

Effects of index-fund investing on commodity futures prices*

James D. Hamilton[†]

Department of Economics

University of California, San Diego

Jing Cynthia Wu[‡]

Booth School of Business

University of Chicago

May 15, 2012

Revised: December 26, 2013

Abstract

The last decade brought substantial increased participation in commodity markets by index funds that maintain long positions in the near futures contracts. Policy makers and academic studies have reached sharply different conclusions about the effects of these funds on commodity futures prices. This paper proposes a unifying framework for examining this question, noting that according to a simple model of futures arbitrage, if index-fund buying influences prices by changing the risk premium, then the notional positions of the index investors should help predict excess returns in these contracts. We find no evidence that the positions of traders in agricultural contracts identified by the CFTC as following an index strategy can help predict returns on the near futures contracts. We review evidence that these positions might help predict changes in oil futures prices, and find that while there is some support for this in the earlier data, this appears to be driven by some of the dramatic features of the 2007-2009 recession, with the relation breaking down out of sample.

*Financial Support from the University of Chicago Booth School of Business is gratefully acknowledged. We thank Jesus Fernandez-Villaverde, Scott Irwin, Lutz Kilian, and anonymous referees for helpful comments on an earlier draft of this paper.

[†]jhamilton@ucsd.edu

[‡]Cynthia.Wu@chicagobooth.edu

1 Introduction

The last decade has seen a phenomenal increased participation by financial investors in commodity futures markets. A typical strategy is to take a long position in a near futures contract, and as the contract nears maturity, sell the position and assume a new long position in the next contract, with the goal being to create an artificial asset that tracks price changes in the underlying commodity. Barclays Capital estimated that exchange traded financial products following such strategies grew from negligible amounts in 2003 to a quarter trillion dollars by 2008 (Irwin and Sanders (2011)). Stoll and Whaley (2010) found that in recent years up to half of the open interest in outstanding agricultural commodity futures contracts was held by institutions characterized by the Commodity Futures Trading Commission (CFTC) as commodity index traders.

This trend has been accompanied by a broad public perception that increased participation by financial institutions in commodity futures markets has made an important contribution to the increase in commodity prices observed since 2004. This position has been championed, for example, by hedge fund manager Michael Masters in testimony before the U.S. Congress (Masters (2008)) and former Congressional Representative Joseph Kennedy (Kennedy (2012)).

Surveys of previous academic studies by Irwin and Sanders (2011) and Fattouh, Kilian, and Mahadeva (forthcoming), as well as our own review in Section 2 below, failed to find much empirical support for these claims. What accounts for their continued prominence in policy discussions? Masters (2009) saw the case as simple and clear-cut:

Buying pressure from Index Speculators overwhelmed selling pressure from pro-

ducers and the result was skyrocketing commodity prices.

This claim involves two separate links: first, that increased volume on the buy side drives up the price of a futures contract, and second, that higher futures prices would be sufficient on their own to produce an increase in spot prices. Possible channels for the second link have been discussed by Hamilton (2009), Kilian and Murphy (forthcoming), Knittel and Pindyck (2013), and Sockin and Xiong (2013). In this paper we focus on the first link— by what mechanism could increased index-fund buying affect the equilibrium price of a futures contract, and how would we look for evidence of such an effect?

After reviewing the previous literature in Section 2, we sketch in Section 3 a simple model in which an increased volume of buy orders could affect futures prices by changing the equilibrium risk premium. We show that an implication of this framework is that regressions similar to those run by previous researchers can be a useful way to measure the impact of index-fund investing on commodity futures prices. We note that according to our theoretical formulation, it would be the notional positions of the index-fund investors rather than the number of contracts or related measures that would help predict log returns if index investing were having a significant effect. In Section 4, we use data on the 12 commodities covered by the Supplemental Commitment of Traders, and find, consistent with most of the earlier literature, that index-fund investing seems to have had little impact on futures prices in these markets. Section 5 examines the evidence on oil markets. We first reproduce some of the findings in Singleton (forthcoming) that are consistent with the claim that index-fund buying affected oil futures prices. We then review the criticisms raised by Irwin and Sanders (2012) about Singleton’s method for inferring the crude oil positions of index-fund traders, and generalize

the method to mitigate some of these criticisms. We find that while such measures still fit in sample, their out-of-sample performance is poor. Our conclusions are summarized in Section 6.

2 Previous literature on the effects of index-fund investing on commodity futures prices.

One piece of evidence sometimes viewed as supportive of the view that financial speculation has played a role in recent commodity price movements is the observation that the correlation between commodity price changes and other financial returns has increased substantially in recent years. Tang and Xiong (2012) found this correlation is stronger among commodities included in the main index funds than for commodities not included, and Buyuksahin and Robe (2010) and Buyuksahin and Robe (2011) related these correlations specifically to micro data on positions of different types of commodity traders. However, as Fattouh, Kilian, and Mahadeva (forthcoming) noted, the increasing correlation could also be attributed to an increasing importance of common factors, such as the growing importance of emerging markets in both commodity markets and global economic activity and the global character of the financial crisis in 2007-2009, factors to which commodity traders would logically respond.

Another approach uses structural vector autoregressions. One common strategy is to interpret a simultaneous unanticipated rise in prices and commodity inventories as reflecting speculative demand pressure. Kilian and Murphy (forthcoming) and Kilian and Lee (forthcoming) concluded that such a model rules out speculative trading as a possible cause of the

2003-2008 surge in oil prices. In related work, Lombardi and van Robays (2011) and Juvenal and Petrella (2011) found only a small role for speculation using alternative specifications. In addition to sensitivity to specification, there is a fundamental identification challenge in using these strategies to distinguish a rise in prices and inventories that results from destabilizing speculation from one that represents a socially optimal response to rationally perceived future market tightness.¹

A third strand in the literature examines whether changes in commodity futures prices could be predicted on the basis of the positions of different types of commodity traders. Here again, the evidence is mostly negative. Brunetti, Buyuksahin, and Harris (2011) used proprietary CFTC data over 2005-2009 on daily positions of traders disaggregated into merchants, manufacturers, floor brokers, swap dealers, and hedge funds. They found that changes in net positions of any of the groups did not help to predict changes in the prices of futures contracts for the three commodities they studied (crude oil, natural gas, and corn). Sanders and Irwin (2011a) used the CFTC's publicly available Disaggregated Commitment of Traders Report on weekly net positions of swap dealers, and found these were no help in predicting returns on 14 different commodity futures contracts over 2006-2009. Sanders and Irwin (2011b) used proprietary CFTC data to extend the public Supplemental Commitment of Traders (SCOT), which categorizes certain participants as commodity index traders, back to 2004. They found that changes in the positions of index traders did not help predict weekly returns for corn or wheat but found some predictability for soybeans under some specifications. Stoll and Whaley (2010) used the public SCOT for 12 agricultural commodities over 2006-2009 and

¹For further discussion see Kilian and Murphy (2012) and Fattouh, Kilian, and Mahadeva (forthcoming).

found that changes in the long positions of commodity index traders predicted weekly returns for cotton contracts but none of the other 11 commodities. Alquist and Gervais (2011) used the public CFTC Commitment of Traders Report to measure net positions of commercial and non-commercial traders, and found that changes in either category could not predict monthly changes in oil prices or the futures-spot spread over 2003-2010, though there was statistically significant predictability when the sample was extended back to 1993. Irwin and Sanders (2012) used the CFTC's Index Investment Data on quarterly positions in 19 commodities held by commodity index funds. They found in a pooled regression that changes in these positions did not predict futures returns over 2008-2011. They also separately analyzed whether changes in futures positions of a particular oil- or gas-specific exchange traded fund could predict daily returns on those contracts over 2006-2011, and again found no predictability. Buyuksahin and Harris (2011) used proprietary CFTC data on daily positions broken down by non-commercials, commercials, swap dealers, hedge funds, and floor broker-dealers. They found the last category could help predict changes in oil futures prices one day ahead, but no predictability for any of the other categories or other horizons. By contrast, Singleton (forthcoming) found that a variety of measures, including a 13-week change in index-fund holdings imputed from the SCOT, could help predict weekly and monthly returns on crude oil futures contracts over September 2006 to January 2010.

