Received
CFTC

2015 FEB 27 PM 12: 56

Office of the

Christopher Kirkpatrick,
Secretary of the Commission,
Commodity Futures Trading
Commission, Three Lafayette Centre,
1155 21st Street NW., Washington, DC 20581

Please consider the enclosed comments on RIN 3038–AD99 for the Position Limits Proposal.

Jing Cynthia Wu
Assistant Professor of Econometrics and Statistics
The University of Chicago Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
Email: *Cynthia.Wu@ChicagoBooth.edu*
Webpage: *http://faculty.chicagobooth.edu/jing.wu/*

INTERNATIONAL ECONOMIC REVIEW
Vol. 56, No. 1, February 2015

EFFECTS OF INDEX-FUND INVESTING ON COMMODITY FUTURES PRICES*

BY JAMES D. HAMILTON AND JING CYNTHIA WU[1]

University of California, San Diego, U.S.A.; University of Chicago, U.S.A.

We develop a simple model of futures arbitrage that implies that if purchases by commodity index funds influence futures prices, then the notional positions of the index investors should help predict excess returns in these contracts. We find no evidence that the positions of index traders in agricultural contracts as identified by the Commodity Futures Trading Commission can help predict returns on the near futures contracts. Although there is some support that these positions might help predict changes in oil futures prices over 2006–2009, the relation breaks down out of sample.

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 et al. (2013), as well as our own review in Section 2, 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 producers 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 (2014), Knittel and Pindyck (2013), and Sockin and Xiong (2013). In this article we focus on the first link—by what mechanism could

*Manuscript received May 2013; revised December 2013.

[1] 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 article. Please address correspondence to: Cynthia Wu, The University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637. E-mail: *cynthia.wu@chicagobooth.edu.*

© (2015) by the Economics Department of the University of Pennsylvania and the Osaka University Institute of Social and Economic Research Association

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 (SCOT) 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 (2014) 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 although 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, 2011) related these correlations specifically to micro data on positions of different types of commodity traders. However, as Fattouh et al. (2013) 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 (2014) and Kilian and Lee (2014) 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 (forthcoming) 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.[2]

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 et al. (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 that these were of no help in predicting returns on 14 different commodity futures

[2] For further discussion see Kilian and Murphy (2012) and Fattouh et al. (2013).

contracts over 2006–2009. Sanders and Irwin (2011b) used proprietary CFTC data to extend the public 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 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 noncommercial 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 that in a pooled regression, 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 noncommercials, 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 (2014) 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, 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 et al. (1983), Chang (1985), Bessembinder (1992), and De Roon et al. (2000) provided empirical support for such an interpretation of the risk premium in commodity futures, whereas 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 et al. (2013), Cheng et al. (2012), and Hamilton and Wu (2014) proposed that the growing volume of commodity index investors could produce hedging price pressure on the buy side, with Hamilton and Wu (2012) finding that the average compensation to the long position in oil futures contracts has decreased, but become substantially more volatile since 2005. Etula (2013), Acharya et al. (2013), Danielsson et al. (2011), and Cheng et al. (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 (2014) in using a simple quadratic objective function as an approximation to a more detailed model of the capital limitations of potential arbitrageurs. Hamilton and Wu (2014) used this framework to infer risk prices from the predictability of futures returns based on their own lagged values. By contrast, in this article 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 that 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 her 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

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

In addition to potential positions in a variety of futures contracts, we presume that the arbitrageur also invests amounts q_{jt} in assets $j = 0, 1 \ldots, J$ (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 $\{q_{0t}, \ldots, q_{Jt}, z_{1t}, \ldots, z_{nt}\}$ so as to maximize

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

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,

$$f_{nt} = \log F_{nt} = \alpha_n + \beta_n' x_t \qquad n = 1, \ldots, N \tag{3}$$

$$r_{jt} = \xi_j + \psi_j' x_t \quad j = 1, \ldots, J,$$

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

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

Under these assumptions, Hamilton and Wu (2014) showed that

$$E_t(W_{t+1}) \approx q_{0t}(1 + r_{0,t+1}) + \sum_{j=1}^{J} q_{jt}\left[1 + \xi_j + \psi_j'(c + \rho x_t) + (1/2)\psi_j'\Sigma\Sigma'\psi_j\right]$$

$$+ \sum_{n=1}^{N} z_{nt}\left[\alpha_{n-1} + \beta_{n-1}'(c + \rho x_t) - \alpha_n - \beta_n' x_t + (1/2)\beta_{n-1}'\Sigma\Sigma'\beta_{n-1}\right], \tag{5}$$

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

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

(7) $$\alpha_{n-1} + \beta_{n-1}'(c + \rho x_t) - \alpha_n - \beta_n' x_t + (1/2)\beta_{n-1}'\Sigma\Sigma'\beta_{n-1} = \beta_{n-1}'\lambda_t,$$

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

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

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

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

then Equations (7) and (9) turn out to imply a recursion that the commodity-futures loadings α_n and β_n would satisfy that they 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 nth contract. In equilibrium, prices must be such so 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 (2014) 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 article, 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

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

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 price lower 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 that index-fund investors always desire a long position with notional exposure K_t in the two-month contract, selling their position as the month comes to a close in order to take a position K_{t+1} in the new two-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 being a function of both fundamentals x_t^* and index-fund buying K_t.[3] 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 $n > 0$.[4]

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.

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). Although 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

[3] In the particular example (11) as well as many instances of the general model (7) to (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.

[4] In the above-mentioned example in which index funds want only the two-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, fundamental uncertainty about the price of the underlying commodity. Assuming these fundamentals are positively serially correlated, if there was zero expected return associated with a one-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 two-period contract. Since there is no one to take to the other side of this one-period contract, equilibrium then requires a negative expected return on a one-period contract purchased at t just as on a two-period contract purchased at t. For details see Hamilton and Wu (2014).

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. Although the CFTC designations are not without some problems (Stoll and Whaley, 2010; Cheng et al., 2012), 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.[5] 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.[6]

Our interest centers on the predictability of r_t, the weekly return of a given contract.[7] 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 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, 2014), 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

[5] Daily futures prices were purchased from Norma's Historical Data (http://www.normashistoricaldata.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.

[6] Note that X_t is measured as 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.

[7] For most weeks, r_t is just $100(f_t - f_{t-1})$, where $f_t = \ln(F_t)$ is 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).

TABLE 1

TESTS FOR PREDICTABILITY OF AGRICULTURAL COMMODITY RETURNS

	Const	r_{t-1}	X_{t-1}	$\bar{R}^2$		Const	r_{t-1}	X_{t-1}	$\bar{R}^2$		Const	r_{t-1}	X_{t-1}	$\bar{R}^2$
Beans					**Wheat**					**Corn**				
Level	10.8033	−0.0032	−0.0056	−0.0026	level	30.8385	0.0014	−0.0166	0.0036	Level	6.4479	−0.0775	−0.0033	0.0005
	(9.6036)	(0.0579)	(0.0051)			(17.5923)	(0.0579)	(0.0094)			(13.5438)	(0.0580)	(0.0071)	
1-week diff	0.2640	−0.0443	0.0321	−0.0064	1-week diff	−0.0489	0.1817	−0.1738	−0.0027	1-week diff	0.1409	0.0691	−0.1352	0.0033
	(0.2306)	(0.1336)	(0.1017)			(0.3091)	(0.1814)	(0.1591)			(0.2959)	(0.1551)	(0.1305)	
13-week diff	0.2303	−0.0191	0.0082	−0.0044	13-week diff	−0.0570	−0.0008	−0.0058	−0.0062	13-week diff	0.1027	−0.1002	0.0147	0.0051
	(0.2345)	(0.0599)	(0.0099)			(0.3099)	(0.0596)	(0.0148)			(0.2967)	(0.0598)	(0.0116)	
Bean Oil					**Cattle**					**Cocoa**				
Level	8.9077	0.0098	−0.0058	−0.0024	Level	4.0293	−0.0628	−0.0024	−0.0013	Level	14.4661	0.0540	−0.0081	0.0067
	(7.7530)	(0.0579)	(0.0052)			(6.8600)	(0.0577)	(0.0042)			(8.1469)	(0.0577)	(0.0046)	
1-week diff	0.1748	−0.1026	0.1034	0.0006	1-week diff	0.0746	−0.1521	0.0717	0.0017	1-week diff	0.0965	0.1002	−0.0373	−0.0014
	(0.2294)	(0.0948)	(0.0704)			(0.1206)	(0.0971)	(0.0643)			(0.2586)	(0.0794)	(0.0442)	
13-week diff	0.1801	0.0046	0.0024	−0.0065	13-week diff	0.0740	−0.0738	0.0069	0.0002	13-week diff	0.1041	0.0580	−0.0033	−0.0034
	(0.2325)	(0.0594)	(0.0098)			(0.1210)	(0.0584)	(0.0077)			(0.2646)	(0.0589)	(0.0094)	
Coffee					**Cotton**					**Fed Cattle**				
Level	3.8530	−0.0329	−0.0024	−0.0050	Level	2.3353	0.0644	−0.0014	−0.0027	Level	5.1860	0.0298	−0.0038	−0.0032
	(9.2521)	(0.0582)	(0.0059)			(11.7720)	(0.0589)	(0.0076)			(5.7449)	(0.0574)	(0.0042)	
1-week-diff	0.1362	−0.0011	−0.0293	−0.0052	1-week diff	0.1601	−0.2249	0.2530*	0.0176	1-week diff	0.0245	0.0436	−0.0142	−0.0050
	(0.2370)	(0.1252)	(0.0960)			(0.2704)	(0.1293)	(0.1018)			(0.1124)	(0.0649)	(0.0271)	
13-week diff	0.1454	−0.0290	−0.0054	−0.0048	13-week diff	0.1482	0.0563	0.0046	−0.0022	13-week diff	0.0205	0.0258	0.0016	−0.0057
	(0.2379)	(0.0594)	(0.0116)			(0.2731)	(0.0603)	(0.0113)			(0.1130)	(0.0580)	(0.0066)	
Hogs					**KC Wheat**					**Sugar**				
Level	−10.9569	0.0530	0.0069	0.0018	Level	9.0292	0.0269	−0.0053	−0.0043	Level	−2.6872	−0.0411	0.0018	−0.0050
	(9.5632)	(0.0587)	(0.0061)			(12.0646)	(0.0582)	(0.0071)			(10.9718)	(0.0586)	(0.0072)	
1-week diff	−0.1522	0.1016	−0.0352	−0.0017	1-week diff	0.0755	−0.0062	0.0263	−0.0059	1-week diff	0.0805	0.2240	−0.2466	0.0077
	(0.2002)	(0.0950)	(0.0726)			(0.2862)	(0.1097)	(0.0840)			(0.3270)	(0.1458)	(0.1255)	
13-week diff	−0.1394	0.0510	0.0104	0.0010	13-week diff	0.0662	0.0152	0.0065	−0.0051	13-week diff	0.0189	−0.0477	0.0071	−0.0041
	(0.2004)	(0.0594)	(0.0102)			(0.2867)	(0.0595)	(0.0113)			(0.3301)	(0.0599)	(0.0125)	

NOTES: 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 one-week diff rows, and $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-14})$ for the 13-week diff rows. All regressions estimated from April 11, 2006, to January 3, 2012. *indicates significant at 5% level.

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 $\hat{\pi}$ is 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 (2014) 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 three-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 indices 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 indices 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 (i) τ is the 5th through the 10th business day of the month and (ii) the Goldman and/or 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 nonroll days. Moreover,

TABLE 2
AVERAGE RETURNS DURING INDEX-FUND ROLL VERSUS NORMAL TIMES

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

NOTES: 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.

because the average return over six days is numerically identical to (1/6) times the six-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 significantly 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 two of the estimated coefficients (namely, those for soybeans and hogs) are statistically significantly different from zero. And although 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

TABLE 3
AVERAGE RETURNS DURING 10-DAY WINDOW AROUND INDEX ROLL VERSUS NORMAL TIMES

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

NOTES: 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.

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 et al. (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 (i) τ is the 1st through the 10th business day of the month and (ii) the Goldman and/or 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 et al. (2012) found price effects that were usually reversed within 15 minutes, whereas 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, whereas Aulerich et al. (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 (2014) reported that a number of interesting variables appear to help predict returns on crude oil futures contracts. Our interest in this article 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} being 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.[8] 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 whereas 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,[9] so that $\tilde{X}_{it}^D = \delta_{it}^D \tilde{X}_t^D$. 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[10]:

$$\tilde{X}_{\text{crude\_oil},t}^{G,[\text{KC\_wheat}]} = \delta_{\text{crude\_oil},t}^G \tilde{X}_{\text{KC\_wheat},t}/\delta_{\text{KC\_wheat},t}^G. \tag{16}$$

[8] Note 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.

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

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

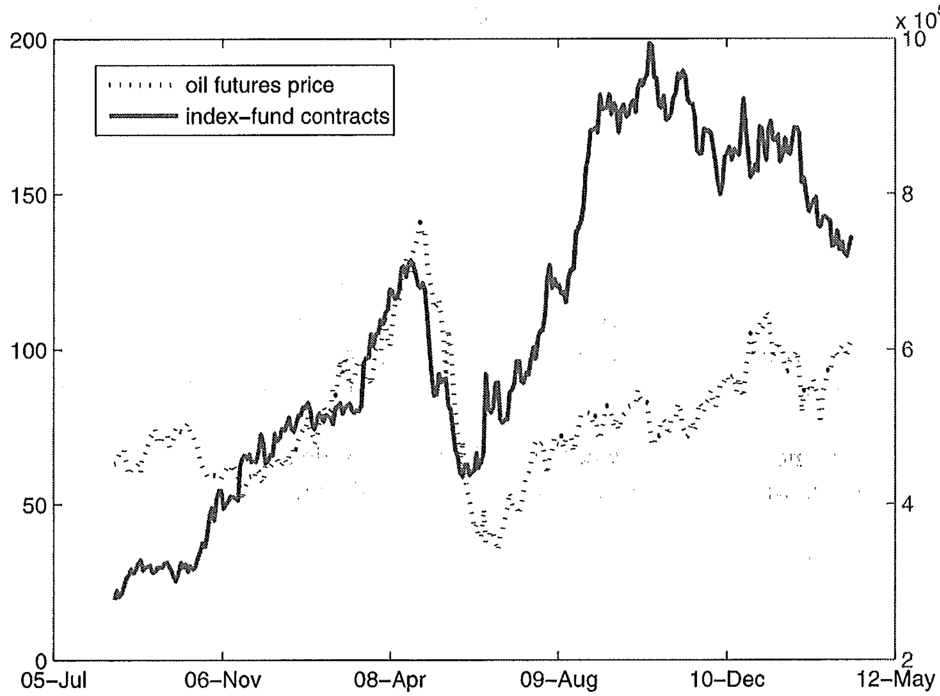

FIGURE 1

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)

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

$$(17) \qquad \tilde{X}^{G,[\text{feeder\_cattle}]}_{\text{crude\_oil},t} = \delta^{G}_{\text{crude\_oil},t}\tilde{X}_{\text{feeder\_cattle},t}/\delta^{G}_{\text{feeder\_cattle},t}.$$

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:

$$(18) \qquad \tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t} = \tilde{X}^{D,[\text{bean\_oil}]}_{\text{crude\_oil},t} + (1/2)\left(\tilde{X}^{G,[\text{KC\_wheat}]}_{\text{crude\_oil},t} + \tilde{X}^{G,[\text{feeder\_cattle}]}_{\text{crude\_oil},t}\right).$$

Figure 1 plots the number of contracts associated with this value for $\tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t}$ against the price of crude oil based on the near contract, updating similar figures in Masters (2008) and Singleton (2014). 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 $\tilde{x}_t = 100 \ln \tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t}$. 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 $\overline{R}^2$ 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.

TABLE 4
TESTS FOR PREDICTABILITY OF CRUDE OIL RETURNS

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

NOTES: OLS coefficients (with standard errors in parentheses) for regressions (12), (13), and (14). Masters block uses $\tilde{x}_t = 100\ln(\tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t})$ 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 one-week diff rows, and $X_{t-1} = (\tilde{x}_{t-1} - \tilde{x}_{t-14})$ for the 13-week diff rows. Regression block uses $\tilde{x}_t = 100\ln(\tilde{X}^{[\text{all}]}_{\text{crude\_oil},t})$ from Equation (21). All regressions are estimated from April 11, 2006, to January 3, 2012. *indicates significant at 5% level.

