correlation with a bandwidth of 12 months. The estimates reported for each commodity group are the slope and t -values when the coefficients of the cubic spline regression are estimated by pooled OLS, which constrains coefficients to be the same across commodities of the same group. The standard errors of the pooled OLS coefficient estimates take into account serial correlation as well as cross-commodity correlation in the error terms. They also take into account the fact that the data is an unbalanced panel, i.e., the starting month differs across commodities. See Appendix C of GHR (2007) for technical details about this joint estimation on an unbalanced panel.

commodity	slope at 1	t	slope at 0.75	t	difference	t	R-sq
Metals group	-0.051	-2.46	-0.118	-6.01	0.067	4.70	
Copper	-0.032	-0.61	-0.153	-2.76	0.121	5.64	0.41
Platinum	-0.034	-1.10	-0.031	-0.93	-0.003	-0.12	0.41
Palladium	-0.045	-1.46	-0.032	-1.26	-0.013	-1.03	0.19
Zinc	-0.019	-0.39	-0.096	-2.22	0.076	3.32	0.32
Lead	-0.146	-2.83	-0.270	-5.57	0.124	4.34	0.54
Nickel	-0.039	-1.06	-0.136	-4.13	0.096	5.95	0.55
Aluminum	-0.057	-1.64	-0.094	-2.86	0.037	2.16	0.25
Tin	-0.001	-0.02	-0.093	-3.03	0.092	5.06	0.40
Softs group	-0.193	-5.65	-0.257	-8.37	0.064	4.93	
Cotton	-0.166	-2.62	-0.244	-3.86	0.078	3.51	0.30
Cocoa	-0.171	-2.47	-0.273	-3.47	0.102	3.30	0.31
OJ	-0.389	-3.74	-0.347	-2.75	-0.043	-0.45	0.25
Lumber	-1.091	-2.33	-5.466	-1.97	4.375	1.42	0.33
Coffee	-0.092	-2.17	-0.162	-4.45	0.070	4.98	0.62
Grains group	-0.214	-5.10	-0.251	-5.02	0.037	1.39	
Wheat	-0.287	-3.10	-0.453	-3.48	0.165	1.78	0.28
Corn	-0.072	-0.93	-0.200	-1.91	0.128	1.61	0.31
Soybeans	-0.249	-4.06	-0.331	-4.46	0.082	2.47	0.27
Soybean Oil	-0.521	-3.74	-0.716	-3.75	0.195	1.05	0.29
Soybean Meal	0.035	0.31	0.186	1.29	-0.150	-1.03	0.16
Oats	-0.126	-1.14	-0.155	-1.25	0.029	0.45	0.18
Meats group	-0.598	-7.03	-0.602	-6.27	0.004	0.12	
Pork Bellies	-0.358	-5.94	-0.392	-5.31	0.034	1.17	0.46
Live Cattle	-0.434	-2.34	-0.093	-0.16	-0.341	-0.51	0.18
Lean Hogs	-1.223	-5.34	-0.644	-1.24	-0.579	-0.91	0.62
Feeder Cattle	-0.142	-0.92	-0.278	-0.61	0.137	0.26	0.16
Milk	-0.638	-0.64	-8.299	-2.08	7.661	1.67	0.43
Butter	-0.517	-3.74	-0.472	-2.99	-0.045	-0.54	0.35
Energies group	-1.546	-7.61	-1.496	-4.15	-0.050	-0.16	
Heating Oil	-1.376	-6.03	-0.993	-1.18	-0.382	-0.40	0.59
Crude Oil	-3.039	-6.06	16.888	0.90	-19.928	-1.05	0.46
Gasoline	-3.596	-4.11	-17.527	-0.40	13.931	0.32	0.50
Propane	-1.410	-5.21	-1.446	-4.61	0.036	0.13	0.56
Natural Gas	-1.749	-2.94	-1.224	-1.93	-0.525	-0.85	0.55
Coal	-0.649	-1.63	-2.108	-0.15	1.459	0.10	0.45

Table 4: Commodity Excess Return and Inventories

The table reports the results of a regression of monthly percentage excess futures returns on de-trended inventories at the end of the previous month, in addition to monthly dummies. De-trended inventories are defined as I/I^* , the ratio of actual to normal inventory levels. Normal inventories are defined as the fitted values of applying a Hodrick-Prescott filter to inventories. The standard errors of the coefficient estimates used for calculating the t -values are calculated using the Newey-West method for correcting error serial correlation with a bandwidth of 12 months. The estimates reported for each commodity group are the coefficient and t -statistics when coefficients are constrained to be the same. For technical details about this constrained estimation, see Appendix C of GHR (2007).