As we will see in Section 3 below, it is possible to motivate regressions similar to these from a simple model of risk premia in commodity futures contracts. Keynes (1930) proposed that risk premia in commodity futures prices could arise from the desire of producers of the physical commodity to hedge their price risk by selling futures contracts. In order to persuade

a counterparty to take the other side, the equilibrium price of a futures contract might be pushed below the expected future spot price to produce a situation sometimes described as “normal backwardation.” Evidence on backwardation is mixed. Carter, Rausser, and Schmitz (1983), Chang (1985), Bessembinder (1992), and De Roon, Nijman, and Veld (2000) provided empirical support for such an interpretation of the risk premium in commodity futures, while Marcus (1984), Hartzmark (1987), and Kolb (1992) concluded that it does not hold as a general characterization.

Cootner (1960) argued more generally that hedging pressure (of which the desire by Keynes’s producers to sell forward is one example) could lead to expected returns for positions in futures contracts as a necessary inducement to potential counterparties to take the other side of the contract from the hedgers. Applying this idea to recent developments in commodity futures markets, Brunetti and Reiffen (2011), Acharya, Lochstoer, and Ramadorai (forthcoming), Cheng, Kirilenko, and Xiong (2012), and Hamilton and Wu (forthcoming) proposed that the growing volume of commodity index investors could produce hedging price pressure on the buy side, with Hamilton and Wu (forthcoming) finding that the average compensation to the long position in oil futures contracts has decreased but become substantially more volatile since 2005. Etula (forthcoming), Acharya, Lochstoer, and Ramadorai (forthcoming), Danielsson, Shin, and Zigrand (2011), and Cheng, Kirilenko, and Xiong (2012) have stressed the role of limited working capital on the part of potential arbitrageurs as the key factor determining how much the futures price might deviate from the expected future spot price. In the following section, we follow Hamilton and Wu (forthcoming) in using a simple quadratic objective function as an approximation to a more detailed model of the capital limi-

tations of potential arbitrageurs. Hamilton and Wu (forthcoming) used this framework to infer risk prices from the predictability of futures returns based on their own lagged values. By contrast, in this paper we study the relation between futures returns and lagged observations on the contract positions of commodity-index traders, and show how the framework can be used to motivate and interpret some simple regression tests of the hypothesis that index-fund investing has an independent effect on commodity futures prices.

3 Index-fund investors and the price of risk.

Let F_{nt} denote the price of a commodity associated with an n -period futures contract agreed upon at date t . Entering such a contract requires maintaining a margin account, to which funds must be added if the market moves against the trader and from which funds can be withdrawn if the market moves in the trader's favor. Following (Duffie, 1992, p. 39), one can think of the initial margin deposit as funds the trader would have held in that form in any case. From that perspective, each unit of the commodity purchased through a long position in the contract is associated with zero initial cost and a cash flow at $t + 1$ of $F_{n-1,t+1} - F_{nt}$. If index-fund buyers want to take the long side of the contract, somebody else must be persuaded to take the short side. We will refer to the index fund's counterparty as an "arbitrageur", and assume that what the arbitrageur cares about is the mean and variance of their composite portfolio. Let z_{nt} denote a representative arbitrageur's notional exposure (with $z_{nt} > 0$ denoting a long position and $z_{nt} < 0$ a short); thus for example z_{nt}/F_{nt} is the number of barrels of oil purchased through an n -period contract. An arbitrageur who takes a position

z_{nt} would experience a cash flow at time $t + 1$ given by

$$z_{nt} \left[\frac{F_{n-1,t+1} - F_{nt}}{F_{nt}} \right]. \quad (1)$$

In addition to potential positions in a variety of futures contracts, we presume that the arbitrageur also invests amounts q_{jt} in assets $j = 0, 1, \dots, J$ (where asset $j = 0$ is presumed to be risk free). If the gross return for asset j between t and $t + 1$ is represented as $\exp(r_{j,t+1})$, then the arbitrageur's wealth at $t + 1$ will be

$$W_{t+1} = \sum_{j=0}^J q_{jt} \exp(r_{j,t+1}) + \sum_{n=1}^N z_{nt} \frac{F_{n-1,t+1} - F_{nt}}{F_{nt}}.$$

We assume that the arbitrageur chooses $\{q_{0t}, \dots, q_{Jt}, z_{1t}, \dots, z_{Nt}\}$ so as to maximize

$$E_t(W_{t+1}) - (\gamma/2)\text{Var}_t(W_{t+1}) \quad (2)$$

subject to $\sum_{j=0}^J q_{jt} = W_t$. We also conjecture that in equilibrium log asset returns and commodity futures prices are affine functions of a vector of factors x_t ,

$$f_{nt} = \log F_{nt} = \alpha_n + \beta'_n x_t \quad n = 1, \dots, N \quad (3)$$

$$r_{jt} = \xi_j + \psi'_j x_t \quad j = 1, \dots, J,$$

where the factors themselves can be described using a Gaussian vector autoregression:

$$x_{t+1} = c + \rho x_t + \Sigma u_{t+1} \quad u_t \sim \text{i.i.d. } N(0, I_m). \quad (4)$$

Under these assumptions, Hamilton and Wu (forthcoming) showed that

$$\begin{aligned} E_t(W_{t+1}) &\approx q_{0t}(1 + r_{0,t+1}) + \sum_{j=1}^J q_{jt} [1 + \xi_j + \psi'_j(c + \rho x_t) + (1/2)\psi'_j \Sigma \Sigma' \psi_j] \\ &\quad + \sum_{n=1}^N z_{nt} [\alpha_{n-1} + \beta'_{n-1}(c + \rho x_t) - \alpha_n - \beta'_n x_t + (1/2)\beta'_{n-1} \Sigma \Sigma' \beta_{n-1}] \end{aligned} \quad (5)$$

$$\text{Var}_t(W_{t+1}) \approx \left(\sum_{j=1}^J q_{jt} \psi'_j + \sum_{n=1}^N z_{nt} \beta'_{n-1} \right) \Sigma \Sigma' \left(\sum_{j=1}^J q_{jt} \psi_j + \sum_{\ell=1}^N z_{\ell t} \beta_{\ell-1} \right). \quad (6)$$

The arbitrageur's first-order condition associated with the choice of z_{nt} is then characterized by

$$\alpha_{n-1} + \beta'_{n-1}(c + \rho x_t) - \alpha_n - \beta'_n x_t + (1/2)\beta'_{n-1} \Sigma \Sigma' \beta_{n-1} = \beta'_{n-1} \lambda_t \quad (7)$$

where λ_t depends on the positions that arbitrageurs take in the various contracts according to

$$\lambda_t = \gamma \Sigma \Sigma' \left(\sum_{j=1}^J q_{jt} \psi_j + \sum_{\ell=1}^N z_{\ell t} \beta_{\ell-1} \right). \quad (8)$$

If we further conjecture that in equilibrium these positions are also affine functions of the underlying factors,

$$\lambda_t = \lambda + \Lambda x_t, \quad (9)$$

then equations (7) and (9) turn out to imply a recursion that the commodity-futures loadings α_n and β_n would satisfy that are very similar to those used in affine models of the term

structure of interest rates (e.g., Ang and Piazzesi (2003)).

The term λ_t allows the possibility of nonzero expected returns in equilibrium and is often referred to as the price of risk. If arbitrageurs are risk neutral, then γ and λ_t in (8) are both 0 and (7) implies that the expected net gain from any futures position (1) is always zero. More generally, with nonzero γ , equation (8) describes how changes in arbitrageurs' risk exposure coming from changes in q_{jt} or $z_{\ell t}$ would be associated with changes in expected returns. In particular, consider the effects of an exogenous increase in index-fund buying pressure on the n th contract. In equilibrium, prices must be such as to persuade arbitrageurs to take the opposite side, that is, the values of α_n and β_n , must be such that (7) holds for z_{nt} given by the necessary negative magnitude. Higher buying must be matched by more negative values for z_{nt} which change the risk premium in (7) through (8).

Hamilton and Wu (forthcoming) estimated values of the structural parameters of this system ($c, \rho, \Sigma, \lambda, \Lambda$) by inferring factors indirectly from the observed time-series properties of oil futures prices. By contrast, in this paper we propose to use direct observations on the positions of index-fund investors. Note that we can substitute (4), (3) and (9) into (7) to deduce

$$f_{n-1,t+1} - f_{nt} = \kappa_{n-1} + \pi'_{n-1}x_t + \varepsilon_{n-1,t+1} \quad (10)$$

where the theory predicts $\kappa_{n-1} = \beta'_{n-1}\lambda - (1/2)\beta'_{n-1}\Sigma\Sigma'\beta_{n-1}$, $\pi'_{n-1} = \beta'_{n-1}\Lambda$, and $\varepsilon_{n-1,t+1} = \beta'_{n-1}\Sigma u_{t+1}$. Thus a core implication of this model of risk aversion is a linear relation between the expected log returns and the notional positions z_{nt} of arbitrageurs in commodity futures contracts. Note that if arbitrageurs are risk neutral ($\gamma = 0$), then we should find $\pi_{n-1} = 0$. On the other hand, risk-averse arbitrageurs would require compensation for taking the opposite

position from index-fund buyers. If index funds are long, arbitrageurs are short, and expect subsequently to close their futures positions at a lower price than the initial contract. This would imply a negative coefficient on the element of x_t corresponding to a measure of the level of index-fund buying— if index buying makes the price of a futures contract higher than it would otherwise be, the expected excess return for a long position on that contract would be negative. The framework thus provides a motivation and interpretation for regressions similar to those reviewed in Section 2.