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 $\tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t}$ 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 $\tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t}$ as a measure of index-fund positions in crude oil futures. Irwin and Sanders noted first that although for some dates the two measures $\tilde{X}^{G,[\text{KC\_wheat}]}_{\text{crude\_oil},t}$ and $\tilde{X}^{G,[\text{feeder\_cattle}]}_{\text{crude\_oil},t}$ 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 indices but not the other. Note that the central claim is that

$$\tilde{X}_{it} = \delta^G_{it}\tilde{X}^G_t + \delta^D_{it}\tilde{X}^D_t, \tag{19}$$

where $\tilde{X}_{it}$, δ^G_{it}, and δ^D_{it} 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 indices, $\tilde{X}^G_t$ and $\tilde{X}^D_t$. Hence for any two arbitrary agricultural commodities i and j there exists a Masters-type estimate of crude oil holdings:

$$\tilde{X}^{[i,j]}_{\text{crude\_oil},t} = \begin{bmatrix} \delta^G_{\text{crude\_oil},t} & \delta^D_{\text{crude\_oil},t} \end{bmatrix} \begin{bmatrix} \delta^G_{it} & \delta^D_{it} \\ \delta^G_{jt} & \delta^D_{jt} \end{bmatrix}^{-1} \begin{bmatrix} \tilde{X}_{it} \\ \tilde{X}_{jt} \end{bmatrix}. \tag{20}$$

To illustrate the variability of such measures, we calculated (20), where j = soybean oil (one of the series used by Masters and Singleton) and i corresponding to any one of the other 11

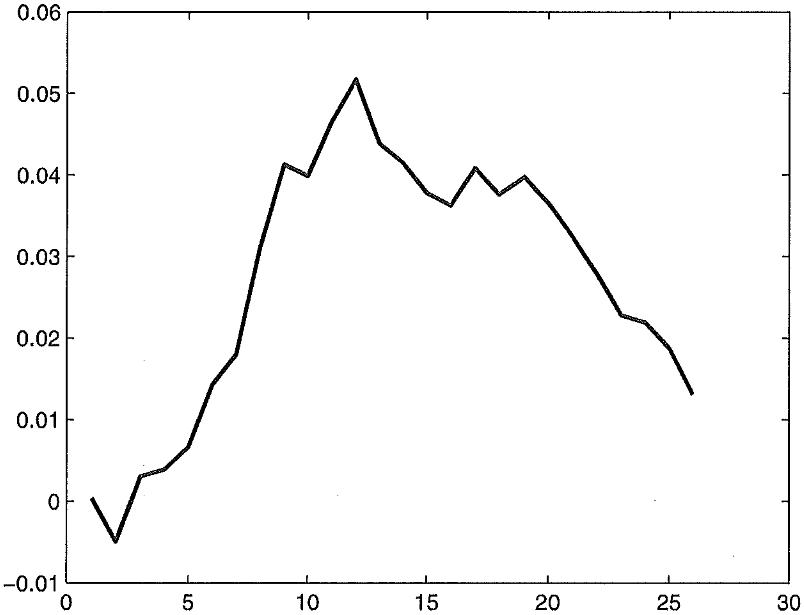

NOTES: Plot of $\bar{R}^2$ for oil return regression as a function of the interval in weeks over which the change in index-fund positions is calculated.

FIGURE 2

ADJUSTED R^2

possible choices. Plots of $\tilde{X}^{[i,\text{ bean\_oil}]}_{\text{crude\_oil},t}$ 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 $\tilde{X}^G_t$ and $\tilde{X}^D_t$ 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 $i = 1, \ldots, 12$ 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^G_{it} and δ^D_{it}:

$$(21) \qquad \tilde{X}^{[\text{all}]}_{\text{crude\_oil},t} = \left[\, \delta^G_{\text{crude\_oil},t} \;\; \delta^D_{\text{crude\_oil},t} \,\right] \begin{bmatrix} \sum_{i=1}^{12} (\delta^G_{it})^2 & \sum_{i=1}^{12} \delta^G_{it}\delta^D_{it} \\ \sum_{i=1}^{12} \delta^D_{it}\delta^G_{it} & \sum_{i=1}^{12} (\delta^D_{it})^2 \end{bmatrix}^{-1} \begin{bmatrix} \sum_{i=1}^{12} \delta^G_{it}\tilde{X}_{it} \\ \sum_{i=1}^{12} \delta^D_{it}\tilde{X}_{it} \end{bmatrix}.$$

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

We repeated our predictive regressions using $\tilde{X}^{[\text{all}]}_{\text{crude\_oil},t}$ in place of $\tilde{X}^{[\text{ Masters}]}_{\text{crude\_oil},t}$ 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 $\overline{R}^2$ are negative for the levels or first-difference specifications, but there is significant predictability from the 13-week change in $\tilde{X}^{[\text{all}]}_{\text{crude\_oil},t}$. 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 are interested in here. However, it is puzzling that agricultural positions cannot predict agricultural

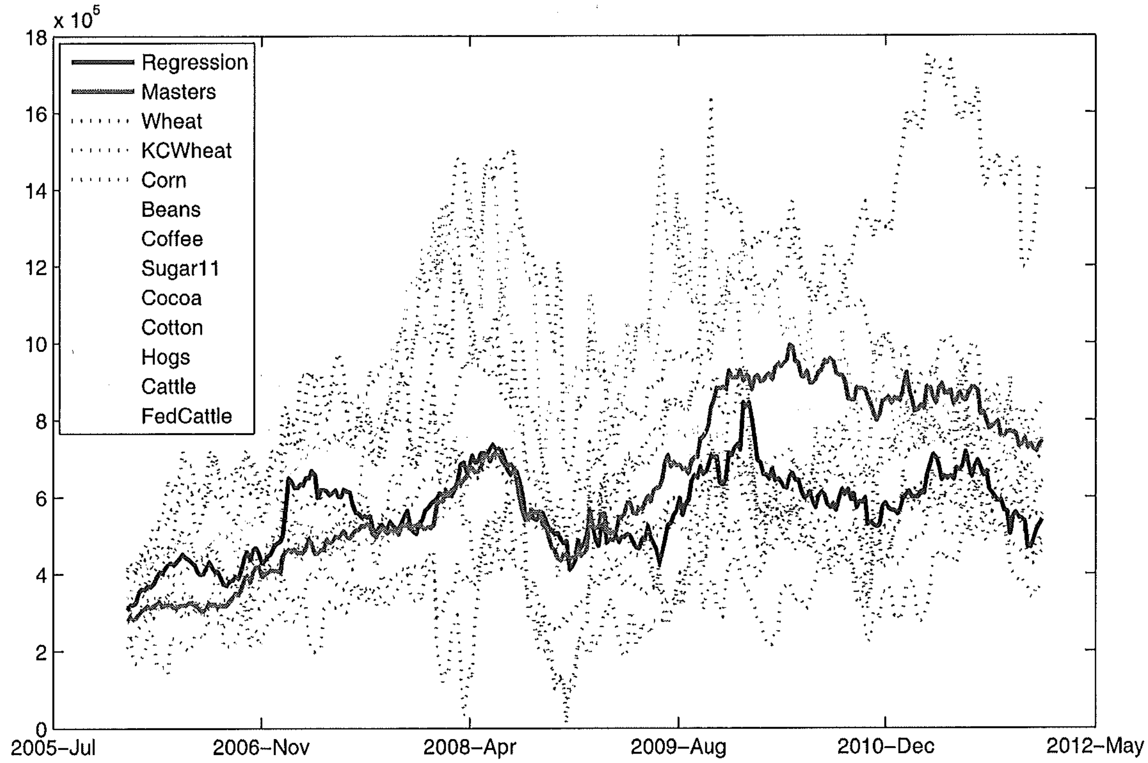

NOTES: Number of oil contracts held by index funds as imputed using Masters' method (Equation (18)), regression method (Equation (21)), and bivariate inferences using soybean oil and one other commodity as specified in Equation (20).

FIGURE 3

HOLDINGS OF CRUDE OIL CONTRACTS HELD BY COMMODITY INDEX TRADERS IMPUTED BY ALTERNATIVE METHODS

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 the time when Singleton's paper was first circulated we have obtained an additional two 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 (MSE) 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

TABLE 5
IN-SAMPLE AND POST-SAMPLE PREDICTABILITY OF CRUDE OIL AND STOCK MARKET RETURNS

	In-Sample				Post-Sample MSE	
	Const	r_{t-1}	X_{t-1}	$\bar{R}^2$	Regression	Random Walk
Oil	−0.4858	−0.1507*	0.0553*	0.0834	24.0113	21.9747
	(0.3968)	(0.0755)	(0.0124)			
S&P500	−0.1332	−0.0637	0.0125*	0.0130	6.5520	6.3688
	(0.2059)	(0.0724)	(0.0060)			

NOTES: 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 $\tilde{x}_t = 100\ln(\tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t})$. 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.

leads to the conclusion that there are two structural breaks[11] 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 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 is 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 $100[\ln \tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t-1} - \tilde{X}^{[\text{Masters}]}_{\text{crude\_oil},t-14}]$, 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.[12]

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,

[11] This 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 two structural breaks emerges from application of the Schwarz (1978) criterion.

[12] A 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.

and Hamilton and Wu (2014) found significant changes in the behavior of the risk premium on oil futures contracts before and after 2005. In this article, 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 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, V. V., L. A. Lochstoer, and T. Ramadorai, "Limits to Arbitrage and Hedging: Evidence from Commodity Markets," *Journal of Financial Economics* 109 (2013), 441–65.

Alquist, R., and O. Gervais, "The Role of Financial Speculation in Driving the Price of Crude Oil," Working Paper, Bank of Canada, 2011.

Ang, A., and M. Piazzesi, "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables," *Journal of Monetary Economics* 50 (2003), 745–87.

Aulerich, N. M., S. H. Irwin, and P. Garcia, "Bubbles, Food Prices, and Speculation: Evidence from the CFTC's Daily Large Trader Data Files," NBER Working Paper 19065, 2013.

Bai, J., and P. Perron, "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica* 66 (1998), 47–78.

Bessembinder, H., "Systemic Risk, Hedging Pressure, and Risk Premiums in Futures Markets," *Review of Financial Studies* 5 (1992), 637–67.

———, A. Carrion, L. Tuttle, and K. Venkataraman, "Predatory or Sunshine Trading? Evidence from Crude Oil ETF Rolls," Working Paper, Southern Methodist University, 2012.

Brunetti, C., and D. Reiffen, "Commodity Index Trading and Hedging Costs," Working Paper, Federal Reserve Board, 2011.

———, B. Buyuksahin, and J. H. Harris, "Speculators, Prices and Market Volatility," Working Paper, Johns Hopkins University, 2011.

Buyuksahin, B., and J. H. Harris, "Do Speculators Drive Crude Oil Futures Prices?" *Energy Journal* 32 (2011), 167–202.

———, and M. A. Robe, "Speculators, Commodities and Cross-Market Linkages," Working Paper, American University, 2010.

———, and ——— "Does 'Paper Oil' Matter?" Working Paper, American University, 2011.

Carter, C. A., G. C. Rausser, and A. Schmitz, "Efficient Asset Portfolios and the Theory of Normal Backwardation," *Journal of Political Economy* 91 (1983), 319–31.

CFTC, "Commitments of Traders: Explanatory Notes," (2012), http://www.cftc.gov/MarketReports/CommitmentsofTraders/ExplanatoryNotes/inde x.htm.

Chang, E. C., "Returns to Speculators and the Theory of Normal Backwardation," *Journal of Finance* 40 (1985), 193–208.

Cheng, I.-H., A. Kirilenko, and W. Xiong, "Convective Risk Flows in Commodity Futures Markets," Working Paper, Princeton University, 2012.

Cootner, P. H., "Returns to Speculators: Telser versus Keynes," *Journal of Political Economy* 68 (1960), 396–404.

Danielsson, J., H. S. Shin, and J. Zigrand, "Balance Sheet Capacity and Endogenous Risk," Working Paper, Princeton University, 2011.

De Roon, F., T. Nijman, and C. Veld, "Hedging Pressure Effects in Futures Markets," *Journal of Finance* 55 (2000), 1437–56.

Duffie, D., *Dynamic Asset Pricing Theory* (Princeton, NJ: Princeton University Press, 1992).
Etula, E., "Broker-Dealer Risk Appetite and Commodity Returns," *Journal of Financial Econometrics* 11 (2013), 486–521.
Fattouh, B., L. Kilian, and L. Mahadeva, "The Role of Speculation in Oil Markets: What Have We Learned So Far?" *Energy Journal* 34 (2013), 7–33.
Hamilton, J. D., "Causes and Consequences of the Oil Shock of 2007-08," *Brookings Papers on Economic Activity* (2009), 215–59.
———, and J. C. Wu, "Identification and Estimation of Gaussian Affine Term Structure Models," *Journal of Econometrics* 168 (2012), 315–31.
———, and ———, "Risk Premia in Crude Oil Futures Prices," *Journal of International Money and Finance* 42 (2014), 9–37.
Hansen, P. R., and A. Timmermann, "Equivalence between Out-of-Sample Forecast Comparisons and Wald Statistics," Working Paper, UCSD, 2013.
Hartzmark, M. L., "Returns to Individual Traders of Futures: Aggregate Results," *Journal of Political Economy* 95 (1987), 1292–1306.
Irwin, S. H., and D. R. Sanders, "Index Funds, Financialization, and Commodity Futures Markets," *Applied Economic Perspectives and Policy* 33 (2011), 1–31.
———, and ———, "Testing the Masters Hypothesis in Commodity Futures Markets," *Energy Economics* 34 (2012), 256–69.
Juvenal, L., and I. Petrella, "Speculation in the Oil Markets," Working Paper, Federal Reserve Bank of St. Louis, 2011.
Kennedy, J. P., "The High Cost of Gambling on Oil,"*New York Times*, April 10, 2012.
Keynes, J. M., *A Treatise on Money*, Volume 2 (London: Macmillan, 1930).
Kilian, L., and T. K. Lee, "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories," *Journal of International Money and Finance* 42 (2014), 71–87.
———, and D. P. Murphy, "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models," *Journal of the European Economic Association* 10 (2012), 1166–88.
———, and ———, "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics* 29 (2014), 454–78.
Knittel, C. R., and R. S. Pindyck, "The Simple Economics of Commodity Price Speculation," Working Paper, MIT, 2013.
Kolb, R. W., "Is Normal Backwardation Normal?" *Journal of Futures Markets* 12 (1992), 75–90.
Lombardi, M. J., and I. van Robays, "Do Financial Investors Destabilize the Oil Price?" Working Paper, European Central Bank, 2011.
Marcus, A. J., "Efficient Asset Portfolios and the Theory of Normal Backwardation: A Comment," *Journal of Political Economy* 92 (1984), 162–64.
Masters, M. W., "Testimony before Committee on Homeland Security and Governmental Affairs of the United States Senate," May 20, 2008.
———, "Testimony before Commodity Futures Trading Commission," August 25, 2009.
Mou, Y., "Limits to Arbitrage and Commodity Index Investment: Front-Running the Goldman Roll," Working Paper, Columbia University, 2010.
Sanders, D. R., and S. H. Irwin, "The Impact of Index Funds in Commodity Futures Markets: A Systems Approach," *Journal of Alternative Investments* 14 (2011a), 40–49.
———, and ———, "New Evidence on the Impact of Index Funds in U.S. Grain Futures Markets," *Canadian Journal of Agricultural Economics* 59 (2011b), 519–32.
Schwarz, G., "Estimating the Dimension of a Model," *Annals of Statistics* 6 (1978), 461–64.
Singleton, K. J., "Investor Flows and the 2008 Boom/Bust in Oil Prices," *Management Science* 60 (2014), 300–18.
Sockin, M., and W. Xiong, "Feedback Effects of Commodity Futures Prices," NBER Working Paper 18906, 2013.
Stoll, H. R., and R. E. Whaley, "Commodity Index Investing and Commodity Futures Prices," *Journal of Applied Finance* 20 (2010), 7–46.
Tang, K., and W. Xiong, "Index Investment and the Financialization of Commodities," *Financial Analysts Journal* 68 (2012), 54–74.