commodity	coefficient of I/I^*	t	R-sq
Metals group	-0.040	-0.09	
Copper	-0.421	-0.85	0.03
Platinum	-1.071	-1.15	0.10
Palladium	0.767	1.03	0.08
Zinc	-0.398	-0.76	0.04
Lead	-0.504	-0.62	0.03
Nickel	-0.160	-0.28	0.03
Aluminum	0.542	1.58	0.05
Tin	0.779	0.94	0.08
Softs group	-0.240	-0.64	
Cotton	-0.933	-1.42	0.04
Cocoa	-0.345	-0.46	0.03
OJ	-3.347	-1.73	0.06
Lumber	-11.839	-2.52	0.08
Coffee	-0.029	-0.08	0.05
Grains group	-0.773	-1.43	
Wheat	-1.850	-1.47	0.03
Corn	0.444	0.37	0.02
Soybeans	-0.333	-0.29	0.02
Soybean Oil	-1.474	-0.79	0.02
Soybean Meal	-0.687	-0.43	0.02
Oats	-0.751	-0.86	0.01
Meats group	-2.819	-2.22	
Pork Bellies	-2.256	-1.98	0.05
Live Cattle	-2.131	-1.20	0.01
Lean Hogs	-4.262	-2.13	0.05
Feeder Cattle	-3.172	-1.84	0.03
Milk	-11.810	-1.57	0.09
Butter	-3.290	-1.19	0.06
Energies group	-8.706	-1.75	
Heating Oil	-6.608	-1.65	0.08
Crude Oil	-7.152	-0.60	0.06
Gasoline	-4.005	-0.24	0.09
Propane	-10.921	-1.78	0.09
Natural Gas	-10.215	-1.17	0.04
Coal	-8.052	-0.52	0.25

Table 5: Returns and Characteristics of Portfolios Sorted on the Inventories

At the end of each month the available commodities are ranked from high to low using normalized inventories, defined as the ratio of the actual level of inventories (I) divided by the fitted value of HP de-trended inventories (I*). The top half of the commodities are assigned to the High inventory portfolio and the commodities with the lowest basis to the Low inventory portfolio. Panel A of the table summarizes the annualized return distributions of the High and Low portfolios in excess of the equally-weighted (EW) index. Average returns and standard deviations are expressed as percent per annum. The bottom panel summarizes information about the average characteristics of the commodities in the High and Low portfolios, as well as the positions of Traders as defined by the CFTC. Characteristics include: average 12-month futures return prior to portfolio formation, the average 12-month prior % change in spot price prior to portfolio formation, the average percentage basis and normalized inventories at the time of ranking, average normalized inventories expressed as the difference, stated in percents, between the log of actual inventories and HP de-trended inventories, and volatility defined as the average % standard deviation of the daily commodity futures returns during the trading month. Positions of Traders are measures as a percent of Open Interest at the time of sorting. The columns measure the characteristics of the commodities in the High portfolio, the Low portfolio, and the *t*-statistic for the difference. Details of the calculation of the *t*-statistics are in Appendix C of GHR (2007).

	1969/12-2006/12			1986/1-2006/12			1990/12-2006/12		
Panel A: Returns Relative to EW Index									
	High	Low	H-L	High	Low	H-L	High	Low	H-L
Mean	-3.85	4.21	-8.06	-3.64	3.61	-7.25	-4.38	4.37	-8.75
Standard Deviation	7.77	7.80	15.48	7.03	7.04	14.02	6.44	6.47	12.84
<i>t</i> -statistic (mean)	-3.03	3.32	-3.19	-2.34	2.33	-2.34	-2.83	2.80	-2.82
% Excess Return>0	42.57	56.53	43.47	41.04	57.37	42.23	41.67	57.29	43.23
Panel B: Average Portfolio Characteristics									
	High	Low	<i>t</i> -stat	High	Low	<i>t</i> -stat	High	Low	<i>t</i> -stat
Prior 12m return	0.41	15.31	-6.45	1.24	12.97	-5.54	0.05	11.20	-5.43
Prior 12m spot return	6.00	9.78	-2.58	5.00	8.85	-2.39	5.33	8.59	-1.95
Basis	-7.78	4.61	-14.51	-6.86	4.51	-11.40	-8.81	2.79	-13.14
Inventories	36.37	-36.15	47.14	37.20	-35.19	37.01	40.80	-31.07	33.52
Volatility (+1)	23.40	23.86	-1.15	23.75	23.90	-0.27	23.84	23.46	0.66
Commercials				-11.71	-7.97	-5.03	-12.33	-8.00	-4.81
Non-Commercials				5.59	5.28	0.58	6.01	5.66	0.53
Non Reportable				6.08	2.75	5.29	6.27	2.41	5.23