As a simple example, suppose that t represents months and index-fund investors always desire a long position with notional exposure K_t in the 2-month contract, selling their position as the month comes to a close in order to take a position K_{t+1} in the new 2-month contract. Then prices must be such that in equilibrium,

$$z_{nt} = \begin{cases} -K_t & \text{if } n = 2 \\ 0 & \text{otherwise} \end{cases} .$$

In the absence of other risk factors, we would then from (8) have the following simple expression determining the price of commodity-futures risk:

$$\lambda_t = -\gamma \Sigma \Sigma' \beta_1 K_t. \tag{11}$$

If we postulate an additional vector of factors x_t^* that matter for commodity price fundamentals in determining the value of f_{0t} , then the complete vector of factors is given by $x_t = (K_t, x_t^*)'$ and we have that $f_{nt} = \alpha_n + \beta_n' x_t$, with futures prices a function of both fundamentals x_t^* and

index-fund buying K_t .² If we were to rule out feedback from futures prices to fundamentals through mechanisms investigated by Knittel and Pindyck (2013), and Sockin and Xiong (2013), the first element of β_0 would be zero but the first element of β_n would generally be nonzero for all $n > 0$.³

In principle, performing the regression (10) would require observation on the demands for futures contracts coming from all the arbitrageurs' counterparties, which would include not just index-fund traders but also commercial hedgers, as well as all other factors influencing risk of any other assets held by arbitrageurs. There is nonetheless a robust implication of the reduced-form equation (10). Unless the positions of commercial hedgers and index-fund traders are perfectly negatively correlated, the positions of index-fund traders will be correlated with the net buying pressure facing arbitrageurs. If the latter is exerting a significant effect on the pricing of futures contracts, we should find a nonzero coefficient on π_{n-1} in (10) when x_t is based on observable measures of the notional positions taken by index-fund traders. We look for empirical evidence of such an effect in the next section.

²In the particular example (11) as well as many instances of the general model (7)-(9), the factors x_t would be spanned by the set of commodity futures and asset prices. However, our empirical specification (10) is also appropriate for the more general case of unspanned factors. The regression (12) looks for a relation between an excess return between dates t and $t+1$ and an observable variable at date t that is not constructed directly from date t prices.

³In the above example in which index funds want only the 2-month contract at time t , an arbitrageur who shorts this contract is exposed to the complete vector of factor risks associated with this contract through $\beta_1' x_{t+1}$, including for example fundamentals uncertainty about the price of the underlying commodity. Assuming these fundamentals are positively serially correlated, if there was zero expected return associated with a 1-period contract purchased at t , an arbitrageur would want to be long this contract to hedge some of the risk associated with being short the 2-period contract. Since there is no one to take the other side of this 1-period contract, equilibrium then requires a negative expected return on a 1-period contract purchased at t just as on a 2-period contract purchased at t . For details see Hamilton and Wu (forthcoming).

4 Predicting futures returns for agricultural commodities.

For 12 agricultural commodities, since 2006 the CFTC has been providing through its Supplemental Commitments of Traders Report weekly positions for traders it characterizes as “replicating a commodity index by establishing long futures positions in the component markets and then rolling those positions forward from future to future using a fixed methodology” (CFTC (2012)). While some of these index traders are pension funds or other managed funds taking a direct position in futures contracts, the majority represent positions by swap dealers, who offer their clients an over-the-counter product that mimics some futures-based index (Stoll and Whaley (2010)). The swap dealers are thus implicitly short in futures agreements arranged over-the-counter, and hedge with an offsetting long position on organized exchanges that get reported to the CFTC. While the CFTC designations are not without some problems (Cheng, Kirilenko, and Xiong (2012), Stoll and Whaley (2010)), these appear to be the best high-frequency data that are publicly available for our purposes.

The SCOT is released on a Friday and reflects positions as of the preceding Tuesday. Let t denote a week ending on a particular Tuesday and X_t the long positions (measured in number of contracts) being held by commodity index traders as identified by SCOT. Let F_t denote the price of the near contract as of the market close on the day for which SCOT reports X_t .⁴

⁴Daily futures prices were purchased from Norma’s Historical Data (<http://www.normahistoricaldata.com/>). For a few days in our sample, SCOT data are reported for days on which we have no closing futures prices. For these observations, we used the futures price as of the following day.

If this is the contract being held by a typical index trader, then

$$\tilde{x}_t = 100(\ln X_t + \ln F_t)$$

would correspond to the log of index traders' notional exposure.⁵

Our interest centers on the predictability of r_t , the weekly return of a given contract.⁶ The first two rows within each block of Table 1 report coefficients and standard errors for the following regression estimated for t running from April 11, 2006 to January 3, 2012:

$$r_t = \alpha_1 + \phi_1 r_{t-1} + \pi_1 \tilde{x}_{t-1} + \varepsilon_{1t}. \quad (12)$$

The coefficient estimates $\hat{\phi}_1$ and $\hat{\pi}_1$ are not statistically significantly different from zero for any of the 12 commodities for which index trader positions are reported, and adjusted R^2 for these regressions are usually negative. This result is consistent with the large number of previous studies discussed in Section 2 that have found limited predictability of commodity futures returns using related regressions.

Equations (4) and (10) can be viewed as the central implications of the model of risk pricing sketched in Section 3. Given observed data on futures prices and x_t , it is straightforward to estimate the unrestricted reduced form implications of these models using ordinary least

⁵Note that X_t is measured in number of contracts, so that the true measure of notional exposure would further multiply X_t by the number of barrels of oil in a single contract. This would simply add a constant to \tilde{x}_t and would have no effect on the slope coefficients in any of the regressions reported in this section, so we have always used the simpler expression given in the text.

⁶For most weeks, r_t is just $100(f_t - f_{t-1})$ where $f_t = \ln(F_t)$ the log of the price of the nearest contract. In other words, r_t is the percent change in price of the near contract. In the case when the near contract as of date $t - 1$ had expired as of date t , we took r_t to be 100 times the change between f_t and the log of the price of that same contract as of $t - 1$ (at which date it was the second available contract).

squares. These coefficients could then be used to infer values for the parameters $(c, \rho, \Sigma, \lambda, \Lambda)$ of the structural model using methods described by Hamilton and Wu (2012) and Hamilton and Wu (forthcoming), and in fact this simple, intuitive approach to estimation turns out to be asymptotically equivalent to full information maximum likelihood estimation of the model. However, when the key reduced-form parameters ϕ and π are statistically indistinguishable from zero, one would have little confidence in any values for the risk-pricing parameters in Λ that one might infer from the data, and for this reason we have chosen not to try to go beyond the simple reduced-form estimates reported here. Our conclusion is that although in principle index-fund buying of commodity futures could influence pricing of risk, we do not find confirmation of that in the week-to-week variability of the notional value of reported commodity index trader positions.

If the factors governing the price of risk are stationary, the dependent variable in (12) would be stationary, in which case one might prefer to use the weekly change in index-trader notional positions rather than the level as the explanatory variable. The third and fourth rows of each block in Table 1 report OLS coefficient estimates and standard errors for the regression

$$r_t = \alpha_2 + \phi_2 r_{t-1} + \pi_2 (\tilde{x}_{t-1} - \tilde{x}_{t-2}) + \varepsilon_{2t}. \quad (13)$$

The coefficients on the change in index notional positions turn out to be statistically significant at the 5% level for cotton and almost statistically significant for sugar. However, the coefficients are of opposite signs, and neither $\hat{\phi}$ nor $\hat{\pi}$ are statistically significantly different from zero for any of the other 10 commodities. We again conclude that there appears to be very little indication in the data that changes in positions of index traders can help explain

risk premia in commodity futures prices.