Journal of International Money and Finance 42 (2014) 9–37

Contents lists available at ScienceDirect

Journal of International Money and Finance

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

Risk premia in crude oil futures prices☆

James D. Hamilton [a,*], Jing Cynthia Wu [b]

[a] Department of Economics, University of California, San Diego, United States
[b] Booth School of Business, University of Chicago, United States

Keywords:
Oil prices
Speculation
Futures risk premium
Affine term structure models

ABSTRACT

If commercial producers or financial investors use futures contracts to hedge against commodity price risk, the arbitrageurs who take the other side of the contracts may receive compensation for their assumption of nondiversifiable risk in the form of positive expected returns from their positions. We show that this interaction can produce an affine factor structure to commodity futures prices, and develop new algorithms for estimation of such models using unbalanced data sets in which the duration of observed contracts changes with each observation. We document significant changes in oil futures risk premia since 2005, with the compensation to the long position smaller on average in more recent data. This observation is consistent with the claim that index-fund investing has become more important relative to commerical hedging in determining the structure of crude oil futures risk premia over time.

© 2013 Elsevier Ltd. All rights reserved.

☆ Financial Support from the University of Chicago Booth School of Business is gratefully acknowledged. We thank Christiane Baumeister, Rob Engle, Lutz Kilian and anonymous reviewers for helpful comments on earlier drafts of this paper, as well as participants in the International Conference on Understanding International Commodity Price Fluctuations and seminar participants at the Bank of Canada, New York University, Econometric Society Annual Meeting at San Diego, Energy Policy Institute at Chicago workshop, Applied Time Series Econometrics Workshop at Federal Reserve Bank of St. Louis, Econometric Society North American Summer Meeting at Chicago, SNDE 21st Annual Symposium at Milan, Econometric Society Australasian Meeting at Melbourne, and Peking University Guanghua.

* Corresponding author.
E-mail addresses: jhamilton@ucsd.edu (J.D. Hamilton), Cynthia.Wu@chicagobooth.edu (J.C. Wu).

0261-5606/$ – see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.jimonfin.2013.08.003

1. Introduction

Volatile oil prices have been drawing a lot of attention in recent years, with Hamilton (2009) for example suggesting that the oil price spike was a contributing factor in the recession of 2007–2009. There has been considerable interest in whether there is any connection between this volatility and the flow of dollars into commodity-index funds that take the long position in crude oil futures contracts. Recent empirical investigations of a possible link include Kilian and Murphy (2013), Tang and Xiong (2012), Buyuksahin and Robe (2011), Alquist and Gervais (2011), Mou (2010), Singleton (2011), Irwin and Sanders (2012), and Fattouh et al. (2013).

A separate question is the theoretical mechanism by which such an effect could arise in the first place. Keynes (1930) theory of normal backwardation proposed that if producers of the physical commodity want to hedge their price risk by selling futures contracts, then the arbitrageurs who take the other side of the contract may be compensated for assuming that risk in the form of a futures price below the expected future spot price. Empirical support for this view has come from Carter et al. (1983), Chang (1985), and De Roon et al. (2000), who interpreted the compensation as arising from the nondiversifiable component of commodity price risk, and from Bessembinder (1992), Etula (2013) and Acharya et al. (2013), who attributed the effect to capital limitations of potential arbitrageurs. In the modern era, buying pressure from commodity-index funds could exert a similar effect in the opposite direction, shifting the receipt of the risk premium from the long side to the short side of the contract.

In this paper we show that if arbitrageurs care about the mean and variance of their futures portfolio, then hedging pressure from commodity producers or index-fund investors can give rise to an affine factor structure to commodity futures prices. We do so by extending the models in Vayanos and Vila (2009) and Hamilton and Wu (2012a), which were originally used to describe how bond supplies affect relative yields, but are adapted in the current context to summarize how hedging demand would influence commodity futures prices. The result turns out to provide a motivation for specifications similar to the class of Gaussian affine term structure models originally developed by Vasicek (1977), Duffie and Kan (1996), Dai and Singleton (2002), Duffee (2002), and Ang and Piazzesi (2003) to characterize the relation between yields on bonds of different maturities. Related affine models have also been used to describe commodity futures prices by Schwartz (1997), Schwartz and Smith (2000), and Casassus and Collin-Dufresne (2006), among others.

In addition, this paper offers a number of methodological advances for use of this class of models to study commodity futures prices. First, we develop the basic relations directly for discrete-time observations, extending the contributions of Ang and Piazzesi (2003) to the setting of commodity futures prices. This allows a much more transparent mapping between model parameters and properties of observable OLS regressions. Second, we show how parameter estimates can be obtained directly from unbalanced data in which the remaining duration of observed contracts changes with each new observation, developing an alternative to the Kalman filter methodology used for this purpose by Cortazar and Naranjo (2006). Third, we show how the estimation method of Hamilton and Wu (2012b) provides diagnostic tools to reveal exactly where the model succeeds and where it fails to match the observed data.

We apply these methods to prices of crude oil futures contracts over 1990–2011. We document significant changes in risk premia in 2005 as the volume of futures trading began to grow significantly. While traders taking the long position in near contracts earned a positive return on average prior to 2005, that premium decreased substantially after 2005, becoming negative when the slope of the futures curve was high. This observation is consistent with the claim that historically commercial producers paid a premium to arbitrageurs for the privilege of hedging price risk, but in more recent periods financial investors have become natural counterparties for commercial hedgers. We also uncover seasonal variation of risk premia over the month, with changes as the nearest contract approaches expiry that cannot be explained from a shortening duration alone.

The plan of the paper is as follows. Section 2 develops the model, and Section 3 describes our approach to empirical estimation of parameters. Section 4 presents results for our baseline specification, while Section 5 presents results for a model allowing for more general variation as contracts near expiration. Conclusions are offered in Section 6.

2. Model

2.1. Role of arbitrageurs

Consider the incentives for a rational investor to become the counterparty to a commercial hedger or mechanical index-fund trader. We will refer to this rational investor as an arbitrageur, so named because the arbitrageur's participation guarantees that risk is priced consistently across all assets and futures contracts in equilibrium. Let F_{nt} denote the price of oil associated with an n-period futures contract entered into at date t. Let z_{nt} denote the arbitrageur's notional exposure (with $z_{nt} > 0$ denoting a long position and $z_{nt} < 0$ for short), so that z_{nt}/F_{nt} is the number of barrels purchased with n-period contracts. Following Duffie (1992, p. 39), we interpret a long position entered into at date t and closed at date $t+1$ as associated with a cash flow of zero at date t and $F_{n-1,t+1} - F_{nt}$ at date $t+1$. The arbitrageur's cash flow for period $t+1$ associated with the contemplated position z_{nt} is then $z_{nt}(F_{n-t,t+1} - F_{nt})/F_{nt}$. We assume the arbitrageur also takes positions q_{jt} in a set of other assets $j = 0, 1..., J$ with gross returns between t and $t+1$ denoted $\exp(r_{j,t+1})$ (so that the net return is approximately $r_{j,t+1}$) and where $r_{0,t+1}$ is assumed to be a risk-free yield. Then the arbitrageur's total wealth at $t+1$ will be

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

The arbitrageur is assumed to choose $\{q_{0t}, ..., q_{Jt}, z_{1t}, ..., z_{nt}\}$ so as to maximize[1]

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

subject to $\sum_{j=0}^{J} q_{jt} = W_t$.

We posit the existence of a vector of factors x_t that jointly determine all returns, which we assume follows a Gaussian vector autoregression (VAR)[2]:

$$x_{t+1} = c + \rho x_t + \Sigma u_{t+1} \quad u_t \sim \text{i.i.d. } N(0, I_m). \tag{3}$$

Log commodity prices and returns are assumed to be affine functions of these factors

$$f_{nt} = \log F_{nt} = \alpha_n + \beta_n' x_t \quad n = 1, ..., N \tag{4}$$

$$r_{jt} = \xi_j + \psi_j' x_t \quad j = 1, ..., J.$$

Using a similar approximation to that in Hamilton and Wu (2012a), we show in Appendix A that under these assumptions,

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

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

[1] It is trivial to extend this to adding positions in futures contracts for a number of alternative commodities. We discuss here the case of the single commodity oil for notational simplicity.

[2] The assumption of Gaussian homoskedastic errors greatly simplifies the estimation because it implies that parameters of the reduced-form representation of the model can be optimally estimated using simple OLS. For an extension of this approach to the case of non-Gaussian factors with time-varying variances, see Creal and Wu (2013).

The first-order conditions for the arbitrageur's positions satisfy

$$\frac{\partial E_t(W_{t+1})}{\partial q_{jt}} = 1 + r_{0,t+1} + (\gamma/2)\frac{\partial \mathrm{Var}_t(W_{t+1})}{\partial q_{jt}} \quad j = 1,\ldots,J$$
$$\frac{\partial E_t(W_{t+1})}{\partial z_{nt}} = (\gamma/2)\frac{\partial \mathrm{Var}_t(W_{t+1})}{\partial z_{nt}} \quad n = 1,\ldots,N.$$

Under (5) and (6) these become

$$\xi_j + \psi_j'(c + \rho x_t) + (1/2)\psi_j'\Sigma\Sigma'\psi_j = r_{0,t+1} + \psi_j'\lambda_t$$
$$\alpha_{n-1} + \beta_{n-1}'(c + \rho x_t) - \alpha_n - \beta_n' x_t + (1/2)\beta_{n-1}'\Sigma\Sigma'\beta_{n-1} = \beta_{n-1}'\lambda_t \tag{7}$$

for

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

Suppose we conjecture that in equilibrium the positions q_{jt}, z_{nt} selected by arbitrageurs are themselves affine functions of the vector of factors, so that

$$\lambda_t = \lambda + \Lambda x_t. \tag{9}$$

Then (7) requires

$$\beta_n' = \beta_{n-1}'\rho - \beta_{n-1}'\Lambda \tag{10}$$

$$\alpha_n = \alpha_{n-1} + \beta_{n-1}'c + (1/2)\beta_{n-1}'\Sigma\Sigma'\beta_{n-1} - \beta_{n-1}'\lambda. \tag{11}$$

From (5), the left side of (7) is the approximate expected return to a \$1 long position in an n-period contract entered at date t. Equation (7) thus characterizes equilibrium expected returns in terms of the price of risk λ_t:

$$E_t\left(\frac{F_{n-1,t+1} - F_{nt}}{F_{nt}}\right) \approx \beta_{n-1}'\lambda_t. \tag{12}$$

In the special case of risk-neutral arbitrageurs ($\gamma = 0$), from (8) we would have $\lambda = 0$ and $\Lambda = 0$ in (9).

Note that this framework allows for all kinds of factors (as embodied in the unobserved values of x_t) to influence commodity futures prices through Equation (4), including interest rates, fundamentals affecting supply and demand, and factors that might influence risk premia in other asset markets. If we consider physical inventory as another possible asset q_{jt}, this may help offset the risks associated with futures positions $z_{\ell t}$ as described in Equation (8) and could also be an element of the hypothesized factor vector x_t. We will demonstrate below that it is not necessary to have direct observations on the factor vector x_t itself in order to make use of the model's primary empirical implications (10) and (11). Instead, these restrictions can be represented solely in terms of implications for the dynamic behavior of the prices of different commodity-futures contracts that have to hold as a result of the factor structure itself and the behavior of the arbitrageurs. Moreover, we will see that it is possible to estimate the risk-pricing parameters λ and Λ solely on the basis of any predictabilities in the returns from positions in commodity-futures contracts.[3]

[3] Alternatively, one can try to make use of direct observations on the positions of commodity index-fund investors as we do in Hamilton and Wu (2012c).

The recursions (10) and (11) can equivalently be viewed as the equilibrium conditions that would result if risk-neutral arbitrageurs were to regard the factor dynamics as being governed not by (3) but instead by

$$x_{t+1} = c^Q + \rho^Q x_t + \Sigma u_{t+1}^Q \tag{13}$$

$$c^Q = c - \lambda \tag{14}$$

$$\rho^Q = \rho - \Lambda \tag{15}$$

$$u_{t+1}^Q \overset{Q}{\sim} N(0, I_m).$$

The recursions (10) and (11) that characterize the relation between the prices of futures contracts of different maturities will be recognized as similar to those that have been developed in the affine term structure literature[4] to characterize the relations that should hold in equilibrium between the interest rates on assets of different maturities. In addition to providing a derivation of how these relations can be obtained in the case of commodity futures contracts, the derivation above demonstrates how commercial hedging or commodity-index funds might be expected to influence commodity futures prices. An increase in the demand for long positions in contract n will require in equilibrium a price process in which arbitrageurs are persuaded to take a corresponding short position in exactly that amount. A larger absolute value of z_{nt} in turn will expose arbitrageurs to different levels of risk which would change the equilibrium compensation to risk λ_t according to equation (8). Again, from (8) and (9), these index traders could be responding through an affine function to interest rates or other economic fundamentals. What matters is that this behavior causes the net risk exposure of arbitrageurs λ_t to be an affine function of the factors in equilibrium. In the following subsection we illustrate this potential effect using a simple example.

2.2. *Example of the potential role of index-fund traders*

Suppose there are some investors who always want to have a long position in the 2-period contract, regardless of anything happening to fundamentals. At the start of each new period, these investors close out their previous position (which is now a 1-period contract) and replace it with a new long position in what is now the current 2-period contract.[5] Let the scalar K_t denote the notional value of 2-period contracts that investors want to buy in period t, and suppose this evolves exogenously according to

$$K_t = c^K + \rho^K K_{t-1} + \Sigma^K u_t^K. \tag{16}$$

If investors and arbitrageurs are the only participants in the market, then equilibrium futures prices must be such as to persuade arbitrageurs to take the opposite side of the investors. Thus arbitrageurs are always short the two-period contract, close that position when it becomes a 1-period contract, take the short side of the new 2-period contract, and have zero net exposure to any other contract in equilibrium. In other words, the process for $\{f_{nt}\}_{n=0}^N$ must be such that (7) and (8) are satisfied with

[4] Our recursions (10) and (11) are essentially the same as Equation (17) in Ang and Piazzesi (2003), with the important difference being that their recursion for the intercept adds a term δ_0 for each n, corresponding to the interest earned each period. No such term appears in our expression because there is no initial capital invested. Another minor notational difference is that our λ corresponds to their $\Sigma\lambda_0$ while our Λ corresponds to their $\Sigma\lambda_1$. An advantage of our notation in the current setting is that our λ_t is then measured in the same units as x_t and is immediately interpreted as the direct adjustment to c and ρ that results from risk aversion by arbitrageurs.

[5] In the case of crude oil contracts, what typically happens is that the commodity-index fund takes a long position from a swap dealer which in turn hedges its exposure by taking a long position in an organized exchange contract. We view the swap fund in such an arrangement as simply an intermediary, with the ultimate demand for the long position (K_t) coming from the commodity-index fund and the index-fund's ultimate counterparty being the short on the organized exchange contract (z_{2t}).