Table 6: Returns and Characteristics of Portfolios Sorted on the Futures Basis

At the end of each month the available commodities are ranked from high to low using the futures basis, defined as the annualized percentage difference between the nearest and next futures price. The top half of the commodities are assigned to the High Basis portfolio and the commodities with the lowest basis to the Low Basis portfolio. Panel A of the table summarizes the annualized return distributions of the High and Low portfolios in excess of the equally-weighted (EW) index. Average returns and standard deviations are expressed as percent per annum. The bottom panel summarizes information about the average characteristics of the commodities in the High and Low portfolios, as well as the positions of Traders as defined by the CFTC. Characteristics include: average 12-month futures return prior to portfolio formation, the average 12-month prior % change in spot price prior to portfolio formation, the average percentage basis and normalized inventories at the time of ranking, average normalized inventories expressed as percentage difference between the log of actual inventories and HP de-trended inventories, and volatility defined as the average % standard deviation of the daily commodity futures returns during the trading month. Positions of Traders are measured as a percent of Open Interest at the time of sorting. The columns measure the characteristics of the commodities in the High portfolio, the Low portfolio, and the t -statistic for the difference. Details of the calculation of the t -statistics are in Appendix C of GHR (2007).

	1969/12-2006/12			1986/1-2006/12			1990/12-2006/12		
Panel A: Returns Relative to EW Index									
	High	Low	H-L	High	Low	H-L	High	Low	H-L
Mean	5.42	-4.82	10.23	5.04	-4.70	9.74	5.71	-5.86	11.57
Standard Deviation	7.76	7.93	15.58	6.87	7.13	13.93	6.08	6.08	12.10
t -statistic (mean)	3.98	-3.44	3.73	3.55	-3.14	3.36	4.04	-4.10	4.08
% Excess Return>0	58.56	42.79	57.88	61.35	39.04	61.35	63.02	37.50	63.02
Panel B: Average Portfolio Characteristics									
	High	Low	t -stat	High	Low	t -stat	High	Low	t -stat
Prior 12m return	21.02	-5.11	12.93	19.68	-5.40	12.99	17.50	-5.93	10.56
Prior 12m spot return	15.61	0.29	10.45	14.39	-0.51	9.51	14.11	0.00	7.16
Basis	15.32	-18.40	46.42	15.44	-17.73	42.10	13.04	-19.01	36.83
Inventories	-14.87	15.31	-17.08	-13.78	15.95	-13.65	-9.34	19.09	-13.76
Volatility (+1)	24.07	23.23	2.13	24.30	23.31	1.72	23.98	23.30	0.99
Commercials				-8.94	-10.34	1.46	-9.87	-10.01	0.13
Non-Commercials				6.89	3.95	4.24	7.78	3.92	4.81
Non Reportable				2.38	6.12	-7.00	2.52	5.73	-5.99

Table 7: Returns and Characteristics of Portfolios Sorted on the Prior 12-month Futures Return

At the end of each month the available commodities are ranked from high to low using prior 12-month futures return. The top half of the commodities are assigned to the High momentum portfolio and the commodities with the lowest basis to the Low Momentum portfolio. Panel A of the table summarizes the annualized return distributions of the High and Low portfolios in excess of the equally-weighted (EW) index. Average returns and standard deviations are expressed as percent per annum. The bottom panel summarizes information about the average characteristics of the commodities in the High and Low portfolios, as well as the positions of Traders as defined by the CFTC. Characteristics include: average 12-month futures return prior to portfolio formation, the average 12-month prior % change in spot price prior to portfolio formation, the average percentage basis and normalized inventories at the time of ranking, average normalized inventories expressed as percentage difference between the log of actual inventories and HP de-trended inventories, and volatility defined as the average % standard deviation of the daily commodity futures returns during the trading month. Positions of Traders are measures as a percent of Open Interest at the time of sorting. The columns measure the characteristics of the commodities in the High portfolio, the Low portfolio, and the t -statistic for the difference. Details of the calculation of the t -statistics are in Appendix C of GHR (2007).