Singleton (forthcoming) has recently suggested that investment flows may matter over longer periods than a week, and instead bases his analysis of crude oil futures returns on the change in notional positions over a 3-month period. The fifth and sixth rows of each block in Table 1 report coefficient estimates and standard errors for our 12 agricultural commodities in which the predictive variable is the lagged 13-week change in the commodity index notional position:

$$r_t = \alpha_3 + \phi_3 r_{t-1} + \pi_3 (\tilde{x}_{t-1} - \tilde{x}_{t-14}) + \varepsilon_{3t}. \quad (14)$$

None of the 24 estimated slope coefficients is statistically distinguishable from zero. To summarize, we find no persuasive evidence that either the level, weekly change, or 13-week change in index-trader positions is related to the risk premium in agricultural commodities.

Finally, we report some simple evidence that is robust with respect to any problems in measuring the volume of index buying itself and makes use of higher-frequency features of the data. The two main indexes that buyers seek to track are the S&P-Goldman Sachs Commodity Index and the Dow Jones-UBS Commodity Index. Each of these has a defined calendar schedule at which a near contract is sold and the next contract is purchased. The Goldman strategy begins the roll into the next contract on the 5th business day of the month and is completed on the 9th, whereas the Dow Jones strategy begins the roll on the 6th and ends on the 10th. It is simple enough to ask whether there is anything special about price movements on these particular days.

We calculated the daily return on the near contract $r_\tau = 100(f_\tau - f_{\tau-1})$ where τ now

indexes business days, and estimated the OLS regression

$$r_\tau = \alpha_4 + \pi_4 S_\tau + \varepsilon_{4\tau}. \quad (15)$$

Here $S_\tau = 1$ if both (1) τ is the 5th through the 10th business day of the month and (2) the Goldman and Dow Jones strategies would be selling the near contract during that period; otherwise, $S_\tau = 0$. In the case of crude oil, for which a new contract exists every month, S_τ is a simple calendar dummy tracking the 5th through 10th business days of the month. For all other commodities, some months are not traded; for example, there is no February contract in soybeans. Thus as of the beginning of January, the near soybean contract would be the March contract. The index funds would not be selling their soybean position during January, so S_τ would remain zero for soybeans throughout the month of January. Note that the OLS estimate $\hat{\pi}_4$ is thus numerically equal to the difference between the average daily return from a long position on the near contract that is sold on days when the index funds are selling that contract compared to the return on the near contract on non-roll days. Moreover, because the average return over 6 days is numerically identical to $(1/6)$ times the 6-day return,

$$N^{-1} \sum_{n=1}^N r_{\tau+n} = (100/N)(f_{\tau+N} - f_\tau),$$

one can equivalently interpret regression (15) as indicating the average result if one buys the near contract on the day before the index funds begin to sell and sells on the day when the funds are finished selling their positions.

Columns 1 and 2 of Table 2 report coefficient estimates and standard errors for (15) fit to

daily returns on the near contract for the commodities we analyze. We find a statistically significant value for $\hat{\pi}_4$ for only 1 of the 12 agricultural commodities (soybeans), and this is of the opposite sign predicted by the simple price impact hypothesis— if index-fund selling is depressing the price on roll days, one would expect the price of the near contract to fall rather than rise on those days.

Columns 3 and 4 of Table 2 report results for the analogous regressions using the next contract. Here only 2 of the estimated coefficients (namely, those for soybeans and hogs) are statistically significantly different from zero. And while these are of the predicted sign (a positive coefficient consistent with the claim that index buying of the next contract is pushing the price of the contract up), note that for beans a statistically significant positive coefficient was found for both the near contract that funds are selling (column 2) as well as the next contract that funds are buying (column 4).

Finally, in columns 5 and 6 we report results when the dependent variable is the return on the next contract minus the return on the near. Here π_4 can be interpreted as the average excess returns for someone who is simultaneously buying the next contract and selling the near. For soybean oil and cotton we find a statistically significant negative value for π_4 and for hogs we find a statistically significant positive value.

Aulerich, Irwin, and Garcia (2013) suggested that the roll period itself might be too narrow to capture all the effects. In Table 3 we repeat the analysis in which $S_\tau = 1$ if both (1) τ is the 1st through the 10th business day of the month and (2) the Goldman and Dow Jones strategies would be selling the near contract during that period. For this specification, none of the coefficients on S_τ are statistically significant for any of the individual agricultural contracts

(columns 2 and 4 of Table 3), and there is only one statistically significant difference between the coefficients on the first two contracts (cotton in column 6).

Other studies have also looked for price patterns associated with the roll window, and have reported mixed results. Mou (2010) used a window beginning 5 to 10 days before the roll and pooled across commodities to find statistically significant results. Bessembinder, Carrion, Tuttle, and Venkataraman (2012) found price effects that were usually reversed within 15 minutes, while Stoll and Whaley (2010) found little impact of the roll. Brunetti and Reiffen (2011) reported that higher index-fund positions raised the return on the second versus the near contract during the roll periods, while Aulerich, Irwin, and Garcia (2013) found the opposite effect.

We conclude that for our sample period and data set, any effects associated with the roll window are too small to show up in simple summary regressions. Our overall conclusion seems to be a robust summary of this data set—changes in the futures positions of commodity-index funds are not associated with significant changes in the expected returns for their counterparties.

5 Predicting returns for crude oil contracts.

Although there appears to be limited relation between commodity index positions and agricultural futures returns, Singleton (forthcoming) reported that a number of interesting variables appear to help predict returns on crude oil futures contracts. Our interest in this paper is the strong predictive power he found for a measure of holdings by index-fund traders in crude

oil futures contracts.

The CFTC does not report a weekly estimate of index-fund positions in crude oil. Singleton’s approach was to use an idea suggested by Masters (2008). Masters claimed that the vast majority of index-fund buyers were following one of two strategies, trying to track either the S&P-Goldman Sachs Commodity Index or the Dow Jones-UBS (formerly Dow Jones-AIG) Commodity Index. Each of these would take positions in a particular set of commodities and roll them over according to a prespecified calendar schedule. Let \tilde{X}_{it} be the notional value of contracts in agricultural commodity i reported by SCOT to have been held by index-fund traders on date t , that is, $\tilde{X}_{it} = X_{it}F_{it}Q_i$ for X_{it} the number of contracts held by index-fund traders, F_{it} the futures price, and Q_i the number of units of the commodity held by a single contract.⁷ Suppose we assume that this is composed entirely of either traders following the Goldman strategy, whose notional holdings we denote by \tilde{X}_{it}^G , or traders following the Dow Jones strategy (\tilde{X}_{it}^D):

$$\tilde{X}_{it} = \tilde{X}_{it}^G + \tilde{X}_{it}^D.$$

Masters (2008) noted that while the Dow Jones index holds soybean oil, the Goldman index does not, in which case

$$\tilde{X}_{\text{bean.oil},t} = \tilde{X}_{\text{bean.oil},t}^D.$$

We further know that the total notional exposure \tilde{X}_t^D of funds replicating the Dow Jones index would be allocated across the commodities according to known weights δ_{it}^D ,⁸ so that

⁷Note that this is a slight change in notation from the previous sections, where we instead measured \tilde{X}_{it} as just $X_{it}F_{it}$. As noted in footnote 5, the slope coefficients in the previous section would be numerically identical regardless of which convention is used, but the added term Q_i needs to be included for the Masters-type calculations used in this section.

⁸We thank Dow Jones for providing us with historical values for the weights δ_{it}^D .

$\tilde{X}_{it}^D = \delta_{it}^D \tilde{X}_t^D$. Hence one could use the SCOT bean oil figures to infer the total notional holdings of Dow Jones investors, $\tilde{X}_t^D = \tilde{X}_{\text{bean.oil},t} / \delta_{\text{bean.oil},t}^D$, and from this impute holdings of Dow Jones investors in crude oil contracts on the basis of SCOT reported holdings of soybean oil:

$$\tilde{X}_{\text{crude.oil},t}^{D,[\text{bean.oil}]} = \delta_{\text{crude.oil},t}^D \tilde{X}_{\text{bean.oil},t} / \delta_{\text{bean.oil},t}^D.$$

We have added the superscript [bean_oil] to this estimate of crude oil notional holdings to clarify that the underlying SCOT data from which it is derived in fact describe soybean oil holdings rather than crude oil. Similarly, the Goldman index holds Kansas City wheat whereas Dow Jones does not, giving an estimate of Goldman holdings of crude oil contracts:⁹

$$\tilde{X}_{\text{crude.oil},t}^{G,[\text{KC.wheat}]} = \delta_{\text{crude.oil},t}^G \tilde{X}_{\text{KC.wheat},t} / \delta_{\text{KC.wheat},t}^G. \quad (16)$$

The Goldman index also holds feeder cattle which Dow Jones does not, affording an alternative estimate

$$\tilde{X}_{\text{crude.oil},t}^{G,[\text{feeder.cattle}]} = \delta_{\text{crude.oil},t}^G \tilde{X}_{\text{feeder.cattle},t} / \delta_{\text{feeder.cattle},t}^G. \quad (17)$$

Masters proposed to use the average of (16) and (17) as an estimate of Goldman crude oil contract holdings. The sum of the oil holdings imputed to these two funds is then his estimate of total index-fund holdings of crude oil contracts:

$$\tilde{X}_{\text{crude.oil},t}^{[\text{Masters}]} = \tilde{X}_{\text{crude.oil},t}^{D,[\text{bean.oil}]} + (1/2) \left(\tilde{X}_{\text{crude.oil},t}^{G,[\text{KC.wheat}]} + \tilde{X}_{\text{crude.oil},t}^{G,[\text{feeder.cattle}]} \right). \quad (18)$$

⁹We thank Standard & Poor's for providing us with historical values for δ_{it}^G .