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

Suppose that arbitrageurs' only risk exposure comes from commodities ($q_{jt} = 0$ for $j = 1,...,J$). Then from (8), in equilibrium we will have

$$\lambda_t = -\gamma\Sigma\Sigma'\beta_1 K_t. \tag{17}$$

Suppose that the spot price depends solely on a scalar "fundamentals" factor x_t^*:

$$f_{0t} = x_t^* \tag{18}$$

$$x_t^* = c^* + \rho^* x_{t-1}^* + \Sigma^* u_t^*.$$

We conjecture that in equilibrium, the factor x_t governing futures prices includes both fundamentals and the level of index-fund investment, $x_t = (x_t^*, K_t)'$, with (18) implying $\beta_0' = (1, 0)$ and the factor evolving according to

$$x_t = c + \rho x_{t-1} + \Sigma u_t$$

or written out explicitly,

$$\begin{bmatrix} x_t^* \\ K_t \end{bmatrix} = \begin{bmatrix} c^* \\ c^K \end{bmatrix} + \begin{bmatrix} \rho^* & 0 \\ 0 & \rho^K \end{bmatrix} \begin{bmatrix} x_{t-1}^* \\ K_{t-1} \end{bmatrix} + \begin{bmatrix} \Sigma^* & 0 \\ 0 & \Sigma^K \end{bmatrix} \begin{bmatrix} u_t^* \\ u_t^K \end{bmatrix}.$$

We can then recognize (17) as a special case of (9) with

$$\lambda = 0$$

$$\underset{(2\times2)}{\Lambda} = \begin{bmatrix} 0 & -\gamma\Sigma\Sigma'\beta_1 \end{bmatrix}.$$

Hence

$$\rho^Q = \begin{bmatrix} \rho^* & 0 \\ 0 & \rho^K \end{bmatrix} + \begin{bmatrix} 0 & \gamma\Sigma\Sigma'\beta_1 \end{bmatrix} \tag{19}$$

$$\begin{aligned} \beta_1' &= \beta_0'\rho^Q \\ &= \begin{bmatrix} 1 & 0 \end{bmatrix}\rho^Q \\ &= \begin{bmatrix} \rho^* & \gamma\begin{bmatrix} (\Sigma^*)^2 & 0 \end{bmatrix}\beta_1 \end{bmatrix} \end{aligned}$$

$$\beta_1 = \begin{bmatrix} \rho^* \\ \gamma\rho^*(\Sigma^*)^2 \end{bmatrix}. \tag{20}$$

Assuming $\rho^* > 0$ and $K_t > 0$, the effect of index-fund buying of the 2-period contract is also to increase the price of a 1-period contract. The reason is that the 2-period contract that the arbitrageurs are currently being induced to short exposes the arbitrageurs to risk associated with uncertainty about the value of x_{t+1}^*. The 1-period contract is also exposed to risk from x_{t+1}^*. If a 1-period contract purchased at t provided zero expected return, arbitrageurs would want to go long the 1-period contract in order to diversify their risk associated with being short the 2-period contract. But there is no counterparty who wants to short the 1-period contract, so equilibrium requires a price f_{1t} such that someone shorting the 1-period contract would also have a positive expected return, earned in the form of a higher price for f_{1t}.

Table 1
Weekly durations associated with monthly contracts at different points in time.

j	$k = 0$	$k = 1$	$k = 2$
1	3	7	11
2	2	6	10
3	1	5	9
4	0	4	8

For specified week of the month j and months until the contract expires k, table entry indicates weeks n remaining until expiry.

Substituting (20) into (19), we now know ρ^Q and can calculate $\beta_n = (\rho^{Q\prime})^n\beta_0$ for each n. Thus investment buying does not matter for f_{0t} but does affect every f_{nt} for $n > 0$, through the same mechanism as operates on the 1-period contract. In particular, from (12),

$$E_t\left(\frac{F_{n-1,t+1} - F_{nt}}{F_{nt}}\right) \approx -\gamma\beta_1'\left(\rho^Q\right)^{n-2}\Sigma\Sigma'\beta_1 K_t,$$

which in general has the opposite sign of K_t for all n; someone would earn a positive expected return by taking the short position in a contract of any duration.

2.3. Empirical implementation

There are two general strategies for empirical implementation of this framework. The first is to make direct use of data on the positions of different types of traders. Hamilton and Wu (2012c) use this approach to study agricultural futures prices. Unfortunately, the data publicly available on trader positions in crude oil futures contracts have some serious problems (see the discussion in Irwin and Sanders (2012) and Hamilton and Wu (2012c)). An alternative approach, which we adopt for purposes of modeling crude oil futures prices in this paper, is to infer the factors x_t based on the behavior of the futures prices themselves. In this case, risk premia are identified from differences between observed futures prices and a rational expectation of future prices. We will use the framework to characterize the dynamic behavior of risk premia and their changes over time.

For purposes of empirical estimation we interpret t as describing weekly intervals. This allows us to capture some key calendar regularities in the data without introducing an excessive number of parameters. NYMEX crude oil futures contracts expire on the third business day prior to the 25th calendar day of the month prior to the month on which the contract is written. To preserve the important calendar structure of the raw data, we divide the "month" leading up to a contract expiry into four "weeks", defined as follows:

week 1 ends on the last business day of the previous calendar month
week 2 ends on the 5th business day of the current calendar month
week 3 ends on the 10th business day of the current calendar month
week 4 ends on the day when the near contract expires

Associated with any week t is an indicator $j_t \in \{1, 2, 3, 4\}$ of where in the month week t falls.

Our estimation uses the nearest three contracts. If we interpret the price at expiry as an $n = 0$ week-ahead contract, the observation y_t for week t would be characterized using the notation of Section 2 as follows:

$$y_t = \begin{cases} \left(f_{3t}, f_{7t}, f_{11,t}\right)' & \text{if } j_t = 1 \\ \left(f_{2t}, f_{6t}, f_{10,t}\right)' & \text{if } j_t = 2 \\ \left(f_{1t}, f_{5t}, f_{9t}\right)' & \text{if } j_t = 3 \\ \left(f_{0t}, f_{4t}, f_{8t}\right)' & \text{if } j_t = 4 \end{cases}.$$

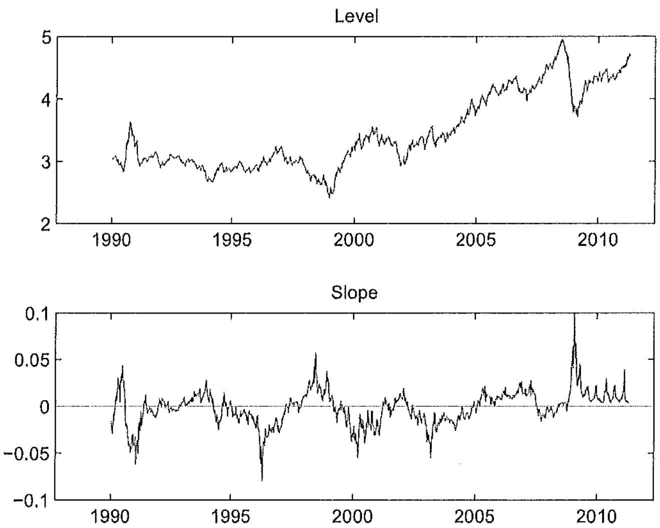

Fig. 1. Data used in the analysis. Weekly observations (with specification of weeks as given in text), January 1990 to June 2011. Top panel: first element of y_{1t} (the average of the log prices of second and third contracts). Bottom panel: second element of y_{1t} (the difference between the log price of third and second contracts).

Table 1 summarizes the relation between the weekly indicator (j), months until expiry is reached (k), and weeks remaining until expiry (n). This feature that the maturity of observed contracts changes with each observation t is one reason that much of the research with commodity futures contracts has used monthly data. However, in our application a key interest is in the higher-frequency movements and specific calendar effects. Fortunately, the framework developed in Section 2 gives us an exact description of the likelihood function for the data as actually observed, as we now describe.

We will assume that there are two underlying factors (that is, x_t is 2×1) Since (4) implies that each element of the (3×1) vector y_t could be written as an exact linear function of x_t, the system as written is stochastically singular – according to the model, the third element of y_t should be given by an exact linear combination of the first two. This issue also commonly arises in studies of the term structure of interest rates. A standard approach in that literature[6] is to assume that some elements or linear combinations of y_t differ from the magnitude predicted in (4) by a measurement or specification error. In the results reported below, we assume that the $k = 1$ - and 2- month contracts are priced exactly as the model predicts. It is helpful for purposes of interpreting parameter estimates to summarize the information in these contracts in terms of the average level of the two prices, which we will associate with the first factor in the system, and spread between them, which we will associate with the second factor:

$$y_{1t} = H_1 y_t$$

$$H_1 = \begin{bmatrix} 0 & (1/2) & (1/2) \\ 0 & -1 & 1 \end{bmatrix}. \tag{21}$$

The two elements of y_{1t} are plotted in Fig. 1.

We assume that the model correctly characterizes these two observed magnitudes. Since $y_t = (f_{4-j_t}, f_{8-j_t}, f_{12-j_t})'$, this implies that

[6] See for example Chen and Scott (1993), Ang and Piazzesi (2003) and Joslin et al. (2011). The observable implications of this assumption are explored in detail in Hamilton and Wu (2013).

$$y_{1t} = A_{1j_t} + B_{1j_t}x_t \tag{22}$$

$$\underset{(2\times 1)}{A_{1j_t}} = H_1\begin{bmatrix}\alpha_{4-j_t}\\ \alpha_{8-j_t}\\ \alpha_{12-j_t}\end{bmatrix} \quad \text{for } j_t = 1,2,3, \text{ or } 4$$

$$\underset{(2\times 2)}{B_{1j_t}} = H_1\begin{bmatrix}\beta'_{4-j_t}\\ \beta'_{8-j_t}\\ \beta'_{12-j_t}\end{bmatrix} \quad \text{for } j_t = 1,2,3, \text{ or } 4. \tag{23}$$

We will use the notational convention that if $j_t = 1$, then $A_{1j_{t-1}} = A_{14}$.

If B_{1j} is invertible, the dynamics of the observed vector y_{1t} can be characterized by substituting (22) into (3):

$$y_{1t} = A_{1j_t} + B_{1j_t}c + B_{1j_t}\rho\left[B^{-1}_{1j_{t-1}}\left(y_{1,t-1} - A_{1j_{t-1}}\right)\right] + B_{1j_t}\Sigma u_t.$$

Since u_t is independent of $\{y_{t-1}, y_{t-2}, \ldots, y_0\}$, this means that the density of y_{1t} conditional on all previous observations is characterized by a VAR(1) with seasonally varying parameters:

$$y_{1t}|y_{t-1}, y_{t-2}, \ldots, y_0 \sim N\left(\phi_{j_t} + \Phi_{j_t}y_{1,t-1}, \Omega_{j_t}\right) \tag{24}$$

$$\Omega_{j_t} = B_{1j_t}\Sigma\Sigma' B'_{1j_t}$$

$$\Phi_{j_t} = B_{1j_t}\rho B^{-1}_{1j_{t-1}}$$

$$\phi_{j_t} = A_{1j_t} + B_{1j_t}c - \Phi_{j_t}A_{1j_{t-1}}.$$

Note that the predicted seasonal parameter variation arises from the fact that the number of weeks remaining until expiry of the observed contracts changes with each new week.

We postulate that the nearest contract, which we write as

$$y_{2t} = H_2 y_t$$

$$H_2 = [1 \quad 0 \quad 0], \tag{25}$$

differs from the value predicted by the framework by a measurement or specification error with mean zero and variance $\sigma^2_{ej_t}$:

$$y_{2t} = A_{2j_t} + B_{2j_t}x_t + \sigma_{ej_t}u_{e,t}$$

$$\underset{(1\times 1)}{A_{2j_t}} = H_2\begin{bmatrix}\alpha_{4-j_t}\\ \alpha_{8-j_t}\\ \alpha_{12-j_t}\end{bmatrix} \quad \text{for } j = 1,2,3,4$$

$$\underset{(1\times 2)}{B_{2j_t}} = H_2\begin{bmatrix}\beta'_{4-j_t}\\ \beta'_{8-j_t}\\ \beta'_{12-j_t}\end{bmatrix} \quad \text{for } j_t = 1,2,3, \text{ or } 4.$$

If the measurement error u_{et} is independent of past observations, this gives the conditional distribution

$$y_{2t}|y_{1t}, y_{t-1}, y_{t-2}, \ldots, y_0 \sim N\left(\gamma_{j_t} + \Gamma_{j_t} y_{1t}, \sigma^2_{e,j_t}\right) \tag{26}$$

$$\Gamma_{j_t} = B_{2,j_t} B^{-1}_{1,j_t} \tag{27}$$

$$\gamma_{j_t} = A_{2,j_t} - \Gamma_{j_t} A_{1,j_t}.$$

The density of y_t conditional on its own past history is thus the product of (24) with (26), meaning that the log likelihood for the full sample of observations $(y'_T, y'_{T-1}, \ldots, y'_1)'$ conditional on the initial observation y_0 is given by

$$\mathscr{L} = \sum_{t=1}^{T} \left[\log g\left(y_{1t}; \phi_{j_t} + \Phi_{j_t} y_{1,t-1}, \Omega_{j_t}\right) + \log g\left(y_{2t}; \gamma_{j_t} + \Gamma_{j_t} y_{1t}, \sigma^2_{e,j_t}\right]\right. \tag{28}$$

where $g(y; \mu, \Omega)$ denotes the multivariate Normal density with mean μ and variance Ω evaluated at the point y.

3. Estimation

3.1. Unrestricted reduced form

The traditional approach to estimation of these kind of models would be to maximize the likelihood function with respect to the unknown structural parameters. However, Hamilton and Wu (2012b) demonstrate that there can be big benefits from using an estimator that turns out to be asymptotically equivalent to MLE but is derived from simple OLS regressions. To understand this estimator, consider first how we would maximize the likelihood if we thought of ϕ_j, Φ_j, Ω_j, γ_j, Γ_j, and σ_{ej} in the above representation as completely unrestricted parameters rather than the particular values implied by the structural model presented above. From this perspective, the log likelihood (28) could be written

$$\mathscr{L}(\phi_1, \Phi_1, \Omega_1, \gamma_1, \Gamma_1, \sigma_{e1}, \ldots, \phi_4, \Phi_4, \Omega_4, \gamma_4, \Gamma_4, \sigma_{e4}) = \sum_{j=1}^{4} \mathscr{L}_{1j}\left(\phi_j, \Phi_j, \Omega_j\right) + \sum_{j=1}^{4} \mathscr{L}_{2j}\left(\gamma_j, \Gamma_j, \sigma_{ej}\right) \tag{29}$$

$$\mathscr{L}_{1j}\left(\phi_j, \Phi_j, \Omega_j\right) = \sum_{t=1}^{T} \delta(j_t = j) \log g\left(y_{1t}; \phi_j + \Phi_j y_{1,t-1}, \Omega_j\right)$$

$$\log g\left(y_{1t}; \phi_j + \Phi_j y_{1,t-1}, \Omega_j\right) = -\log 2\pi - (1/2)\log|\Omega_j| - (1/2)\left(y_{1t} - \phi_j - \Phi_j y_{1,t-1}\right)' \Omega_j^{-1} \left(y_{1t} - \phi_j - \Phi_j y_{1,t-1}\right)$$

$$\mathscr{L}_{2j}\left(\gamma_j, \Gamma_j, \sigma_{ej}\right) = \sum_{t=1}^{T} \delta(j_t = j) \log g\left(y_{2t}; \gamma_j + \Gamma_j y_{1t}, \sigma^2_{ej}\right)$$

$$\log g\left(y_{2t}; \gamma_j + \Gamma_j y_{1t}, \sigma^2_{ej}\right) = -(1/2)\log 2\pi - (1/2)\log \sigma^2_{ej} - \frac{\left(y_{2t} - \gamma_j - \Gamma_j y_{1t}\right)^2}{2\sigma^2_{ej}}$$

where for example $\delta(j_t = 1)$ is 1 if t is in the first week of the month and is zero otherwise. It is clear that the unconstrained likelihood function is in fact maximized by a series of OLS regressions. To estimate the parameters in block j, we collect all observations whose left-hand variable is in the jth week of the month, and simply perform OLS regressions on what now looks like a monthly data set.