	1969/12-2006/12			1986/1-2006/12			1990/12-2006/12		
Panel A: Returns Relative to EW Index									
	High	Low	H-L	High	Low	H-L	High	Low	H-L
Mean	6.54	-6.82	13.36	6.81	-7.03	13.84	7.69	-7.67	15.36
Standard Deviation	8.52	8.62	16.99	7.80	7.90	15.53	6.84	6.83	13.64
t -statistic (mean)	4.82	-4.95	4.93	4.24	-4.35	4.34	4.56	-4.62	4.60
% Excess Return>0	58.78	42.34	58.11	61.35	39.44	60.96	64.58	35.42	64.58
Panel B: Average Portfolio Characteristics									
	High	Low	t -stat	High	Low	t -stat	High	Low	t -stat
Prior 12m return	32.62	-16.65	21.34	31.57	-17.14	27.03	29.40	-17.79	25.54
Prior 12m spot return	26.22	-10.43	23.52	25.54	-11.70	24.16	25.37	-11.23	20.33
Basis	6.73	-9.96	19.15	6.94	-9.30	17.97	5.03	-11.08	14.73
Inventories	-9.30	9.88	-8.26	-7.29	9.44	-6.07	-3.51	13.29	-5.74
Volatility (+1)	24.10	23.28	1.71	24.43	23.24	1.83	24.37	22.97	1.83
Commercials				-11.57	-8.01	-2.73	-12.53	-7.46	-3.61
Non-Commercials				9.02	1.58	9.81	10.11	1.24	11.72
Non Reportable				2.74	6.18	-4.31	2.67	5.89	-3.67

Table 8: Returns and Characteristics of Portfolios Sorted on the Prior 12-month Spot Return

At the end of each month the available commodities are ranked from high to low using prior 12-month spot return, defined as the percentage change in the spot price. The top half of the commodities are assigned to the High momentum portfolio and the commodities with the lowest basis to the Low Momentum portfolio. Panel A of the table summarizes the annualized return distributions of the High and Low portfolios in excess of the equally-weighted index (EW) index. Average returns and standard deviations are expressed as percent per annum. The bottom panel summarizes information about the average characteristics of the commodities in the High and Low portfolios, as well as the positions of Traders as defined by the CFTC. Characteristics include: average 12-month futures return prior to portfolio formation, the average 12-month prior % change in spot price prior to portfolio formation, the average percentage basis and normalized inventories at the time of ranking, average normalized inventories expressed as percentage difference between the log of actual inventories and HP de-trended inventories, and volatility defined as the average % standard deviation of the daily commodity futures returns during the trading month. Positions of Traders are measures as a percent of Open Interest at the time of sorting. The columns measure the characteristics of the commodities in the High portfolio, the Low portfolio, and the t -statistic for the difference. Details of the calculation of the t -statistics are in Appendix C of GHR (2007).

	1969/12-2006/12			1986/1-2006/12			1990/12-2006/12		
Panel A: Returns Relative to EW Index									
	High	Low	H-L	High	Low	H-L	High	Low	H-L
Mean	6.73	-7.12	13.85	8.55	-8.82	17.37	7.87	-8.16	16.03
Standard Deviation	8.69	8.58	17.19	8.53	8.34	16.83	6.71	6.78	13.44
t -statistic (mean)	4.77	-5.09	4.95	4.79	-5.07	4.94	4.36	-4.55	4.47
% Excess Return>0	56.76	41.67	57.88	59.76	38.25	60.96	61.46	36.98	61.98
Panel B: Average Portfolio Characteristics									
	High	Low	t -stat	High	Low	t -stat	High	Low	t -stat
Prior 12m return	28.61	-12.79	18.13	27.98	-13.58	22.84	25.99	-14.41	20.56
Prior 12m spot return	29.78	-13.87	28.85	28.60	-14.67	28.66	28.15	-14.01	24.84
Basis	3.94	-7.08	11.57	4.71	-7.05	12.40	3.00	-9.00	10.22
Inventories	-3.00	3.27	-2.77	-2.25	4.09	-2.51	1.56	8.17	-2.57
Volatility (+1)	24.18	23.25	1.82	24.35	23.33	1.40	24.43	22.90	1.91
Commercials				-13.02	-6.45	-6.29	-14.03	-5.83	-7.80
Non-Commercials				9.60	1.14	13.68	10.59	0.95	16.12
Non Reportable				3.68	5.04	-1.87	3.78	4.53	-0.93