Figure 1 plots the number of contracts associated with this value for $\tilde{X}_{\text{crude.oil},t}^{[\text{Masters}]}$ against the price of crude oil based on the near contract, updating similar figures in Masters (2008) and Singleton (forthcoming). The figure suggests a strong connection between these two series, particularly during 2008 and 2009.

We repeated our basic regressions (12)-(14) for r_t now the weekly return on the near crude oil futures contract and $\tilde{x}_t = 100 \ln \tilde{X}_{\text{crude.oil},t}^{[\text{Masters}]}$. These results are reported in the first block of Table 4. Both the levels and weekly-difference regression results are similar to those for agricultural commodities, with negative values for \bar{R}^2 and statistically insignificant coefficients.

We also find using daily data (last rows of Table 2) no evidence of excess returns from buying the near or next crude oil contract during the period in which many index traders are rolling contracts. We do find statistically significant coefficients when a 10-day window is used (last rows of Table 3), though the fact that there is the same negative coefficient on the first and second contracts is again inconsistent with the claim that this correlation arises as a consequence of index funds selling the near contract and buying the next.

In contrast to the broad lack of evidence of an effect of commodity-index investing based on the tests reported so far, the last row of the first block of Table 4 reproduces Singleton's finding that the 13-week change in $\tilde{X}_{\text{crude.oil},t}^{[\text{Masters}]}$ appears to be quite helpful for predicting changes in crude oil futures prices. Since the use of a 13-week window appears to increase the variable's predictive power for oil returns but not for any of the other 12 commodities, we were curious

to look at the performance of the oil regression across all possible window choices:

$$r_t = \alpha_n + \phi_n r_{t-1} + \pi_n (\tilde{x}_{t-1} - \tilde{x}_{t-1-n}) + \varepsilon_{nt}.$$

Figure 2 plots the adjusted R^2 for each possible choice of n between 1 and 26. The biggest \bar{R}^2 turns out to be obtained by setting $n = 12$, close to the value $n = 13$ proposed by Singleton.

Irwin and Sanders (2012) have recently raised several strong criticisms of using $\tilde{X}_{\text{crude.oil},t}^{[\text{Masters}]}$ as a measure of index-fund positions in crude oil futures. Irwin and Sanders noted first that although for some dates the two measures $\tilde{X}_{\text{crude.oil},t}^{G, [\text{KC-wheat}]}$ and $\tilde{X}_{\text{crude.oil},t}^{G, [\text{feeder.cattle}]}$ are reasonably close, for other dates they can differ greatly. In terms of this concern, we would point out that in fact there is no need to restrict the inference as Masters did only to commodities that appear in one of the indexes but not the other. Note that the central claim is that

$$\tilde{X}_{it} = \delta_{it}^G \tilde{X}_t^G + \delta_{it}^D \tilde{X}_t^D \tag{19}$$

where \tilde{X}_{it} , δ_{it}^G , and δ_{it}^D are all observed directly. Given any two arbitrary agricultural commodities i and j , one can use the two equations in (19) to solve for the implied total holdings of the two indexes, \tilde{X}_t^G and \tilde{X}_t^D . Hence for any two arbitrary agricultural commodities i and j there exists a Masters-type estimate of crude oil holdings:

$$\tilde{X}_{\text{crude.oil},t}^{[i,j]} = \begin{bmatrix} \delta_{\text{crude.oil},t}^G & \delta_{\text{crude.oil},t}^D \end{bmatrix} \begin{bmatrix} \delta_{it}^G & \delta_{it}^D \\ \delta_{jt}^G & \delta_{jt}^D \end{bmatrix}^{-1} \begin{bmatrix} \tilde{X}_{it} \\ \tilde{X}_{jt} \end{bmatrix}. \tag{20}$$

To illustrate the variability of such measures, we calculated (20) where $j = \text{soybean oil}$

(one of the series used by Masters and Singleton) and i corresponding to any one of the other 11 possible choices. Plots of $\tilde{X}_{\text{crude.oil},t}^{[i,\text{bean.oil}]}$ for different i are plotted in Figure 3. These series indeed appear to be fairly sensitive to the choice of i .

One way to deal with this issue is to generalize Masters' averaging idea. There's no reason in fact not to use all 12 agricultural commodities together, choosing \tilde{X}_t^G and \tilde{X}_t^D so as to minimize the sum of squared discrepancies in predicting the SCOT reported value for \tilde{X}_{it} across the 12 commodities. This amounts to treating the collection of equations (19) for $i = 1, \dots, 12$ and for a given t as a sample of size 12 in which the dependent variable is \tilde{X}_{it} and explanatory variables are δ_{it}^G and δ_{it}^D :

$$\tilde{X}_{\text{crude.oil},t}^{[\text{all}]} = \begin{bmatrix} \delta_{\text{crude.oil},t}^G & \delta_{\text{crude.oil},t}^D \end{bmatrix} \begin{bmatrix} \sum_{i=1}^{12} (\delta_{it}^G)^2 & \sum_{i=1}^{12} \delta_{it}^G \delta_{it}^D \\ \sum_{i=1}^{12} \delta_{it}^D \delta_{it}^G & \sum_{i=1}^{12} (\delta_{it}^D)^2 \end{bmatrix}^{-1} \begin{bmatrix} \sum_{i=1}^{12} \delta_{it}^G \tilde{X}_{it} \\ \sum_{i=1}^{12} \delta_{it}^D \tilde{X}_{it} \end{bmatrix}. \quad (21)$$

These regression-based estimates of index-fund holdings are also plotted in Figure 3.

We repeated our predictive regressions using $\tilde{X}_{\text{crude.oil},t}^{[\text{all}]}$ in place of $\tilde{X}_{\text{crude.oil},t}^{[\text{Masters}]}$ in the second block of Table 4. The results turn out to be quite similar to those based on Masters' original series. Nothing is significant and \bar{R}^2 are negative for the levels or first-difference specifications, but there is significant predictability from the 13-week change in $\tilde{X}_{\text{crude.oil},t}^{[\text{all}]}$. Interestingly, we found the same is true for every one of the bivariate estimates plotted in Figure 3: the 13-week change in crude oil index-trader positions, inferred from soybean oil and any other arbitrary agricultural commodity, appear to help predict crude oil returns.

A second concern that Irwin and Sanders (2012) raised about Masters' methodology is that the index oil holdings as imputed from agricultural commodity positions are quite different

from direct estimates of crude oil positions that the CFTC reports in its new Index Investment Data report. Unfortunately, the Index Investment Data report is only available at a quarterly frequency and for a shorter period than SCOT, so is not usable for the kind of regressions we're interested in here. However, it is puzzling that agricultural positions cannot predict agricultural prices but do predict crude oil prices, and that a direct measure of crude oil index holding would do a poorer job at predicting oil prices than does an index imputed from agricultural holdings.

One way to shed further light on these issues is to investigate whether the in-sample success of the oil regression (14) translates into useful out-of-sample forecasts. Since Singleton's paper was first circulated we have obtained an additional 2 years of data, which allow us to see whether the 13-week change has predictive power outside of the sample for which it was originally proposed. We estimated regression (14) for a sample ending January 12, 2010, which was the end date for Singleton's analysis. If we use those coefficients to predict oil returns over January 17, 2010 through January 3, 2012, the out-of-sample mean squared error is 24.01. That compares with an out-of-sample MSE of 21.97 if we had instead simply always forecast $r_t = 0$ (see the first row of Table 5).