Specifically, to estimate $(\phi_j, \Phi_j, \Omega_j)$ for a particular j, we associate month τ with an observed monthly-frequency vector $y^{\dagger}_{1,j,\tau}$ defined as follows. For illustration, consider $j = 1$ and suppose that τ corresponds to the month spanned by the last week of December and first 3 weeks of January. The first element of $y^{\dagger}_{1,1,\tau|\text{month}(\tau)=\text{Jan}}$ is the average of the log prices of the March and April contracts as of the last business day of December. The second element of $y^{\dagger}_{1,1,\tau|\text{month}(\tau)=\text{Jan}}$ is based on the log price of the April contract on the last day of December minus the log price of March contract. For $j = 1$ and general τ,

$$y^{\dagger}_{1,1,\tau} = H_1 \begin{bmatrix} f_{3,t(\tau)} \\ f_{7,t(\tau)} \\ f_{11,t(\tau)} \end{bmatrix}$$

for H_1 given in (21) and where $t(\tau)$ denotes the week t associated with month τ. The explanatory variables in these $j = 1$ block regressions consist of a constant, the average log prices of the February and March contracts on the day in December when the January contract expired, and the spread between the March price and February price at the December expiry of the January contract:

$$x^{\dagger}_{1,1,\tau} = \begin{bmatrix} 1 \\ H_1 \begin{bmatrix} f_{0,t(\tau)-1} \\ f_{4,t(\tau)-1} \\ f_{8,t(\tau)-1} \end{bmatrix} \end{bmatrix}. \tag{30}$$

Consider the estimates from OLS regression of $y^{\dagger}_{1,1,\tau}$ on $x^{\dagger}_{1,1,\tau}$,

$$\begin{bmatrix} \hat{\phi}_1 & \hat{\Phi}_1 \end{bmatrix} = \left(\sum_{\tau=1}^{\mathcal{T}} y^{\dagger}_{1,1,\tau} x^{\dagger\prime}_{1,1,\tau} \right) \left(\sum_{\tau=1}^{\mathcal{T}} x^{\dagger}_{1,1,\tau} x^{\dagger\prime}_{1,1,\tau} \right)^{-1}$$

$$\hat{\Omega}_1 = \mathcal{T}^{-1} \sum_{\tau=1}^{\mathcal{T}} \left(y^{\dagger}_{1,1,\tau} - \begin{bmatrix} \hat{\phi}_1 & \hat{\Phi}_1 \end{bmatrix} x^{\dagger}_{1,1,\tau} \right) \left(y^{\dagger}_{1,1,\tau} - \begin{bmatrix} \hat{\phi}_1 & \hat{\Phi}_1 \end{bmatrix} x^{\dagger}_{1,1,\tau} \right)'$$

where $\mathcal{T}$ denotes the number of months in the sample. These estimates maximize the log likelihood (29) with respect to $\{\phi_1, \Phi_1, \Omega_1\}$

For $j = 2$ we regress $y^{\dagger}_{1,2,\tau}$ (whose first element, for example, would be the average of the March and April contracts as of the fifth business day in January) on $x^{\dagger}_{1,2,\tau}$ (e.g., a constant and the level and spread as of the last day of December),

$$\begin{bmatrix} \hat{\phi}_2 & \hat{\Phi}_2 \end{bmatrix} = \left(\sum_{\tau=1}^{\mathcal{T}} y^{\dagger}_{1,2,\tau} x^{\dagger\prime}_{1,2,\tau} \right) \left(\sum_{\tau=1}^{\mathcal{T}} x^{\dagger}_{1,2,\tau} x^{\dagger\prime}_{1,2,\tau} \right)^{-1}$$

$$\hat{\Omega}_2 = \mathcal{T}^{-1} \sum_{\tau=1}^{\mathcal{T}} \left(y^{\dagger}_{1,2,\tau} - \begin{bmatrix} \hat{\phi}_2 & \hat{\Phi}_2 \end{bmatrix} x^{\dagger}_{1,2,\tau} \right) \left(y^{\dagger}_{1,2,\tau} - \begin{bmatrix} \hat{\phi}_2 & \hat{\Phi}_2 \end{bmatrix} x^{\dagger}_{1,2,\tau} \right)',$$

to obtain $\hat{\phi}_2$, $\hat{\Phi}_2$, and $\hat{\Omega}_2$. Similar separate monthly regressions of the 1- and 2-month prices in the third or fourth week of each month on their values the week before produce $\{\hat{\phi}_j, \hat{\Phi}_j, \hat{\Omega}_j\}$ for $j = 3$ or 4.

Likewise, note that the components of $\sum_{j=1}^{4} \mathcal{L}_{2j}(\gamma_j, \Gamma_j, \sigma_e)$ take the form of regressions in which the residuals are uncorrelated across blocks, meaning full-information maximum likelihood estimates of γ_j

and Γ_j are obtained by OLS regressions for individual j. For example, for $j = 1$ and τ corresponding to December–January, $y^{\dagger}_{2,j,\tau}$ is the price of the February contract on the last day of December,

$$y^{\dagger}_{2,1,\tau} = H_2 \begin{bmatrix} f_{3,t(\tau)} \\ f_{7,t(\tau)} \\ f_{11,t(\tau)} \end{bmatrix},$$

for H_2 in (25) and explanatory variables the level and slope as of the last day of December:

$$x^{\dagger}_{2,1,\tau} = \begin{bmatrix} 1 \\ H_1 \begin{bmatrix} f_{3,t(\tau)} \\ f_{7,t(\tau)} \\ f_{11,t(\tau)} \end{bmatrix} \end{bmatrix}.$$

The maximum likelihood estimates are given by

$$\begin{aligned} \begin{bmatrix} \widehat{\gamma}_j & \hat{\Gamma}_j \end{bmatrix} &= \left(\sum_{\tau=1}^{\mathcal{T}} y^{\dagger}_{2,j,\tau} x^{\dagger\prime}_{2,j,\tau} \right) \left(\sum_{\tau=1}^{\mathcal{T}} x^{\dagger}_{2,j,\tau} x^{\dagger\prime}_{2,j,\tau} \right)^{-1} \quad \text{for } j = 1,2,3,4 \\ \widehat{\sigma}^2_{ej} &= \mathcal{T}^{-1} \sum_{\tau=1}^{\mathcal{T}} \left(y^{\dagger}_{2,j,\tau} - \begin{bmatrix} \widehat{\gamma}_j & \hat{\Gamma}_j \end{bmatrix} x^{\dagger}_{2,j,\tau} \right)^2. \end{aligned} \tag{31}$$

3.2. *Structural estimation of the baseline model*

Now consider estimation of the underlying structural parameters of the model presented in Section 2. The key point to note is that the above OLS estimates $\{\widehat{\phi}_1, \widehat{\Phi}_1, \widehat{\Omega}_1, \widehat{\gamma}_1, \hat{\Gamma}_1, \widehat{\sigma}_{e1}, \ldots, \widehat{\phi}_4, \widehat{\Phi}_4, \widehat{\Omega}_4, \widehat{\gamma}_4, \hat{\Gamma}_4, \widehat{\sigma}_{e4}\}$ are sufficient statistics for inference about these parameters – anything that the full sample of data is able to tell us about the model parameters can be summarized by the values of these OLS estimates. The idea behind the minimum-chi-square estimation proposed by Hamilton and Wu (2012b) is to choose structural parameters that would imply reduced-form coefficients as close as possible to the unrestricted estimates, an approach that turns out to be asymptotically equivalent to full MLE.

Note that the model developed here specifies observed prices in terms of an unobserved factor vector x_t. There is an arbitrary normalization in any such system, in that if we were to multiply x_t by a nonsingular matrix and add a constant, the result would be observationally equivalent in terms of the implied likelihood for observed y_t.[7] Since we have treated the factors x_t as directly inferable from the values of y_{1t}, we normalize the factors so that they could be interpreted as the level and slope as of the date of expiry of the near-term contract:

$$x_t = H_1 y_t \quad \text{for } j_t = 4. \tag{32}$$

Recalling (22), this would be the case if

$$H_1 y_t = H_1 \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix} + H_1 \begin{bmatrix} \beta'_0 \\ \beta'_4 \\ \beta'_8 \end{bmatrix} x_t \quad \text{for } j_t = 4. \tag{33}$$

Substituting (32) into (33), our chosen normalization thus calls for

$$x_t = H_1 \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix} + H_1 \begin{bmatrix} \beta'_0 \\ \beta'_4 \\ \beta'_8 \end{bmatrix} x_t \quad \text{for } j_t = 4$$

[7] For further discussion of identification and normalization, see Hamilton and Wu (2012b).

$$H_1 \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} = I_2 \tag{34}$$

$$H_1 \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix} = 0. \tag{35}$$

Since

$$\begin{bmatrix} f_{0t} \\ f_{4t} \\ f_{8t} \end{bmatrix} = \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix} + \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} x_t,$$

our normalization could alternatively be described as $x_t = H_1(f_{0t}, f_{4t}, f_{8t})'$ for all t. Following Joslin et al. (2011) and Hamilton and Wu (2013), this can be implemented by defining ξ_1 and ξ_2 to be the eigenvalues of $\rho^Q = \rho - \Lambda$. Given this normalization and values for ξ_1, ξ_2, Σ, and α_0, we can then determine the values for ρ^Q, c^Q, and $\{\beta_n, \alpha_n\}_{n=0}^N$; details are provided in Appendix B. These along with ρ, c, and σ_e then provide everything we need to evaluate the likelihood function or to calculate what the predicted values for any of the unrestricted reduced-form coefficients ought to be.

Let θ denote the vector of unknown structural parameters, that is, the 16 elements of $\{\xi_1, \xi_2, \Sigma, \alpha_0, \rho, c, \sigma_{e1}, \sigma_{e2}, \sigma_{e3}, \sigma_{e4}\}$ for Σ lower triangular. Collect elements of the unrestricted OLS estimates in a vector $\widehat{\pi}$:

$$\widehat{\pi} = (\widehat{\pi}_\Phi', \widehat{\pi}_\Omega', \widehat{\pi}_\Gamma', \widehat{\pi}_\sigma)'$$

$$\widehat{\pi}_\Phi = \left(\left[\text{vec}\left(\left[\widehat{\phi}_1 \quad \widehat{\Phi}_1\right]'\right)\right]', \ldots, \left[\text{vec}\left(\left[\widehat{\phi}_4 \quad \widehat{\Phi}_4\right]'\right)\right]'\right)'$$

$$\widehat{\pi}_\Omega = \left(\left[\text{vech}\left(\widehat{\Omega}_1\right)\right]', \ldots, \left[\text{vech}\left(\widehat{\Omega}_4\right)\right]'\right)'$$

$$\widehat{\pi}_\Gamma = \left(\left[\text{vec}\left(\left[\widehat{\gamma}_1 \quad \hat{\Gamma}_1\right]'\right)\right]', \ldots, \left[\text{vec}\left(\left[\widehat{\gamma}_4 \quad \hat{\Gamma}_4\right]'\right)\right]'\right)'$$

$$\widehat{\pi}_\sigma = (\widehat{\sigma}_{e1}, \widehat{\sigma}_{e2}, \widehat{\sigma}_{e3}, \widehat{\sigma}_{e4})'.$$

Let $g(\theta)$ denote the corresponding predicted values for those coefficients from the model; specific values for the elements of $g(\theta)$ are summarized in Appendix C. The minimum-chi-square (MCS) estimate of θ is the value that minimizes

$$T[\widehat{\pi} - g(\theta)]'\widehat{R}[\widehat{\pi} - g(\theta)] \tag{36}$$

for $\widehat{R}$ the information matrix associated with the OLS estimates $\widehat{\pi}$, which is also detailed in Appendix C. The MCS estimator has the same asymptotic distribution as the maximum likelihood estimator, but has a number of computational and interpretive advantages over MLE discussed in Hamilton and Wu (2012b). Because $\widehat{R}$ is block-diagonal with respect to π_σ, the MCS estimates of these parameters are given immediately by the OLS estimates (31). Hamilton and Wu (2012b) show that asymptotic standard errors can be estimated using

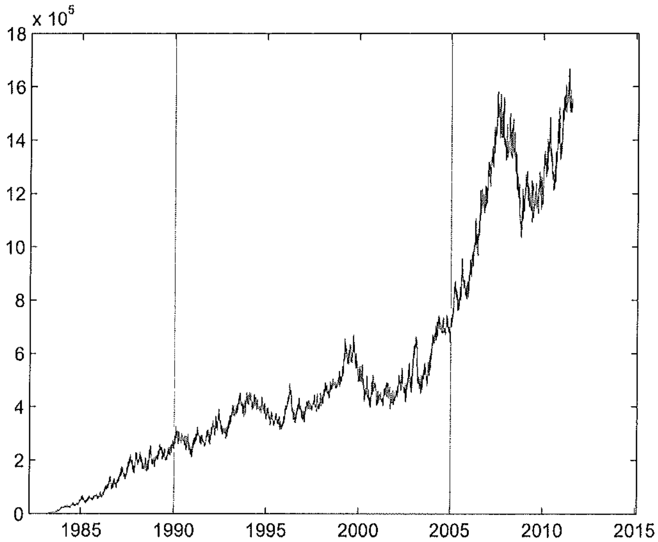

Fig. 2. Open Interest. Total number of outstanding crude oil futures contracts of all maturities, daily observations, March 3, 1983 to July 17, 2011. Vertical lines drawn at January 1, 1990 and January 1, 2005.

$$E\left(\hat{\theta} - \theta_0\right)\left(\hat{\theta} - \theta_0\right)' \simeq T^{-1}\left(\hat{G}'\hat{R}\hat{G}\right)^{-1}$$
$$\hat{G} = \left.\frac{\partial g(\theta)}{\partial \theta'}\right|_{\theta=\hat{\theta}},$$

which are identical to the usual asymptotic errors that would be obtained by taking second derivatives of the log likelihood function (28) with respect to θ.

4. Empirical results for the baseline model

Crude oil futures contracts were first traded on the New York Mercantile Exchange (NYMEX) in 1983. In the first few years, volume was much lighter than the more recent data, and we choose to begin our empirical analysis in January, 1990. Fig. 2 plots the total open interest on all NYMEX light sweet crude contracts. Volume expanded very quickly after 2004, in part in response to the increased purchases of futures contracts as a vehicle for financial diversification. Some researchers have suggested that participation in the markets by this new class of traders resulted in significant changes in the dynamic behavior of crude oil futures prices.[8] A likelihood ratio test (e.g., Hamilton (1994, p. 296)) of the null hypothesis that the coefficients of the unrestricted reduced form are constant over time against the alternative that all 52 parameters changed in January 2005 produces a $\chi^2(52)$ statistic of 181.96, which calls for dramatic rejection of the null hypothesis (p value of 2.2×10^{-16}). Since one of our interests in this paper is to document how futures price dynamics have changed over time, we conduct our analysis on two subsamples, the first covering January 1990 through December 2004, and the second January 2005 through June 2011.

The left panel of Table 2 reports minimum-chi-square estimates of the 16 elements of θ based on the first subsample. The eigenvalues of ρ, the matrix summarizing the objective P-measure persistence of

[8] See for example Alquist and Kilian (2007), Singleton (2011), Tang and Xiong (2012), Mou (2010), and Buyuksahin and Robe (2011).

Table 2
Pre-2005 parameter estimates for baseline model.

Estimated parameters			Implied parameters		
c	0.0102	0.0016	λ	0.0104	−0.0061*
	(0.0209)	(0.0029)		(0.0209)	(0.0029)
ρ	0.9974*	0.1839*	Λ	−0.0023	−0.0591
	(0.0067)	(0.0918)		(0.0067)	(0.0918)
	−0.0006	0.9301*		0.0020*	−0.0050
	(0.0009)	(0.0134)		(0.0009)	(0.0136)
ξ	0.9876*	0.9473*	$\lambda + \Lambda\bar{x}$	0.0037*	3.53e-005
	(0.0010)	(0.0044)		(0.0018)	(2.63e-004)
Σ	0.0449*	0			
	(0.0012)				
	−0.0038*	0.0047*			
	(0.0002)	(0.0001)			
α_0	0.0357*				
	(0.0007)				
π_σ	0.0099*	0.0081*			
	(0.0005)	(0.0004)			
	0.0105*	0.0201*			
	(0.0006)	(0.0011)			

Left panel: MCS estimates of elements of θ for data from January 1990 through December 2004 (asymptotic standard errors in parentheses). Right panel: assorted magnitudes of interest implied by value of $\hat{\theta}$ (asymptotic standard errors in parentheses). * denotes statistically significant at the 5% level.

factors, are 0.9956 and 0.9319, implying that both level and slope are highly persistent, with similar estimates for their Q-measure counterparts (ξ_1 and ξ_2).