Table 9: Average Commodity Return Volatility of Sorted Portfolios

At the end of each month the available commodities are ranked from high to low using various characteristics: normalized inventories, futures basis, prior 12-month futures returns and prior 12-month spot return. The table reports the t -statistic of the null hypothesis that the difference between the volatility of the commodities in the High portfolio and the volatility of the commodities in the Low Characteristic Portfolios is equal to zero. In panel A, the time series mean of is removed from the time series of all commodity returns and characteristics before conducting the portfolio sorts. Panel B collects the t -statistics of a similar test on the raw data, which are reported in Tables 5-9.

Characteristic	1969/12-2006/12	1986/1-2006/12	1990/12-2006/12
Panel A: De-meaned data			
Inventories	-1.75	-1.59	-0.78
Basis	3.33	5.61	4.68
Prior 12-month futures return	2.77	3.42	3.49
Prior 12-month spot return	2.82	2.75	3.07
Panel B: Raw data			
Inventories	-1.15	-0.27	0.66
Basis	2.13	1.72	0.99
Prior 12-month futures return	1.71	1.83	1.83
Prior 12-month spot return	1.82	1.4	1.91

Table 10: Difference between the High Basis and Low Basis Return Distributions

Bootstrap test over sample period: 1990/12 – 2006/12

The table summarizes the bootstrapped tests of significance between the sample moments of the empirical conditional distributions of pooled standardized commodity futures returns. For each commodity of the 28 sample commodities (31 commodities, excluding Milk, Butter, and Coal), standardized futures returns are calculated by re-normalizing the time series of excess returns to have zero mean and unit standard deviation. Next, the standardized returns are divided into quintiles corresponding to the highest (lowest) beginning of month basis. Finally the lowest (highest) quintile observations of each of the commodity are assigned to the Low (High) Basis empirical distribution. These two distributions are illustrated in Figure 5. The Table evaluates the difference between the observed sample moments of the two conditional distributions against those of bootstrapped distributions under the 4 different assumptions. The bootstrapped distributions are constructed by randomizing the date assignments of signals (or blocks of signals) that are used to construct the Basis sorted portfolios. Block bootstraps of 3, 6, and 12 months are used to capture the time series properties of the signals. The p -value measures the probability of obtaining a larger value than the sample statistic under the bootstrapped distribution.

	Moment Difference between High and Low conditional distributions			
	Average	Standard Deviation	Skewness	Kurtosis
Sample Statistic	14.372	29.134	-0.0231	0.6864
i.i.d. Bootstrap				
p-value	0.002	<0.001	0.543	0.18
90% Confidence Level	[-8.009, 7.980]	[-7.297, 7.348]	[-0.389, 0.385]	[-1.190, 1.227]
Block 3 Month				
p-value	<0.001	<0.001	0.351	0.207
90% Confidence Level	[-7.259, 7.531]	[-7.129, 7.298]	[-0.477, 0.254]	[-1.146, 1.312]
Block 6 Month				
p-value	0.001	<0.001	0.297	0.276
90% Confidence Level	[-6.611, 7.980]	[-7.642, 6.768]	[-0.524, 0.224]	[-0.955, 1.486]
Block 12 Month				
p-value	<0.001	<0.001	0.198	0.284
90% Confidence Level	[-6.675, 6.964]	[-4.979, 6.481]	[-0.573, 0.138]	[-0.874, 1.466]

Table 11: Summary of Positions of Traders 1986 – 2006

The table summarizes the positions of traders in commodity futures markets according to the classifications employed in the Commitment of Traders Report published by the CFTC: For each category (Commercials, Non-Commercials, and Non-Reportables) positions are measured as net long and expressed as a percentage of Open Interest. The columns report the sample average position, the standard deviation of the position, the fraction of the months the position is long, and the first-order autocorrelation (rho) of the position.