Note that the above calculations represent a true out-of-sample evaluation, namely, a calculation of how well a proposed empirical relation describes data that came in after the initial study has been released. It is in this sense a more meaningful exercise than the pseudo out-of-sample evaluations that are popularly reported. Hansen and Timmermann (2013) demonstrate that the popular practice of calculating an "out-of-sample" MSE of a set of recursive regressions that end at all points between some sample dates T_0 and T is

asymptotically equivalent to looking at the difference between two simple Wald tests, the first statistic using just the subsample of observations from 1 to T_0 and the second statistic using all data from 1 to T . If our goal is to evaluate whether a variable belongs in a forecasting relation or to test whether the regression relation is stable, there are a number of alternative tests that are much more appropriate than these artificial “out-of-sample” forecast evaluation exercises. For example, making efficient use of the full sample of data from April 11, 2006 to January 3, 2012, the Bai and Perron (1998) test leads to the conclusion that there are two structural breaks¹⁰ in the oil return regression (14) dated at September 30, 2008 and January 13, 2009— both inside the original Singleton sample of data. The coefficient on $\tilde{x}_{t-1} - \tilde{x}_{t-14}$ is positive over the 2006-2008 subsample, as it was found to be in Singleton’s regressions and in our regressions using the full sample of observations from 2006 to 2012 as well as the 2006-2010 regression in row 1 of Table 5. However, when estimated with the two indicated break points, the coefficient turns out to be negative over both the second and third subsamples, with a t statistic of -0.75 for data from January 30, 2009 through January 3, 2012. The correlation identified by Singleton thus has no success at describing data since his paper was written, and indeed seems not to have captured a stable predictive relation even within the sample that he analyzed.

Returning to Figure 1, the striking feature of the Masters indicator is that it collapses as the recession worsened in 2008 but began to rebound sharply before the recovery began, key movements that precede equally dramatic parallel moves in oil prices. The close fit over

¹⁰This is based on evaluation of the sequential F test in Bai and Perron (1998) Table II imposing the restriction that breaks must be separated by at least 5% of the sample size (in this case, 15 observations) and that the maximum number of breaks is $m = 4$. The identical conclusion of 2 structural breaks emerges from application of the Schwarz (1978) criterion.

Singleton’s original sample period thus seems to result from the broad comovement of the series during the first phase of the Great Recession. It’s interesting to note that if we replace r_t in expression (14) with the weekly return on the U.S. S&P500 stock price index, but with the second explanatory variable still $100[\ln \tilde{X}_{\text{crude.oil},t-1}^{\text{Masters}} - \tilde{X}_{\text{crude.oil},t-14}^{\text{Masters}}]$, for a sample that ends at Singleton’s January 2010 endpoint, the Masters variable would also appear to be positive and statistically significant (see the last rows of Table 5). In other words, if we used only data from the recession, we would conclude that index-trader positions in soybean oil, Kansas City wheat, and feeder cattle could also be used to predict stock prices. Once again, however, a relation estimated over this period has a bigger out-of-sample MSE than the simple no-change forecast.¹¹

6 Conclusion.

The increased participation by financial investors in commodity futures markets over the last decade has been quite substantial. In principle this could have influenced the risk premium, and Hamilton and Wu (forthcoming) found significant changes in the behavior of the risk premium on oil futures contracts before and after 2005. In this paper we studied data since 2006 to look for a systematic relation between the notional value of commodity futures contracts held on behalf of index-fund investors and expected returns on futures contracts.

We found essentially no relation for the 12 agricultural commodities for which the CFTC

¹¹A referee notes that if $\tilde{x}_t - \tilde{x}_{t-13}$ were indeed capturing the market pricing kernel, then according to the theory sketched in Section 3 it should also help predict returns on all assets. To us a more natural reading of the above results is that the magnitude never had any true ability to predict returns on oil contracts or any other asset, but instead is just a constructed variable whose movements for a brief period turned out to be in the same direction as other important market developments.

reports such positions. We reviewed evidence that positions in crude oil contracts imputed from the reported agricultural holdings could help predict crude oil futures returns, and noted that the methodology for such imputation could be generalized to make use of all the available data. We confirmed that these imputed holdings appear to help predict crude oil returns over 2006-2009, though this is closely related to the dynamics of index investing during the Great Recession, and indeed the same imputed holdings also appear to predict stock returns over that period. We found, however, that both relations broke down when trying to describe the data since 2009.

Our overall conclusion is thus consistent with most of the previous literature—there seems to be little evidence that index-fund investing is exerting a measurable effect on commodity futures prices. As noted in the introduction, even if one could demonstrate an effect of index-fund buying on commodity futures prices, it would be a separate challenge to explain how this could also end up changing the equilibrium spot price. We conclude that it is difficult to find much empirical foundation for a view that continues to have a significant impact on policy decisions.

References

- Acharya, Viral V., Lars A. Lochstoer, and Tarun Ramadorai (forthcoming) “Limits to Arbitrage and Hedging: Evidence from Commodity Markets” *Journal of Financial Economics*.
- Alquist, Ron, and Olivier Gervais (2011) “The Role of Financial Speculation in Driving the Price of Crude Oil” Working paper, Bank of Canada.

- Ang, Andrew, and Monika Piazzesi (2003) “A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables” *Journal of Monetary Economics* 50, 745–787.
- Aulerich, Nicole M., Scott H. Irwin, and Philip Garcia (2013) “Bubbles, Food Prices, and Speculation: Evidence from the CFTC’s Daily Large Trader Data Files” NBER Working paper 19065.
- Bai, Jushan, and Pierre Perron (1998) “Estimating and Testing Linear Models with Multiple Structural Changes” *Econometrica* 66, 47–78.
- Bessembinder, Hendrik (1992) “Systemic Risk, Hedging Pressure, and Risk Premiums in Futures Markets” *Review of Financial Studies* 5, 637–667.
- Bessembinder, Hendrik, Allen Carrion, Laura Tuttle, and Kumar Venkataraman (2012) “Predatory or Sunshine Trading? Evidence from Crude Oil ETF Rolls” Working paper, Southern Methodist University.
- Brunetti, Celso, Bahattin Buyuksahin, and Jeffrey H. Harris (2011) “Speculators, Prices and Market Volatility” Working paper, Johns Hopkins University.
- Brunetti, Celso, and David Reiffen (2011) “Commodity Index Trading and Hedging Costs” Working paper, Federal Reserve Board.
- Buyuksahin, Bahattin, and Jeffrey H. Harris (2011) “Do Speculators Drive Crude Oil Futures Prices?” *Energy Journal* 32, 167–202.

- Buyuksahin, Bahattin, and Michel A. Robe (2010) “Speculators, Commodities and Cross-Market Linkages” Working paper, American University.
- Buyuksahin, Bahattin, and Michel A. Robe (2011) “Does ‘Paper Oil’ Matter?” Working paper, American University.
- Carter, Colin A., Gordon C. Rausser, and Andrew Schmitz (1983) “Efficient Asset Portfolios and the Theory of Normal Backwardation” *Journal of Political Economy* 91, 319–331.
- CFTC (2012) “Commitments of Traders: Explanatory Notes” [Http://www.cftc.gov/MarketReports/CommitmentsofTraders/ExplanatoryNotes/index.htm](http://www.cftc.gov/MarketReports/CommitmentsofTraders/ExplanatoryNotes/index.htm).
- Chang, Eric C. (1985) “Returns to Speculators and the Theory of Normal Backwardation” *Journal of Finance* 40, 193–208.
- Cheng, Ing-Haw, Andrei Kirilenko, and Wei Xiong (2012) “Convective Risk Flows in Commodity Futures Markets” Working paper, Princeton University.
- Cootner, Paul H. (1960) “Returns to Speculators: Telser versus Keynes” *Journal of Political Economy* 68, 396–404.
- Danielsson, Jon, Hyun Song Shin, and JeanPierre Zigrand (2011) “Balance Sheet Capacity and Endogenous Risk” Working paper, Princeton University.
- De Roon, Frans, Theo Nijman, and Chris Veld (2000) “Hedging Pressure Effects in Futures Markets” *Journal of Finance* 55, 1437–1456.
- Duffie, Darrell (1992) *Dynamic Asset Pricing Theory* Princeton University Press, Princeton, NJ.