The differences between the P- and Q-measures, or implied characterization of λ_t, are reported in the right panel of Table 2.[9] The individual elements of λ and Λ are generally small and statistically insignificant. The last two entries of Table 2 report the elements of $\lambda + \Lambda\bar{x}$, where $\bar{x}$ is the average value for the level and spread over the sample. The positive value of 0.0037 for the first element of this vector suggests that an investor who was always long in the two contracts would on average have come out ahead over this period, an estimate that is just statistically significant at the 5% level.[10]

Table 3 reports parameter estimates for the later subsample, in which there appear to be significant differences in risk pricing from the earlier data. Most noteworthy is the large negative value for Λ_{12}. This signifies that when the spread (the second element of x_t) gets sufficiently high, a long position in the 1- and 2-month contracts would on average lose money. We also see from the last entry of Table 3 that the first element of $\lambda + \Lambda\bar{x}$ is smaller in the second subsample than in the first, and is no longer statistically significant. The average reward for taking long positions in the second subsample is not as evident as in the first subsample.

Fig. 3 plots our estimated values for $\lambda_t = \lambda + \Lambda x_t$ for each week t in our sample, along with 95% confidence intervals. The price of level risk (top panel) was uniformly positive up until 2006, but has often been negative since 2008. By contrast, slope risk (bottom panel) was typically not priced before 2004, whereas going long the 2-month contract and short the 1-month has frequently been associated with positive expected returns since then.[11]

[9] Standard errors for λ and Λ were obtained by reparameterizing the MCS estimation in terms of λ and Λ instead of c and ρ. The values of λ and Λ can be obtained analytically from $\lambda = c - c^Q$ and $\Lambda = \rho - \rho^Q$ with c^Q given by equation (47) and ρ^Q given by equation (42).

[10] Note from (12) and (34) that the expected return on a portfolio with equal weights on the second and third contracts is given by $(1/2)(\beta_4' + \beta_8')\lambda_t = [1 \quad 0]\lambda_t$ whose average value is the first element of $\lambda + \Lambda\bar{x}$.

[11] From (12) and (34), the expected return on a portfolio that is long the third contract and short the second is given by $(\beta_8' - \beta_4')\lambda_t = [0 \quad 1]\lambda_t$ whose average value is the second element of $\lambda + \Lambda\bar{x}$.

Table 3
Post-2005 parameter estimates for baseline model.

Estimated parameters			Implied parameters		
c	0.1802*	0.0164*	λ	0.1813*	0.0179*
	(0.0574)	(0.0070)		(0.0574)	(0.0070)
ρ	0.9600*	−0.3487	Λ	−0.0400*	−0.5892*
	(0.0131)	(0.2018)		(0.0131)	(0.2018)
	−0.0035*	0.8629*		−0.0039*	−0.0311
	(0.0016)	(0.0241)		(0.0016)	(0.0243)
ξ	1.0010*	0.8931*	$\lambda + \Lambda\bar{x}$	0.0028	0.0009*
	(0.0001)	(0.0047)		(0.0026)	(0.0003)
Σ	0.0439*	0			
	(0.0019)				
	−0.0021*	0.0049*			
	(0.0003)	(0.0002)			
α_0	−0.0086*				
	(0.0009)				
π_σ	0.0059*	0.0086*			
	(0.0005)	(0.0007)			
	0.0087*	0.0223*			
	(0.0007)	(0.0018)			

Left panel: MCS estimates of elements of θ for data from January 2005 through June 2011 (asymptotic standard errors in parentheses). Right panel: assorted magnitudes of interest implied by value of $\hat{\theta}$ (asymptotic standard errors in parentheses). * denotes statistically significant at the 5% level.

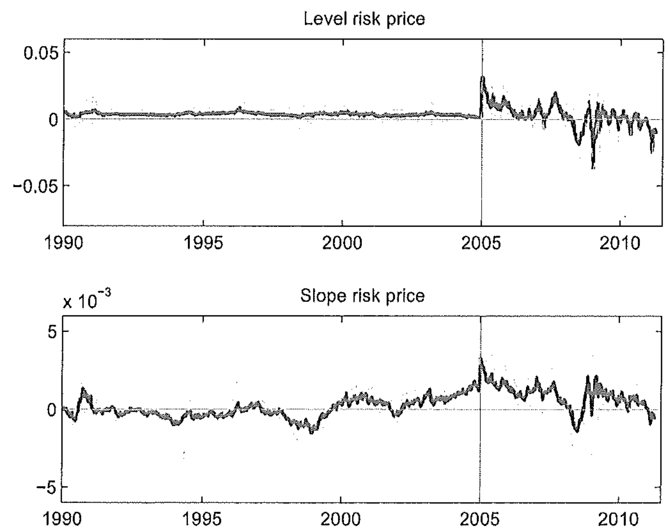

Fig. 3. Prices of factor risk. Top panel: first element of $\lambda + \Lambda x_t$ as estimated from baseline model, with sample split in 2005. Bottom panel: second element. Dashed lines indicate 95% confidence intervals.

Following Cochrane and Piazzesi (2009) and Bauer et al. (2012), another way to summarize the implications of these results is to calculate how different the log price of a given contract would be if there was no compensation for risk. To get this number, we calculate $\tilde{f}_{nt} = \tilde{\alpha}_n + \tilde{\beta}_n x_t$, where $\tilde{\beta}_n$ and $\tilde{\alpha}_n$ denote the values that would be obtained from the recursions (10) and (11) if Λ and λ were both set to zero. The value for the difference $\tilde{f}_{nt} - f_{nt}$ for $n = 8$ weeks is plotted in Fig. 4. In the absence of risk

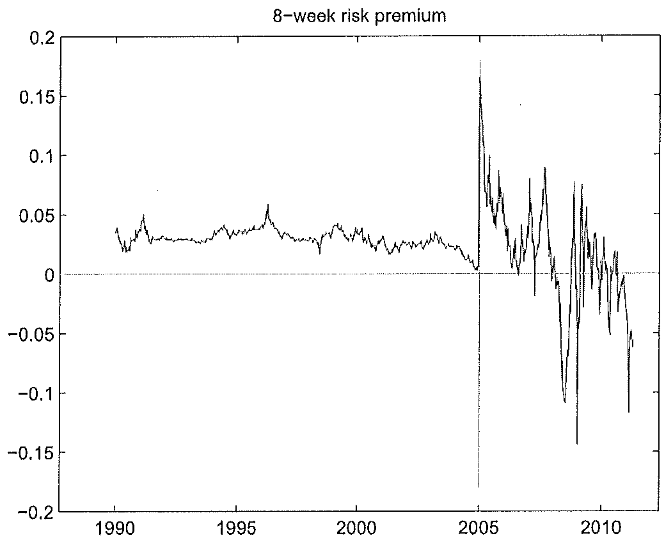

Fig. 4. Risk premium on 8-week futures contract. Plot of $\tilde{f}_{8t} - f_{8t}$ as estimated from the baseline model with sample split in 2005.

effects, an 8-week contract price would have been a few percent higher on average over the 1990–2004 subsample.[12] Since 2005, risk aversion has made a more volatile contribution, though the average effect is significantly smaller.

In terms of the framework proposed in Section 2 for interpreting these results, the positive average value for the first element of λ_t in the first subsample suggests that arbitrageurs were on average long in crude oil futures contracts over this period, accepting the positive expected earnings from their positions as compensation for providing insurance to sellers, who were presumably commercial producers who wanted to hedge their price risks by selling futures contracts. From that perspective, an increase in index fund buying could have been one explanation for why a long position in futures contracts no longer has a statistically significant positive return. In effect, index-fund buyers are serving as counterparty for commercial hedgers, and are willing to do so without the risk compensation that the position earned on average in the first subsample. The emerging positive return to a spreading position (positive average second element of λ_t in the second subsample) would be consistent with the view that arbitrageurs are buying and holding 2-month futures from oil producers, but then selling these positions and going short 1-month futures as they sell to index-fund investors.

As noted by Hamilton and Wu (2012b), another benefit of estimation by minimum chi square is that the optimized value for the objective function provides an immediate test of the overall framework. Under the null hypothesis that the model is correctly specified, the minimum value achieved for (36) has an asymptotic χ^2 distribution with degrees of freedom given by the number of overidentifying restrictions. The first column of Table 4 reports the value of this statistic for each of the two subsamples. The model is overwhelmingly rejected in either subsample.

Because the weighting matrix $\hat{R}$ in (49) is block-diagonal, it is easy to decompose these test statistics into components coming from the respective elements of π, as is done in subsequent columns of Table

[12] The average size of this risk premium is 2.9% for the first sample. This compares with an average realized 2-month ex post return over this period of 2.0% for the long position on a 3-month contract (that is, the average log value of the first contract minus the average log value of the third contract two months earlier), and an average difference between the first contract and third contract at the same date of 1.2% (that is, the futures curve sloped down on average with a slope of −1.2%). The last number is similar to the value reported by Alquist and Kilian (2010), who noted that the 3-month futures price was 1.1% below the spot price on average over 1987–2007. The difference between the average ex post return to the long position and the negative of the average slope results from the significantly higher price of oil at the end of the sample than at the beginning.

Table 4
χ^2 specification test and breakdown by individual components.

	χ^2	$d.f.$	p-value	π_Φ	π_Ω	π_Γ
Before 2005	86.57	36	4.73e-6	25.61	43.98	16.99
Since 2005	151.87	36	3.33e-16	120.01	17.22	14.64

χ^2: minimum value achieved for MCSE objective function. $d.f.$: degrees of freedom. p-value: probability of observing $\chi^2(d.f.)$ value this large. Last 3 columns: contribution to χ^2 of individual parameter blocks.

4. In the first subsample, about half of the value of the test statistic comes from the π_Ω block – the differences in the variability of the level and slope across different weeks of the month is more than can be explained by the fact that the maturities of observed contracts are changing week to week. The biggest problem in the second subsample come from the π_Φ block – unrestricted forecasts of the level and slope vary more week-to-week than is readily explained by differences in the maturities of the contracts.

It is also possible to look one parameter at a time at where the structural model misses. For each of the unrestricted reduced-form parameters π there is a corresponding prediction from the model $g(\theta)$ for what that value is supposed to be if the model is correct. Figs. 5 and 6 plot the unrestricted OLS estimates of the various elements of π along with their 95% confidence intervals for the first subsample. The thick red lines indicate the value the coefficient is predicted to have according to the structural parameters reported in Table 2. The biggest problems come from the fact that the model underpredicts the difficulty of forecasting the spread in weeks 1 and 3 (the lower left panel of Fig. 6). Figs. 7 and 8 provide the analogous plots for the second subsample. Here the biggest problems come from the fact that the equations one would want to use to forecast the spread in weeks 3 and 4 are quite different

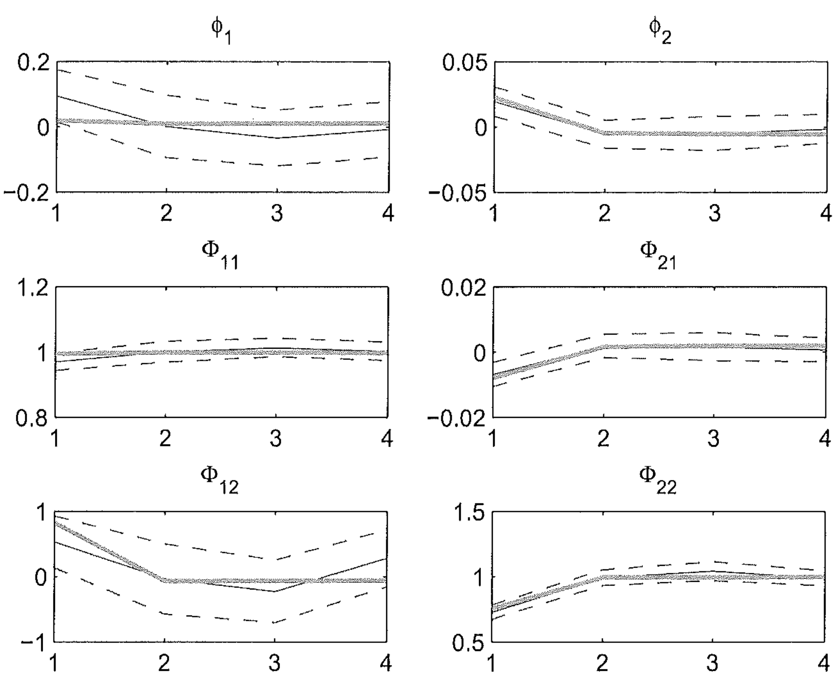

Fig. 5. $\pi_\Phi$before 2005. Light blue line: Unrestricted OLS estimates of coefficients for regression in which y_{1t} is the dependent variable, plotted as a function of week of the month. Dashed blue lines: 95% confidence intervals for unrestricted OLS estimates. Bold red line: predicted values for coefficients derived from baseline model. All estimates based on data January 1990 to December 2004. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

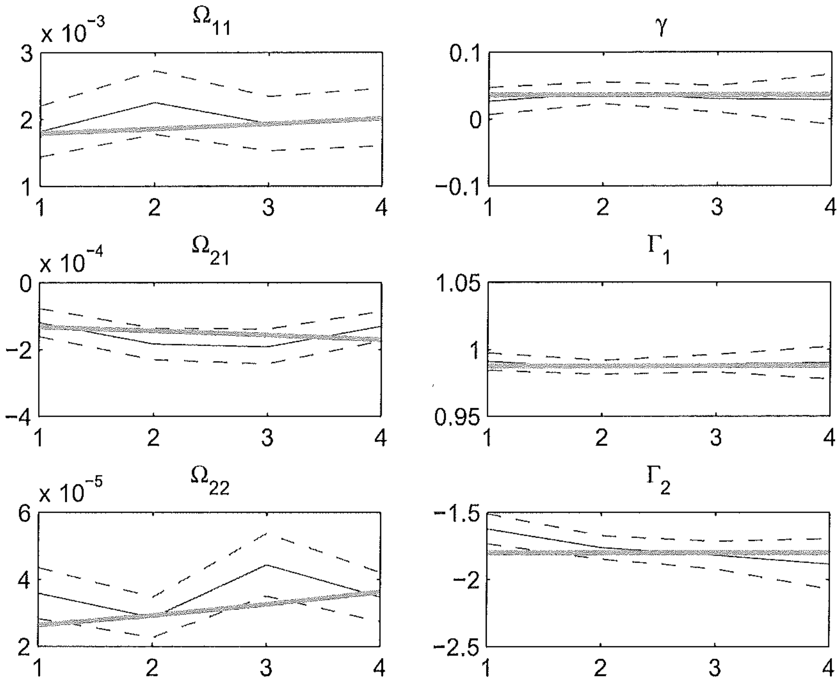

Fig. 6. π_Ω and π_Γ before 2005. First column: estimated elements of variance-covariance matrix for regression in which y_{1t} is the dependent variable, plotted as a function of week of the month. Second column: Estimated values of coefficients for regression in which y_{2t} is the dependent variable, plotted as a function of week of the month. In each panel, light blue lines are unrestricted OLS estimates, dashed blue lines are 95% confidence intervals for unrestricted OLS estimates, and bold red lines are predicted values for coefficients derived from baseline model. All estimates based on data January 1990 to December 2004. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

from those found for weeks 1 and 2 (see the right-hand column of Fig. 7). These considerations suggest looking at models that allow for more general seasonal variation than our baseline specification, which we explore in the next section.