		Net Long Positions of Traders as Percent of Open Interest											
		Commercials				Non-Commercials				Non Reportable			
	Commodity	Average	Stdev	%Long	rho	Average	Stdev	%Long	rho	Average	Stdev	%Long	rho
Metals	Copper	-16.67	22.70	26.19	0.76	8.28	17.01	67.86	0.74	8.39	8.42	85.32	0.81
	Platinum	-38.93	24.02	7.14	0.71	23.99	22.00	83.73	0.74	14.94	7.83	97.62	0.79
	Palladium	-30.48	30.15	22.62	0.92	17.33	18.70	76.59	0.88	13.15	14.72	82.14	0.92
Softs	Cotton	-4.02	23.11	42.06	0.71	-1.41	19.93	49.60	0.73	5.42	6.32	83.73	0.76
	Cocoa	-8.77	16.14	28.97	0.78	2.40	12.61	56.35	0.74	6.38	5.74	89.29	0.88
	Sugar	-20.72	21.66	22.62	0.73	9.43	14.85	72.22	0.72	11.30	9.04	90.08	0.77
	Orange Juice	-15.06	25.57	26.19	0.77	6.38	17.41	64.29	0.70	8.68	13.65	83.73	0.86
	Lumber	-10.50	18.62	32.14	0.74	4.57	15.21	66.67	0.62	5.93	12.00	69.84	0.74
	Coffee	-17.41	15.38	16.67	0.59	6.49	13.65	69.84	0.56	10.92	4.76	100.00	0.76
Grains	Wheat	-9.35	15.77	30.95	0.73	4.60	12.74	59.52	0.73	4.75	8.54	68.25	0.80
	Corn	1.01	13.81	51.59	0.76	5.69	10.97	66.27	0.74	-6.70	5.97	11.11	0.83
	Soybeans	-10.73	17.61	27.38	0.87	6.67	12.68	70.24	0.80	4.06	7.68	68.65	0.89
	Soybean Oil	-13.11	18.28	28.97	0.74	5.17	12.94	63.49	0.75	7.94	7.23	87.70	0.72
	Soybean Meal	-13.72	14.89	21.43	0.70	4.67	10.25	67.06	0.70	9.04	5.85	94.05	0.69
	Oats	-37.15	15.92	1.19	0.71	11.95	11.51	90.87	0.77	25.20	13.49	98.02	0.82
	Rough Rice	-7.43	21.14	37.07	0.85	2.72	13.35	53.88	0.83	4.71	13.99	56.90	0.82
Meats	Pork Bellies	-0.84	14.41	43.65	0.76	-1.91	18.82	44.84	0.68	2.75	18.76	53.17	0.80
	Live Cattle	-8.31	11.34	26.98	0.85	8.05	10.25	75.40	0.73	0.26	10.21	48.02	0.88
	Lean Hogs	0.59	12.02	46.83	0.68	5.81	14.47	66.67	0.64	-6.40	7.99	17.46	0.56
	Feeder Cattle	8.79	11.90	75.00	0.75	8.86	12.96	76.19	0.70	-17.65	13.99	14.29	0.87
	Milk	10.94	16.42	76.58	0.85	1.12	10.89	45.05	0.75	-12.06	8.83	4.50	0.75
Energies	Heating Oil	-9.00	9.75	18.65	0.61	1.80	6.26	59.92	0.55	7.20	5.41	90.87	0.72
	Crude Oil	-0.10	8.43	47.62	0.66	0.39	6.28	50.79	0.68	-0.29	3.39	46.83	0.58
	Unleaded Gas	-8.76	11.43	23.81	0.60	6.54	8.58	76.19	0.65	2.22	4.50	73.02	0.38
	Propane	-9.82	11.83	19.74	0.71	-0.61	6.08	28.29	0.71	10.43	10.35	82.24	0.65
	Natural Gas	-7.01	8.22	22.00	0.63	0.76	7.21	56.00	0.65	6.25	3.47	98.00	0.79

Table 12: Hedging Pressure and Futures Returns 1986/12 – 2006/12

The table summarizes the results of a simple regression of futures returns realized at the end of month t on commercial positions measured at time t (contemporaneous) and measured at the end of month $t-1$ (lagged). Commercial Positions are defined as the net long position in a commodity future expressed as a percent of the open interest in that commodity using data obtained from the Report of Traders of the CFTC. The table reports the slope coefficient and the associated t -statistic, and the R-squared of the regression.