- Etula, Erkki (forthcoming) “Broker-Dealer Risk Appetite and Commodity Returns” *Journal of Financial Econometrics*.
- Fattouh, Bassam, Lutz Kilian, and Lavan Mahadeva (forthcoming) “The Role of Speculation in Oil Markets: What Have We Learned So Far?” *Energy Journal*.
- Hamilton, James D. (2009) “Causes and Consequences of the Oil Shock of 2007-08” *Brookings Papers on Economic Activity* Spring 2009, 215–259.
- Hamilton, James D., and Jing Cynthia Wu (2012) “Identification and Estimation of Gaussian Affine Term Structure Models” *Journal of Econometrics* 168, 315–331.
- Hamilton, James D., and Jing Cynthia Wu (forthcoming) “Risk Premia in Crude Oil Futures Prices” *Journal of International Money and Finance*.
- Hansen, Peter R., and Allan Timmermann (2013) “Equivalence Between Out-of-Sample Forecast Comparisons and Wald Statistics” Working paper, UCSD.
- Hartzmark, Michael L. (1987) “Returns to Individual Traders of Futures: Aggregate Results” *Journal of Political Economy* 95, 1292–1306.
- Irwin, Scott H., and Dwight R. Sanders (2011) “Index Funds, Financialization, and Commodity Futures Markets” *Applied Economic Perspectives and Policy* 33, 1–31.
- Irwin, Scott H., and Dwight R. Sanders (2012) “Testing the Masters Hypothesis in Commodity Futures Markets” *Energy Economics* 34, 256–269.
- Juvenal, Luciana, and Ivan Petrella (2011) “Speculation in the Oil Markets” Working paper, Federal Reserve Bank of St. Louis.

Kennedy, Joseph P. (2012) “The High Cost of Gambling on Oil” *New York Times*, April 10.

Keynes, John M. (1930) *A Treatise on Money, Volume 2* Macmillan, London.

Kilian, Lutz, and Thomas K. Lee (forthcoming) “Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories” *Journal of International Money and Finance*.

Kilian, Lutz, and Daniel P. Murphy (2012) “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models” *Journal of the European Economic Association*.

Kilian, Lutz, and Daniel P. Murphy (forthcoming) “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil” *Journal of Applied Econometrics*.

Knittel, Christopher R., and Robert S. Pindyck (2013) “The Simple Economics of Commodity Price Speculation” Working paper, MIT.

Kolb, Robert W. (1992) “Is Normal Backwardation Normal?” *Journal of Futures Markets* 12, 75–90.

Lombardi, Marco J., and Ine van Robays (2011) “Do Financial Investors Destabilize the Oil Price?” Working paper, European Central Bank.

Marcus, Alan J. (1984) “Efficient Asset Portfolios and the Theory of Normal Backwardation: A Comment” *Journal of Political Economy* 92, 162–164.

Masters, Michael W. (2008) “Testimony before Committee on Homeland Security and Governmental Affairs of the United States Senate” May 20.

- Masters, Michael W. (2009) “Testimony before Commodity Futures Trading Commission”
August 25.
- Mou, Yiqun (2010) “Limits to Arbitrage and Commodity Index Investment: Front-Running
the Goldman Roll” Working paper, Columbia University.
- Sanders, Dwight R., and Scott H. Irwin (2011a) “The Impact of Index Funds in Commodity
Futures Markets: A Systems Approach” *Journal of Alternative Investments* 14, 40–49.
- Sanders, Dwight R., and Scott H. Irwin (2011b) “New Evidence on the Impact of Index Funds
in U.S. Grain Futures Markets” *Canadian Journal of Agricultural Economics* 59, 519–532.
- Schwarz, Gideon (1978) “Estimating the Dimension of a Model” *Annals of Statistics* 6, 461–
464.
- Singleton, Kenneth J. (forthcoming) “Investor Flows and the 2008 Boom/Bust in Oil Prices”
Management Science.
- Sockin, Michael, and Wei Xiong (2013) “Feedback Effects of Commodity Futures Prices”
NBER Working Paper 18906.
- Stoll, Hans R., and Robert E. Whaley (2010) “Commodity Index Investing and Commodity
Futures Prices” *Journal of Applied Finance* 20, 7–46.
- Tang, Ke, and Wei Xiong (2012) “Index Investment and the Financialization of Commodities”
Financial Analysts Journal 68, 54–74.

Table 1: Tests for predictability of agricultural commodity returns

Beans		Wheat			Corn			
	const	r_{t-1}	X_{t-1}	\bar{R}^2	const	r_{t-1}	X_{t-1}	\bar{R}^2
level	10.8033 (9.6036)	-0.0032 (0.0579)	-0.0056 (0.0051)	-0.0026	30.8385 (17.5923)	0.0014 (0.0579)	-0.0166 (0.0094)	0.0036
1-wk diff	0.2640 (0.2306)	-0.0443 (0.1336)	0.0321 (0.1017)	-0.0064	-0.0489 (0.3091)	0.1817 (0.1814)	-0.1738 (0.1591)	-0.0027
13-wk diff	0.2303 (0.2345)	-0.0191 (0.0599)	0.0082 (0.0099)	-0.0044	-0.0570 (0.3099)	-0.0008 (0.0596)	-0.0058 (0.0148)	-0.0062
BeanOil		Cattle			Cocoa			
	const	r_{t-1}	X_{t-1}	\bar{R}^2	const	r_{t-1}	X_{t-1}	\bar{R}^2
level	8.9077 (7.7530)	0.0098 (0.0579)	-0.0058 (0.0052)	-0.0024	4.0293 (6.8600)	-0.0628 (0.0577)	-0.0024 (0.0042)	-0.0013
1-wk diff	0.1748 (0.2294)	-0.1026 (0.0948)	0.1034 (0.0704)	0.0006	0.0746 (0.1206)	-0.1521 (0.0971)	0.0717 (0.0643)	0.0017
13-wk diff	0.1801 (0.2325)	0.0046 (0.0594)	0.0024 (0.0098)	-0.0065	0.0740 (0.1210)	-0.0738 (0.0584)	0.0069 (0.0077)	0.0002
Coffee		Cotton			FedCattle			
	const	r_{t-1}	X_{t-1}	\bar{R}^2	const	r_{t-1}	X_{t-1}	\bar{R}^2
level	3.8530 (9.2521)	-0.0329 (0.0582)	-0.0024 (0.0059)	-0.0050	2.3353 (11.7720)	0.0644 (0.0589)	-0.0014 (0.0076)	-0.0027
1-wk diff	0.1362 (0.2370)	-0.0011 (0.1252)	-0.0293 (0.0960)	-0.0052	0.1601 (0.2704)	-0.2249 (0.1293)	0.2530* (0.1018)	0.0176
13-wk diff	0.1454 (0.2379)	-0.0290 (0.0594)	-0.0054 (0.0116)	-0.0048	0.1482 (0.2731)	0.0563 (0.0603)	0.0046 (0.0113)	-0.0022
Hogs		KCWheat			Sugar			
	const	r_{t-1}	X_{t-1}	\bar{R}^2	const	r_{t-1}	X_{t-1}	\bar{R}^2
level	-10.9569 (9.5632)	0.0530 (0.0587)	0.0069 (0.0061)	0.0018	9.0292 (12.0646)	0.0269 (0.0582)	-0.0053 (0.0071)	-0.0043
1-wk diff	-0.1522 (0.2002)	0.1016 (0.0950)	-0.0352 (0.0726)	-0.0017	0.0755 (0.2862)	-0.0062 (0.1097)	0.0263 (0.0840)	-0.0059
13-wk diff	-0.1394 (0.2004)	0.0510 (0.0594)	0.0104 (0.0102)	0.0010	0.0662 (0.2867)	0.0152 (0.0595)	0.0065 (0.0113)	-0.0051

Notes to Table 1. OLS coefficients (with standard errors in parentheses) for regressions (12), (13), and (14), for $X_{t-1} = \tilde{x}_{t-1}$ for the level rows, $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-2})$ for the 1-wk diff rows, and $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-14})$ for the 13-wk diff rows. All regressions estimated April 11, 2006 to January 3, 2012. * indicates significant at 5% level.