5. Less restrictive seasonal models

5.1. Structural estimation

Here we consider a system in which the dynamic process followed by the factors is itself dependent on which week of the month we are looking at:

$$x_{t+1} = c_{j_t} + \rho_{j_t} x_t + \Sigma_{j_t} u_{t+1}.$$

If we hypothesize that the risk-pricing parameters also vary with the season,

$$\lambda_t = \lambda_{j_t} + \Lambda_{j_t} x_t$$

then the no-arbitrage conditions (10) and (11) generalize to

$$\beta_n' = \beta_{n-1}' \rho_{j(n)}^Q \tag{37}$$

$$\alpha_n = \alpha_{n-1} + \beta_{n-1}' c_{j(n)}^Q + (1/2)\beta_{n-1}' \Sigma_{j(n)} \Sigma_{j(n)}' \beta_{n-1}.$$

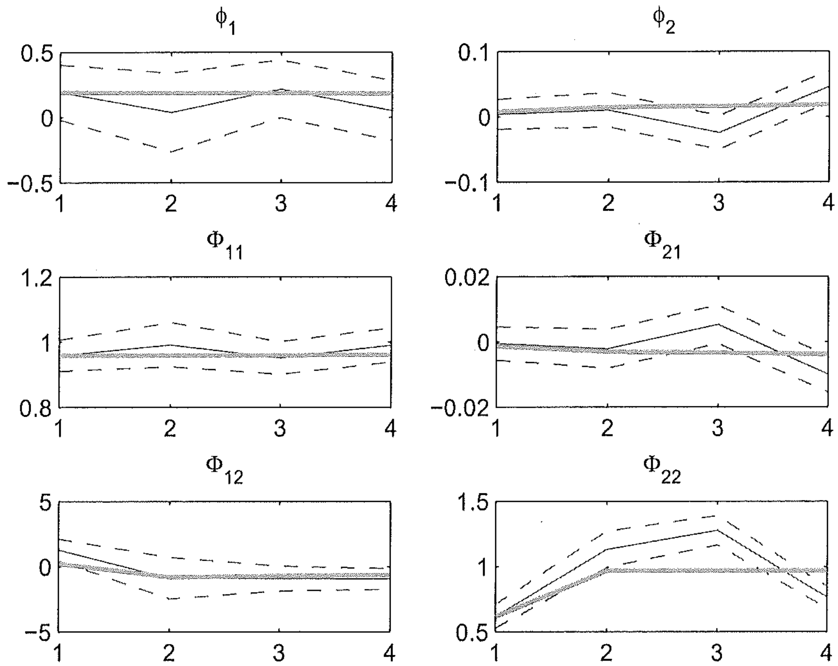

Fig. 7. π_{Φ} since 2005. Light blue line: Unrestricted OLS estimates of coefficients for regression in which y_{1t} is the dependent variable, plotted as a function of week of the month. Dashed blue lines: 95% confidence intervals for unrestricted OLS estimates. Bold red line: predicted values for coefficients derived from baseline model. All estimates based on data January 2005 to June 2011. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

where which observation week j is associated with a given maturity n can be read off of Table 1 and where we have defined $\rho_j^Q = \rho_j - \Lambda_j$ and $c_j^Q = c_j - \lambda_j$. Unfortunately, if all the parameters were allowed to vary with the week j in this way, the model would be unidentified. The reason is that even if one hypothesizes different values of ρ_j^Q for different j, a generalization of the algebra in (48) still implies that Γ_j should be the same for all j:

$$\Gamma_j = H_2 \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} \quad \text{for } j = 1,2,3,4.$$

Since $\beta_4' = \beta_0'\rho_3^Q\rho_2^Q\rho_1^Q\rho_4^Q$ and $\beta_8' = \beta_4'\rho_3^Q\rho_2^Q\rho_1^Q\rho_4^Q$, the only information available from the regressions in which y_{2t} is the dependent variable (26) is about the product $\rho_3^Q\rho_2^Q\rho_1^Q\rho_4^Q$, which does not allow identification of the individual terms. In the next subsection we report estimates for a system in which although c_j, ρ_j, λ_j, and Λ_j all vary with j, the differences $c^Q = c_j - \lambda_j$ and $\rho^Q = \rho_j - \Lambda_j$ do not. For this system, the flexibility of the c_j and ρ_j parameters allows us to fit the unrestricted OLS values for ϕ_j and Φ_j perfectly. Details of the normalization and estimation for this less restrictive specification are reported in Appendices B and C.

5.2. Empirical results for the less restrictive seasonal model

Empirical estimates for the parameters of the above system for each of the two subsamples are reported in Tables 5 and 6. In the first subsample, the main differences are that the specification allows

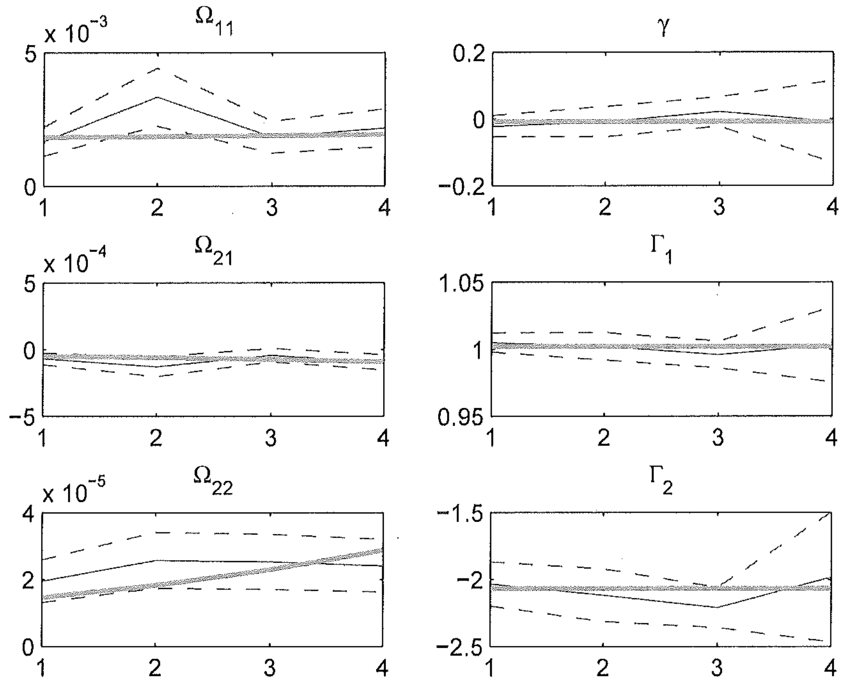

Fig. 8. π_Ω and π_Γ since 2005. First column: estimated elements of variance-covariance matrix for regression in which y_{1t} is the dependent variable, plotted as a function of week of the month. Second column: Estimated values of coefficients for regression in which y_{2t} is the dependent variable, plotted as a function of week of the month. In each panel, light blue lines are unrestricted OLS estimates, dashed blue lines are 95% confidence intervals for unrestricted OLS estimates, and bold red lines are predicted values for coefficients derived from baseline model. All estimates based on data January 2005 to June 2011. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the spread to become harder to forecast as the near contract approaches expiry (that is, the (2,2) element of Σ_j increases in j) and the level and slope at the end of the month are less related to their values at expiry than is typical of the relation between y_{1t} and $y_{1,t-1}$ at other times (that is, diagonal elements of ρ_j are smaller for $j = 4$). Although implied values for λ and Λ are estimated with much less precision, the overall conclusion that individual elements are small and statistically insignificant applies across individual weeks as well.

For the second subsample (Table 6), the dependence of Λ_{12j} on week j is very dramatic, with an average value of -0.78 for $j = 1, 2$, or 3 but an estimated value of $+0.46$ for $j = 4$. A high spread signals lower returns to the long position during weeks 1–3, but this effect completely disappears, and may even take on the opposite sign, during expiry week 4. This may be related to the strong weekly pattern to index-fund strategies. For example, to replicate the crude oil holdings of the Goldman Sachs Commodity Index, an index fund would be selling the $k = 0$ contract and buying the $k = 1$ contract during week $j = 3$. It is interesting that we also find strong weekly patterns in the pricing of risk in data since 2005, though trying to interpret those changes in detail is beyond the scope of this paper.

Although our more general specification can fit the unrestricted OLS estimates $\widehat{\phi}_j$ and $\widehat{\Phi}_j$ perfectly, it still imposes testable overidentifying restrictions on other parameters, essentially using the 3 parameters in $\{\alpha_0, \xi_1, \xi_2\}$ to fit the 12 values for $\{\widehat{\gamma}_j, \widehat{\Gamma}_j\}_{j=1}^{4}$. The resulting $\chi^2(9)$ MCS test statistic for the first subsample is 13.86, which with a p-value of 0.13 is consistent with the null hypothesis that the

Table 5
Pre-2005 parameter estimates for seasonal model.

Estimated parameters					Implied parameters				
c^{base}	0.0102	0.0016			λ^{base}	0.0104	−0.0061*		
c_1	0.0042	0.0005			λ_1	0.0044	−0.0068		
	(0.0467)	(0.0056)				(0.0467)	(0.0055)		
c_2	−0.0359	0.0023			λ_2	−0.0356	−0.0050		
	(0.0422)	(0.0066)				(0.0422)	(0.0065)		
c_3	−0.0054	0.0053			λ_3	−0.0051	−0.0020		
	(0.0414)	(0.0055)				(0.0414)	(0.0054)		
c_4	0.0871*	−0.0003			λ_4	0.0874*	−0.0077		
	(0.0422)	(0.0073)				(0.0421)	(0.0079)		
ρ^{base}	0.9974*	0.1839*	−0.0006	0.9301*	Λ^{base}	−0.0023	−0.0591	0.0020*	−0.0050
ρ_1	0.9992*	0.2166	−0.0003	0.9369*	Λ_1	−0.0005	−0.0271	0.0022	−0.0061
	(0.0152)	(0.2386)	(0.0018)	(0.0285)		(0.0152)	(0.2386)	(0.0018)	(0.0281)
ρ_2	1.0137*	0.0371	−0.0010	0.9825*	Λ_2	0.0139	−0.2067	0.0015	0.0395
	(0.0138)	(0.2190)	(0.0022)	(0.0343)		(0.0138)	(0.2190)	(0.0021)	(0.0339)
ρ_3	1.0011*	0.5076*	−0.0018	0.9310*	Λ_3	0.0014	0.2639	0.0007	−0.0120
	(0.0135)	(0.2061)	(0.0018)	(0.0274)		(0.0135)	(0.2061)	(0.0018)	(0.0270)
ρ_4	0.9726*	−0.0718	0.0000	0.8690*	Λ_4	−0.0272*	−0.3155	0.0025	−0.0740
	(0.0138)	(0.2063)	(0.0024)	(0.0350)		(0.0138)	(0.2063)	(0.0026)	(0.0375)
Σ^{base}	0.0449*	−0.0038*	0.0047*						
Σ_1	0.0491*	−0.0042*	0.0039*						
	(0.0026)	(0.0004)	(0.0002)						
Σ_2	0.0447*	−0.0046*	0.0052*						
	(0.0024)	(0.0005)	(0.0003)						
Σ_3	0.0447*	−0.0029*	0.0051*						
	(0.0024)	(0.0004)	(0.0003)						
Σ_4	0.0446*	−0.0035*	0.0059*						
	(0.0024)	(0.0005)	(0.0003)						
ξ^{base}	0.9876*	0.9473*							
ξ	0.9854*	0.9574*							
	(0.0029)	(0.0070)							
α_0^{base}	0.0357*								
α_0	0.0331*								
	(0.0051)								
π_σ	0.0099*	0.0081*	0.0105*	0.0201*					

Left panel: MCS estimates for elements of θ (asymptotic standard errors in parentheses) for the unrestricted seasonal model, with estimates from baseline model also reported for comparison, with all estimates based on data from January 1990 through December 2004. Elements of matrices reported as first row, then second row. Right panel: values implied by the reported estimates of θ. * denotes statistically significant at the 5% level.

model has adequately captured all the week-to-week variations in parameters. The second subsample ($\chi^2(9) = 13.25$, $p = 0.15$) also passes this specification test.

6. Conclusions

In this paper, we studied the interaction between hedging demands from commercial producers or financial investors and risk aversion on the part of the arbitrageurs who are persuaded to be the hedgers' counterparties. We demonstrated that this interaction can produce an affine factor structure for the log prices of futures contracts in which expected returns depend on the arbitrageurs' net exposure to nondiversifiable risk. We developed new algorithms for estimation and diagnostic tools for testing this class of models appropriate for an unbalanced data set in which the duration of observed contracts changes with each observation.

Prior to 2005, we found that someone who consistently took the long side of nearby oil futures contracts received positive compensation on average, with relatively modest variation of this risk premium over time, consistent with the interpretation that the primary source of this premium was

Table 6
Post-2005 parameter estimates for seasonal model.

Estimated parameters					Implied parameters				
c^{base}	0.1802*	0.0164*			λ^{base}	0.1813*	0.0179*		
c_1	0.0361	0.0106			λ_1	0.0374	0.0119		
	(0.1528)	(0.0169)				(0.1528)	(0.0167)		
c_2	0.2252*	−0.0299*			λ_2	0.2264*	−0.0286		
	(0.1094)	(0.0147)				(0.1094)	(0.0144)		
c_3	0.0508	0.0443*			λ_3	0.0520	0.0456*		
	(0.1127)	(0.0120)				(0.1127)	(0.0119)		
c_4	0.1852	0.0102			λ_4	0.1865	0.0115		
	(0.1081)	(0.0175)				(0.1081)	(0.0184)		
ρ^{base}	0.9600*	−0.3487	−0.0035*	0.8629*	Λ^{base}	−0.0400*	−0.5892*	−0.0039*	−0.0311
ρ_1	0.9918*	−0.4445	−0.0021	1.0026*	Λ_1	−0.0082	−0.6847	−0.0025	0.1141
	(0.0346)	(0.5708)	(0.0038)	(0.0627)		(0.0346)	(0.5708)	(0.0038)	(0.0622)
ρ_2	0.9488*	−0.5659	0.0066*	1.1369*	Λ_2	−0.0512*	−0.8061*	0.0062	0.2484*
	(0.0247)	(0.3810)	(0.0033)	(0.0507)		(0.0247)	(0.3811)	(0.0033)	(0.0500)
ρ_3	0.9895*	−0.6038	−0.0097*	0.6759*	Λ_3	−0.0105	−0.8440*	−0.0100*	−0.2126*
	(0.0257)	(0.3454)	(0.0027)	(0.0367)		(0.0257)	(0.3454)	(0.0027)	(0.0364)
ρ_4	0.9592*	0.7020	−0.0021	0.8796*	Λ_4	−0.0409	0.4618	−0.0025	−0.0089
	(0.0245)	(0.4244)	(0.0040)	(0.0661)		(0.0245)	(0.4244)	(0.0042)	(0.0684)
Σ^{base}	0.0439*	−0.0021*	0.0049*						
Σ_1	0.0579*	−0.0031*	0.0055*						
	(0.0047)	(0.0007)	(0.0005)						
Σ_2	0.0423*	−0.0013*	0.0054*						
	(0.0035)	(0.0006)	(0.0004)						
Σ_3	0.0456*	−0.0020*	0.0044*						
	(0.0037)	(0.0005)	(0.0004)						
Σ_4	0.0417*	−0.0029*	0.0055*						
	(0.0034)	(0.0007)	(0.0005)						
ξ^{base}	1.0010*	0.8931*							
ξ	1.0008*	0.8877*							
	(0.0010)	(0.0059)							
α_0^{base}	−0.0086*								
α_0	−0.0077								
	(0.0106)								
π_σ	0.0059*	0.0086*	0.0087*	0.0223*					

Left panel: MCS estimates for elements of θ (asymptotic standard errors in parentheses) for the unrestricted seasonal model, with estimates from baseline model also reported for comparison, with all estimates based on data from January 2005 through June 2011. Elements of matrices reported as first row, then second row. Right panel: values implied by the reported estimates of θ. * denotes statistically significant at the 5% level.

hedging by commercial producers. However, we uncovered significant changes in the pricing of risk after the volume of trading in these contracts increased significantly in 2005. The expected compensation from a long position is lower on average in the recent data, often significantly negative when the futures curve slopes upward. We suggest that increased participation by financial investors in oil futures markets may have been a factor in changing the nature of risk premia in crude oil futures contracts.

Appendix A. Approximations to portfolio mean and variance

We first note that if $\log X \sim N(\mu, \sigma^2)$ then $E(X) = \exp(\mu + \sigma^2/2)$. Taking a first-order Taylor approximation around $\mu = \sigma^2 = 0$ we have $E(X) \approx 1 + \mu + \sigma^2/2$. Thus in particular since

$$\log(F_{n-1,t+1}/F_{nt}) \sim N\left(\mu_{n-1,t}, \sigma^2_{n-1}\right)$$

$$\mu_{n-1,t} = \alpha_{n-1} + \beta'_{n-1}(c + \rho x_t) - \alpha_n - \beta'_n x_t \qquad (38)$$

$$\sigma_{n-1}^2 = \beta_{n-1}'\Sigma\Sigma'\beta_{n-1} \tag{39}$$

we have the approximations

$$E_t[(F_{n-1,t+1} - F_{nt})/F_{nt}] \approx \mu_{n-1,t} + \sigma_{n-1}^2/2$$

$$E_t[\exp(r_{j,t+1})] \approx 1 + \xi_j + \psi_j'(c + \rho x_t) + \psi_j'\Sigma\Sigma'\psi_j/2.$$

Substituting these into (1) gives (5).