		Contemporaneous			Lagged		
	Commodity	slope	t -stat	R-sq	slope	t -stat	R-sq
Metals	Copper	-0.13	-4.95	0.13	-0.02	-0.91	0.00
	Platinum	-0.10	-7.49	0.16	0.00	-0.03	0.00
	Palladium	-0.06	-2.54	0.04	-0.03	-1.22	0.01
Softs	Cotton	-0.16	-8.82	0.22	-0.02	-1.15	0.01
	Cocoa	-0.15	-4.50	0.09	0.00	0.13	0.00
	Sugar	-0.19	-7.24	0.18	0.00	-0.08	0.00
	Orange Juice	-0.07	-3.35	0.05	-0.02	-0.76	0.00
	Lumber	-0.11	-3.70	0.05	-0.03	-0.97	0.00
	Coffee	-0.31	-6.85	0.17	0.04	0.86	0.00
Grains	Wheat	-0.15	-6.12	0.15	0.01	0.49	0.00
	Corn	-0.23	-8.32	0.21	-0.01	-0.30	0.00
	Soybeans	-0.10	-5.19	0.08	0.01	0.27	0.00
	Soybean Oil	-0.16	-8.31	0.18	0.00	0.13	0.00
	Soybean Meal	-0.20	-7.40	0.17	0.00	-0.04	0.00
	Oats	-0.01	-0.23	0.00	0.09	2.17	0.02
	Rough Rice	-0.09	-3.49	0.05	-0.06	-2.43	0.02
Meats	Pork Bellies	0.02	0.35	0.00	0.04	0.79	0.00
	Live Cattle	-0.09	-4.22	0.06	-0.03	-1.38	0.01
	Lean Hogs	-0.19	-5.63	0.09	0.01	0.22	0.00
	Feeder Cattle	-0.03	-1.23	0.01	0.05	2.39	0.02
	Milk	-0.14	-2.41	0.08	-0.08	-1.23	0.02
Energies	Heating Oil	-0.46	-8.22	0.22	-0.05	-0.75	0.00
	Crude Oil	-0.49	-6.73	0.17	-0.12	-1.60	0.01
	Unleaded Gas	-0.35	-6.70	0.15	-0.05	-0.95	0.00
	Propane	0.10	1.03	0.01	-0.12	-1.22	0.01
	Natural Gas	-0.77	-6.13	0.17	-0.10	-0.71	0.00

Figure 1

The figure plots the sample average basis against the sample average collateralized futures return for individual commodity futures between 1990/12 and 2006/12. The basis is measured as the relative price difference between the two closest to maturity contracts, expressed as a percent per annum





Figure 2: Seasonal Variation of Inventories

The figure graphs the estimated coefficients of a regression of log of inventories, measured in deviation from HP filtered inventories, on monthly dummies. Panel A plots the seasonal coefficients for Natural Gas Inventories, Panel B shows the seasonal variation for inventories of Wheat, and Panel C is for Corn.


Panel A

Seasonal Variation of Natural Gas Inventories
Deviation of Inventories from Trend 1975/9-2006/12



Panel B

Seasonal Variation of Wheat Inventories Deviation of Inventories from Trend 1970/6-2006/12



Panel C

Seasonal Variation of Corn Inventories Deviation of Inventories from Trend 1974/6-2006/12