Table 2: Average returns during index-fund roll versus normal times

	1 st contract		2 nd contract		spread	
	constant	S_τ	constant	S_τ	constant	S_τ
Beans	0.0195 (0.0523)	0.3710* (0.1514)	-0.0015 (0.0481)	0.3720* (0.1393)	-0.0211 (0.0178)	0.0010 (0.0515)
Wheat	-0.0069 (0.0701)	-0.0619 (0.2030)	-0.0116 (0.0664)	-0.0861 (0.1922)	-0.0046 (0.0113)	-0.0242 (0.0328)
Corn	0.0164 (0.0626)	0.0509 (0.1812)	0.0112 (0.0602)	0.0550 (0.1741)	-0.0052 (0.0118)	0.0041 (0.0343)
BeanOil	0.0131 (0.0481)	0.2267 (0.1391)	0.0060 (0.0473)	0.1851 (0.1368)	-0.0071 (0.0059)	-0.0415* (0.0172)
Cattle	0.0260 (0.0268)	-0.0574 (0.0713)	-0.0111 (0.0259)	-0.0445 (0.0688)	-0.0371* (0.0127)	0.0129 (0.0337)
Cocoa	0.0043 (0.0566)	0.0958 (0.1634)	0.0056 (0.0538)	0.0406 (0.1554)	0.0013 (0.0159)	-0.0552 (0.0458)
Coffee	0.0339 (0.0522)	-0.0663 (0.1507)	0.0184 (0.0508)	-0.0529 (0.1469)	-0.0155* (0.0057)	0.0133 (0.0164)
Cotton	-0.0035 (0.0579)	0.0731 (0.1882)	0.0227 (0.0526)	-0.0987 (0.1706)	0.0249 (0.0225)	-0.1920* (0.0728)
FedCattle	0.0146 (0.0230)	-0.0646 (0.0526)	0.0166 (0.0264)	-0.0761 (0.0605)	0.0019 (0.0114)	-0.0115 (0.0261)
Hogs	-0.0542 (0.0385)	0.0896 (0.0943)	-0.0752 (0.0411)	0.2428* (0.1006)	-0.0210 (0.0256)	0.1532* (0.0627)
KCWheat	0.0135 (0.0624)	-0.0748 (0.1808)	0.0086 (0.0601)	-0.0924 (0.1741)	-0.0049 (0.0115)	-0.0176 (0.0334)
Sugar	-0.0309 (0.0691)	0.4393 (0.2242)	-0.0255 (0.0624)	0.3293 (0.2026)	0.0053 (0.0175)	-0.1099 (0.0570)
Oil	0.0282 (0.0808)	-0.1863 (0.1509)	0.0127 (0.0749)	-0.1425 (0.1400)	-0.0155 (0.0213)	0.0438 (0.0399)

Notes to Table 2. Coefficient estimates (standard errors in parentheses) from OLS estimation of (15) using daily data from April 11, 2006 to December 30, 2011. * indicates significant at 5% level.

Table 3: Average returns during 10-day window around index roll versus normal times

	1 st contract		2 nd contract		spread	
	constant	S_τ	constant	S_τ	constant	S_τ
Beans	0.0201 (0.0548)	0.2220 (0.1235)	-0.0009 (0.0504)	0.2218 (0.1136)	-0.0209 (0.0186)	-0.0001 (0.0420)
Wheat	-0.0394 (0.0735)	0.1272 (0.1654)	-0.0440 (0.0695)	0.1122 (0.1566)	-0.0046 (0.0119)	-0.0150 (0.0267)
Corn	0.0069 (0.0656)	0.0794 (0.1477)	-0.0015 (0.0630)	0.0978 (0.1419)	-0.0083 (0.0124)	0.0184 (0.0279)
BeanOil	0.0109 (0.0503)	0.1482 (0.1134)	0.0035 (0.0495)	0.1244 (0.1115)	-0.0074 (0.0062)	-0.0238 (0.0140)
Cattle	0.0335 (0.0284)	-0.0660 (0.0585)	-0.0061 (0.0274)	-0.0477 (0.0565)	-0.0396* (0.0135)	0.0182 (0.0277)
Cocoa	0.0341 (0.0592)	-0.0928 (0.1332)	0.0338 (0.0563)	-0.1179 (0.1266)	-0.0003 (0.0166)	-0.0251 (0.0374)
Coffee	0.0093 (0.0547)	0.0839 (0.1228)	-0.0074 (0.0533)	0.0981 (0.1197)	-0.0167* (0.0059)	0.0142 (0.0134)
Cotton	-0.0121 (0.0600)	0.0997 (0.1519)	0.0194 (0.0545)	-0.0388 (0.1375)	0.0303 (0.0233)	-0.1504* (0.0587)
FedCattle	0.0242 (0.0250)	-0.0693 (0.0444)	0.0305 (0.0287)	-0.0902 (0.0511)	0.0063 (0.0124)	-0.0208 (0.0220)
Hogs	-0.0448 (0.0414)	0.0198 (0.0785)	-0.0663 (0.0442)	0.1138 (0.0839)	-0.0216 (0.0276)	0.0941 (0.0522)
KCWheat	-0.0184 (0.0654)	0.1164 (0.1473)	-0.0211 (0.0630)	0.0947 (0.1419)	-0.0027 (0.0121)	-0.0217 (0.0272)
Sugar	0.0081 (0.0716)	0.0173 (0.1812)	0.0059 (0.0647)	-0.0015 (0.1637)	-0.0021 (0.0182)	-0.0188 (0.0460)
Oil	0.1334 (0.0941)	-0.3333* (0.1364)	0.1000 (0.0873)	-0.2693* (0.1266)	-0.0334 (0.0249)	0.0641 (0.0361)

Notes to Table 3. Coefficient estimates (standard errors in parentheses) from OLS estimation of (15) using daily data from April 11, 2006 to December 30, 2011. * indicates significant at 5% level.

Table 4: Tests for predictability of crude oil returns

Oil: Masters				
	const	r_{t-1}	X_{t-1}	\bar{R}^2
level	-5.4784 (11.1640)	-0.0314 (0.0585)	0.0030 (0.0063)	-0.0052
1-wk diff	-0.0760 (0.3132)	0.1006 (0.1157)	-0.1340 (0.1045)	-0.0004
13-wk diff	-0.2298 (0.3064)	-0.1171 (0.0609)	0.0440* (0.0112)	0.0438
Oil: regression				
	const	r_{t-1}	X_{t-1}	\bar{R}^2
level	-4.4562 (14.9609)	-0.0306 (0.0587)	0.0025 (0.0085)	-0.0057
1-wk diff	-0.1147 (0.3128)	-0.0472 (0.0969)	0.0191 (0.0770)	-0.0057
13-wk diff	-0.1337 (0.3040)	-0.1240* (0.0609)	0.0448* (0.0107)	0.0499

Notes to Table 4. OLS coefficients (with standard errors in parentheses) for regressions (12), (13), and (14). Masters block uses $\tilde{x}_{t-1} = \ln(\tilde{X}_{\text{crude.oil},t}^{\text{Masters}})$ from equation(18) and $X_{t-1} = \tilde{x}_{t-1}$ for the level rows, $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-2})$ for the 1-wk diff rows, and $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-14})$ for the 13-wk diff rows. Regression block uses $\tilde{x}_{t-1} = \ln(\tilde{X}_{\text{crude.oil},t}^{\text{all}})$ from equation(21). All regressions estimated April 11, 2006 to January 3, 2012. * indicates significant at 5% level.

Table 5: In-sample and post-sample predictability of crude oil and stock market returns

	const	In-sample			Post-sample MSE	
		r_{t-1}	X_{t-1}	\bar{R}^2	regression	random walk
Oil	-0.4858 (0.3968)	-0.1507* (0.0755)	0.0553* (0.0124)	0.0834	24.0113	21.9747
S&P500	-0.1332 (0.2059)	-0.0637 (0.0724)	0.0125* (0.0060)	0.0130	6.5520	6.3688

Notes to Table 5. In-sample: OLS coefficients (standard errors in parentheses) for regression (14) as estimated over April 11, 2006 to January 12, 2010, with $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-14})$ for $\tilde{x}_{t-1} = \ln(\tilde{X}_{\text{crude_oil},t}^{\text{Masters}})$. Oil regression uses weekly percentage change in near crude oil contract for r_t and r_{t-1} . S&P500 uses percentage change in S&P500 for r_t and r_{t-1} . Post-sample MSE for regression reports mean squared error over January 17, 2010 to January 3, 2012 using the historically estimated regression. Post-sample MSE for random walk reports the average squared value of r_t over January 17, 2010 to January 3, 2012. * indicates significant at 5% level.

Figure 1: Price of near crude oil contract (left scale) and number of crude oil contracts held by index traders as imputed by Masters' method (right scale).






Figure 2: Adjusted R^2



Notes to Figure 2. Plot of \bar{R}^2 for oil return regression as a function of the interval in weeks over which the change in index-fund positions is calculated.

Figure 3: Holdings of crude oil contracts held by commodity index traders imputed by alternative methods



Notes to Figure 3. Number of oil contracts held by index funds as imputed using Masters' method (equation (18)), regression method (equation(21)), and bivariate inferences using soybean oil and one other commodity as specified in equation(20).