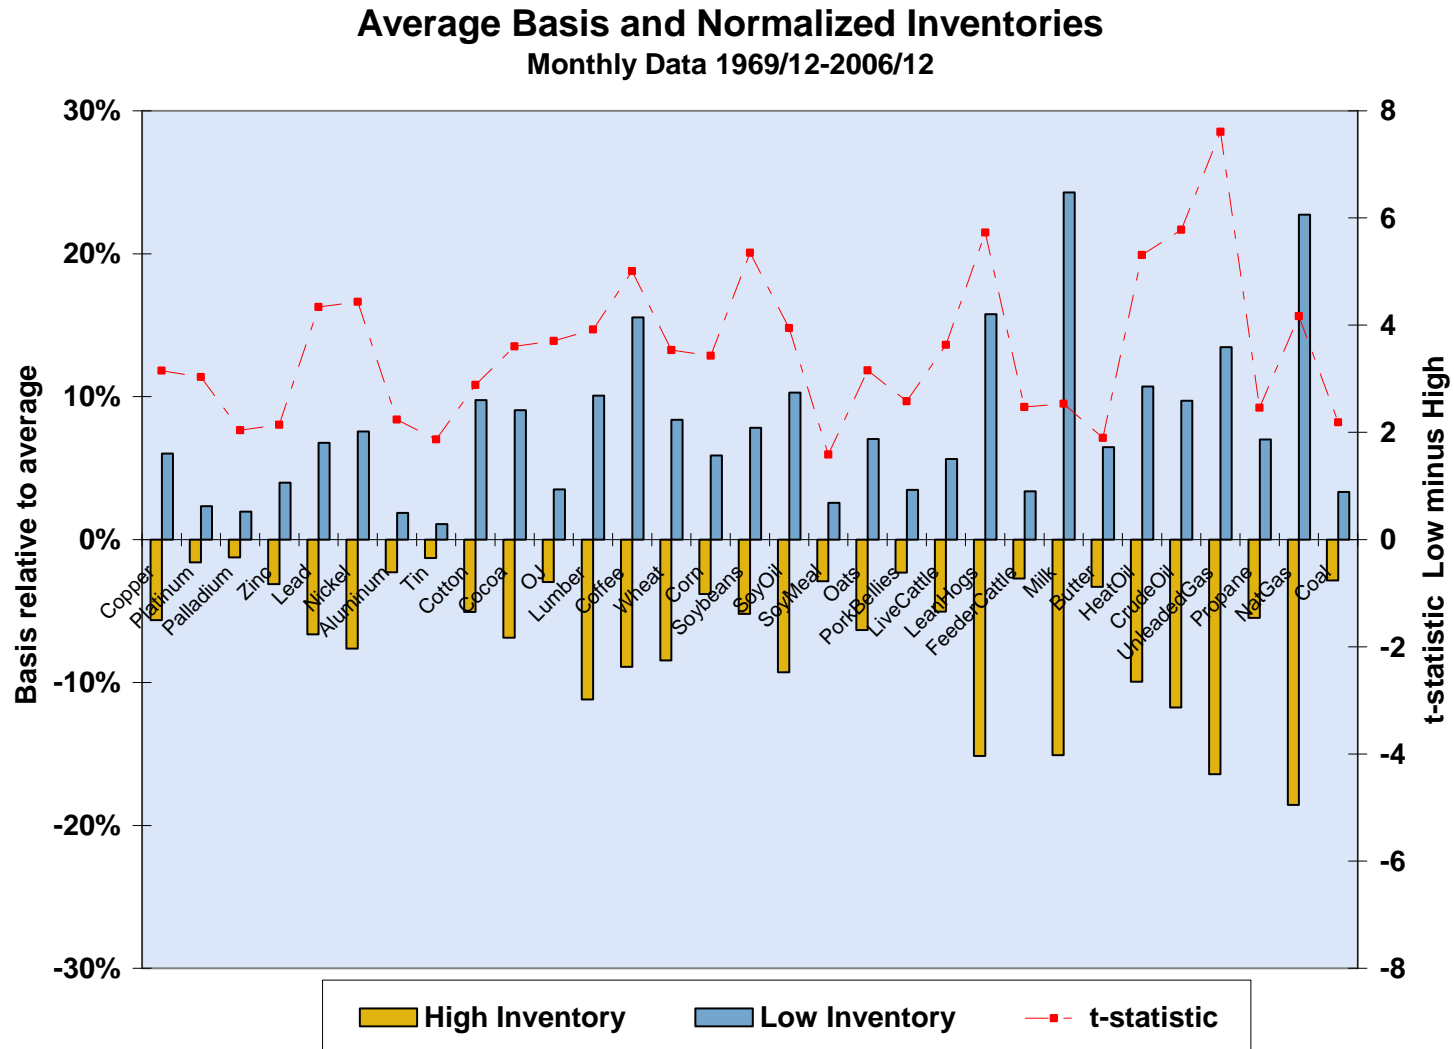


Figure 3

For each commodity we divide the sample in months when actual relative to normalized inventories is above unity (High) and when it is below unity (Low). In Panel A, we plot for each commodity the average basis in High and Low inventory months, expressed in deviation from the full sample mean. In Panel B, we show for each commodity the prior 12-month futures returns in High and Low inventory months, expressed in deviation from the annualized sample average 12-m return. The t-statistics of a test for the difference of the characteristics is given in red.

Panel A



Panel B

Prior 12-month Return and Normalized Inventories
Monthly Data 1969/12-2006/12




Figure 4

The figure shows a scatter plot of the monthly observations of the futures basis against the ratio of inventories relative to trend (I/I^*) for Copper and Crude Oil. The basis is net of seasonal effect, i.e., after subtracting the estimated linear function of monthly dummies in the cubic spline regression. In addition (in red) we give the fitted values of a cubic spline regression of the basis on inventories.

Panel A




Figure 5

The figure shows the distributions of the normalized (to have zero mean and unit standard deviation) excess futures returns for the 20% of the observations with the highest Basis and the 20% of the observations with the lowest Basis.