Likewise, if

$$\begin{bmatrix} \log X_i \\ \log X_j \end{bmatrix} \sim N\left(\begin{bmatrix} \mu_i \\ \mu_j \end{bmatrix}, \begin{bmatrix} \sigma_{ii} & \sigma_{ij} \\ \sigma_{ji} & \sigma_{jj} \end{bmatrix}\right),$$

then

$$\mathrm{Cov}(X_i, X_j) = \exp\left[\left(\mu_i + \mu_j\right) + (\sigma_{ii} + \sigma_{jj})/2\right]\left[\exp(\sigma_{ij}) - 1\right].$$

A first-order Taylor expansion around $\mu_i = \mu_j = \sigma_{ii} = \sigma_{jj} = \sigma_{ij} = 0$ gives

$$\mathrm{Cov}(X_i, X_j) \approx \sigma_{ij}.$$

To use this result we define

$$\underset{(L\times 1)}{y_{t+1}} = \left(r_{1,t+1}, \ldots, r_{J,t+1}, f_{0,t+1} - f_{1t}, f_{1,t+1} - f_{2t}, \ldots, f_{N-1,t+1} - f_{Nt}\right)'$$

for $L = J + N$. Notice that conditional on information at date t, $y_{t+1} \sim N(\mu_t, H'\Sigma\Sigma'H)$ for[13]

$$H = \begin{bmatrix} \psi_1 & \cdots & \psi_J & \beta_0 & \cdots & \beta_{N-1} \end{bmatrix}.$$

Notice further that (1) can be written $W_{t+1} = k_t + \sum_{\ell=1}^{L} h_{\ell t} \exp(y_{\ell,t+1})$ for $h_{\ell t} = q_{\ell t}$ for $\ell = 1, \ldots, J$ and $h_{\ell t} = z_{\ell-J,t}$ for $\ell = J+1, \ldots, L$. Thus for $h_t = (h_{1t}, \ldots, h_{Lt})'$,

$$\begin{aligned} \mathrm{Var}_t(W_{t+1}) &\approx h_t' H'\Sigma\Sigma' H h_t \\ &= \left(\sum_{j=1}^{J} q_{jt}\psi_j' + \sum_{n=1}^{N} z_{nt}\beta_{n-1}'\right)\Sigma\Sigma'\left(\sum_{j=1}^{J} q_{jt}\psi_j + \sum_{\ell=1}^{N} z_{\ell t}\beta_{\ell-1}\right). \end{aligned}$$

Appendix B. Normalization

Baseline model. Let $\rho^Q = \rho - \Lambda$, and notice from (10) that

$$\beta_n = \left(\rho^{Q\prime}\right)^n \beta_0. \tag{40}$$

We will parameterize this in terms of $\xi = (\xi_1, \xi_2)'$ where ξ_i denotes an eigenvalue of $\rho^{Q\prime}$. We could calculate the following matrix as a function of those eigenvalues:

$$\underset{(2\times 3)}{K(\xi)} = \begin{bmatrix} \xi_1^0 & \xi_1^4 & \xi_1^8 \\ \xi_2^0 & \xi_2^4 & \xi_2^8 \end{bmatrix}. \tag{41}$$

The claim is that if we specify

[13] Here $\mu_t = (\mu_{1t}, \ldots, \mu_{Lt})'$ for $\mu_{\ell t} = \xi_\ell + \psi_\ell'(c + \rho x_t)$ for $\ell = 1, \ldots, J$ and $\mu_{\ell t} = \alpha_{\ell-J-1} + \beta_{\ell-J-1}'(c + \rho x_t) - \alpha_{\ell-J} - \beta_{\ell-J}' x_t$ for $\ell = J+1, \ldots, L$.

$$\underset{(2\times2)}{\rho^{Q\prime}} = \begin{bmatrix} \underset{(2\times3)}{K(\xi)} & \underset{(3\times2)}{H_1'} \end{bmatrix}^{-1} \begin{bmatrix} \xi_1 & 0 \\ 0 & \xi_2 \end{bmatrix} \begin{bmatrix} \underset{(2\times3)}{K(\xi)} & \underset{(3\times2)}{H_1'} \end{bmatrix} \tag{42}$$

$$\underset{(2\times1)}{\beta_0} = \begin{bmatrix} \underset{(2\times3)}{K(\xi)} & \underset{(3\times2)}{H_1'} \end{bmatrix}^{-1} \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \tag{43}$$

then (34), the desired condition for β_n, would be satisfied. To prove this, observe from (40) that

$$\begin{aligned} \beta_n &= [K(\xi)H_1']^{-1} \begin{bmatrix} \xi_1^n & 0 \\ 0 & \xi_2^n \end{bmatrix} [K(\xi)H_1'] [K(\xi)H_1']^{-1} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \\ &= [K(\xi)H_1']^{-1} \begin{bmatrix} \xi_1^n \\ \xi_2^n \end{bmatrix} \end{aligned} \tag{44}$$

so that

$$H_1 \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} = H_1 \begin{bmatrix} \xi_1^0 & \xi_2^0 \\ \xi_1^4 & \xi_2^4 \\ \xi_1^8 & \xi_2^8 \end{bmatrix} [H_1 K(\xi)']^{-1}. \tag{45}$$

Substituting (41) into (45) produces (34), as claimed. Thus if we know ξ, we can use (41) and (44) to calculate the value of β_n for any n as well as ρ^Q from (42).

To achieve the separate condition (35) on α_n, notice from (11) that

$$\alpha_n = \alpha_0 + (\beta_{n-1}' + \beta_{n-2}' + \cdots + \beta_0')c^Q + (1/2)(\beta_{n-1}'\Sigma\Sigma'\beta_{n-1} + \beta_{n-2}'\Sigma\Sigma'\beta_{n-2} + \cdots + \beta_0'\Sigma\Sigma'\beta_0). \tag{46}$$

Define

$$\begin{aligned} \underset{(1\times2)}{\zeta_n(\xi)} &= \beta_{n-1}' + \beta_{n-2}' + \cdots + \beta_0' \\ \underset{(1\times1)}{\psi_n(\xi,\alpha_0,\Sigma)} &= \alpha_0 + (1/2)(\beta_{n-1}'\Sigma\Sigma'\beta_{n-1} + \beta_{n-2}'\Sigma\Sigma'\beta_{n-2} + \cdots + \beta_0'\Sigma\Sigma'\beta_0) \end{aligned}$$

so that (46) can be written

$$\alpha_n = \zeta_n(\xi)c^Q + \psi_n(\xi,\alpha_0,\Sigma)$$

where for $n = 0$ we have $\zeta_0(\xi) = 0$ and $\psi_0(\xi,\alpha_0,\Sigma) = \alpha_0$. We claim that if we choose

$$c^Q = -\left(H_1 \begin{bmatrix} \zeta_0(\xi) \\ \zeta_4(\xi) \\ \zeta_8(\xi) \end{bmatrix}\right)^{-1} \left(H_1 \begin{bmatrix} \psi_0(\xi,\alpha_0,\Sigma) \\ \psi_4(\xi,\alpha_0,\Sigma) \\ \psi_8(\xi,\alpha_0,\Sigma) \end{bmatrix}\right), \tag{47}$$

then (35) would be satisfied. This is demonstrated as follows:

$$\begin{aligned} H_1 \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix} &= H_1 \begin{bmatrix} \zeta_0(\xi)c^Q + \psi_0(\xi,\alpha_0,\Sigma) \\ \zeta_4(\xi)c^Q + \psi_4(\xi,\alpha_0,\Sigma) \\ \zeta_8(\xi)c^Q + \psi_8(\xi,\alpha_0,\Sigma) \end{bmatrix} \\ &= -H_1 \begin{bmatrix} \psi_0(\xi,\alpha_0,\Sigma) \\ \psi_4(\xi,\alpha_0,\Sigma) \\ \psi_8(\xi,\alpha_0,\Sigma) \end{bmatrix} + H_1 \begin{bmatrix} \psi_0(\xi,\alpha_0,\Sigma) \\ \psi_4(\xi,\alpha_0,\Sigma) \\ \psi_8(\xi,\alpha_0,\Sigma) \end{bmatrix} \\ &= 0. \end{aligned}$$

Thus if we know ξ, α_0, and Σ, we can use (46) and (47) to calculate the value of α_n for any n. Also ξ, α_0, and Σ allow calculation of c^Q from (47).

Seasonal model. Just as in the baseline model, we let (ξ_1, ξ_2) denote the ordered eigenvalues of ρ^Q and write

$$\beta_n' = \begin{bmatrix} \xi_1^n & \xi_2^n \end{bmatrix} \left[H_1 K(\xi)'\right]^{-1}$$

which achieves the normalization (34). Likewise for α_n we again use (47) where now

$$\zeta_n(\xi) = \beta_{n-1}' + \beta_{n-2}' + \cdots + \beta_0'$$
$$\psi_n(\xi, \alpha_0, \Sigma) = \alpha_0 + (1/2)\left(\beta_{n-1}'\Sigma_{j(n)}\Sigma_{j(n)}'\beta_{n-1} + \beta_{n-2}'\Sigma_{j(n-1)}\Sigma_{j(n-1)}'\beta_{n-2} + \cdots + \beta_0'\Sigma_{j(1)}\Sigma_{j(1)}'\beta_0\right).$$

Appendix C. Mapping from structural to reduced-form parameters

Baseline model. Expressions involving B_{1j} in (23) can be simplified by noting from (40) and (34) that

$$B_{1j} = H_1 \begin{bmatrix} \beta_{4-j}' \\ \beta_{8-j}' \\ \beta_{12-j}' \end{bmatrix} = H_1 \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} \left(\rho^Q\right)^{4-j} = \left(\rho^Q\right)^{4-j}.$$

Thus for example (27) simplifies to

$$\begin{aligned} \Gamma_j &= H_2 \begin{bmatrix} \beta_{4-j}' \\ \beta_{8-j}' \\ \beta_{12-j}' \end{bmatrix} \left(H_1 \begin{bmatrix} \beta_{4-j}' \\ \beta_{8-j}' \\ \beta_{12-j}' \end{bmatrix} \right)^{-1} \\ &= H_2 \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} \left(\rho^Q\right)^{4-j} \left[\left(\rho^Q\right)^{4-j}\right]^{-1} \\ &= H_2 \begin{bmatrix} \beta_0' \\ \beta_4' \\ \beta_8' \end{bmatrix} \quad \text{for } j = 1,2,3,4. \end{aligned} \tag{48}$$

The population magnitudes corresponding to the other reduced-form OLS coefficients are as follows:

$$\Omega_j = \left(\rho^Q\right)^{4-j} \Sigma\Sigma' \left(\rho^{Q\prime}\right)^{4-j} \quad \text{for } j = 1,2,3,4$$

$$\Phi_1 = \left(\rho^Q\right)^3 \rho$$

$$\Phi_j = \left(\rho^Q\right)^{4-j} \rho \left[\left(\rho^Q\right)^{-1}\right]^{4-j+1} \quad \text{for } j = 2,3,4$$

$$\phi_1 = H_1 \begin{bmatrix} \alpha_3 \\ \alpha_7 \\ \alpha_{11} \end{bmatrix} + H_1 \begin{bmatrix} \beta_3' \\ \beta_7' \\ \beta_{11}' \end{bmatrix} c - \Phi_1 H_1 \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix}$$

$$\phi_j = H_1 \begin{bmatrix} \alpha_{4-j} \\ \alpha_{8-j} \\ \alpha_{12-j} \end{bmatrix} + H_1 \begin{bmatrix} \beta'_{4-j} \\ \beta'_{8-j} \\ \beta'_{12-j} \end{bmatrix} c - \Phi_j H_1 \begin{bmatrix} \alpha_{5-j} \\ \alpha_{9-j} \\ \alpha_{13-j} \end{bmatrix} \quad \text{for } j = 2,3,4$$

$$\gamma_j = H_2 \begin{bmatrix} \alpha_{4-j} \\ \alpha_{8-j} \\ \alpha_{12-j} \end{bmatrix} - \Gamma_j H_1 \begin{bmatrix} \alpha_{4-j} \\ \alpha_{8-j} \\ \alpha_{12-j} \end{bmatrix} \quad \text{for } j = 1,2,3,4.$$

Given the scalars $\{\xi_1, \xi_2\}$ (corresponding to the eigenvalues of $\rho^Q = \rho - \Lambda$), we can calculate β_n from (44) and (41). These β_n give us predicted values for $\{\Gamma_j\}_{j=1}^4$, and the β_n along with Σ give predicted values for $\{\Omega_j\}_{j=1}^4$. Note $\{\xi_1, \xi_2\}$ also gives us ρ^Q, and this plus ρ gives predicted values for $\{\Phi_j\}_{j=1}^4$. From β_n, Σ, and α_0 we can calculate c^Q from (47) and α_n from (46). Using these along with c we then obtain the predicted values for $\{\phi_j\}_{j=1}^4$ and $\{\gamma_j\}_{j=1}^4$.

The information matrix for the OLS estimates $\hat{\pi} = (\hat{\pi}'_\Phi, \hat{\pi}'_\Omega, \hat{\pi}'_\Gamma \hat{\pi}_\sigma)'$ is given by

$$\hat{R}_\pi = \begin{bmatrix} \hat{R}_\Phi & 0 & 0 & 0 \\ 0 & \hat{R}_\Omega & 0 & 0 \\ 0 & 0 & \hat{R}_\Gamma & 0 \\ 0 & 0 & 0 & \hat{R}_\sigma \end{bmatrix} \tag{49}$$

$$\hat{R}_\Phi = \begin{bmatrix} \hat{R}_{\Phi 1} & 0 & 0 & 0 \\ 0 & \hat{R}_{\Phi 2} & 0 & 0 \\ 0 & 0 & \hat{R}_{\Phi 3} & 0 \\ 0 & 0 & 0 & \hat{R}_{\Phi 4} \end{bmatrix}$$

$$\hat{R}_{\Phi j} = \hat{\Omega}_j^{-1} \otimes T^{-1} \sum_{\tau=1}^{T} x^{\dagger}_{1j,\tau} x^{\dagger\prime}_{1j,\tau}$$

$$\hat{R}_\Omega = \begin{bmatrix} \hat{R}_{\Omega 1} & 0 & 0 & 0 \\ 0 & \hat{R}_{\Omega 2} & 0 & 0 \\ 0 & 0 & \hat{R}_{\Omega 3} & 0 \\ 0 & 0 & 0 & \hat{R}_{\Omega 4} \end{bmatrix}$$

$$\hat{R}_{\Omega j} = (1/2) D'_2 \left(\hat{\Omega}_j^{-1} \otimes \hat{\Omega}_j^{-1} \right) D_2$$

$$D_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$\hat{R}_\Gamma = \begin{bmatrix} \hat{R}_{\Gamma 1} & 0 & 0 & 0 \\ 0 & \hat{R}_{\Gamma 2} & 0 & 0 \\ 0 & 0 & \hat{R}_{\Gamma 3} & 0 \\ 0 & 0 & 0 & \hat{R}_{\Gamma 4} \end{bmatrix}$$

$$\widehat{R}_{\Gamma j} = \widehat{\sigma}_{ej}^{-2}\left(\mathcal{T}^{-1}\sum_{\tau=1}^{\mathcal{T}} x_{2j,\tau}^{\dagger} x_{2j,\tau}^{\dagger\prime}\right)$$

$$\widehat{R}_{\sigma} = \begin{bmatrix} (1/2)\widehat{\sigma}_{e1}^{-4} & 0 & 0 & 0 \\ 0 & (1/2)\widehat{\sigma}_{e2}^{-4} & 0 & 0 \\ 0 & 0 & (1/2)\widehat{\sigma}_{e3}^{-4} & 0 \\ 0 & 0 & 0 & (1/2)\widehat{\sigma}_{e4}^{-4} \end{bmatrix}$$

where D_2 is the duplication matrix satisfying D_2 vech(Ω) = vec(Ω) for 2×2 symmetric matrix Ω. Note that for all models considered the MCS estimate of σ_{ej} is always equal to the unconstrained MLE $\widehat{\sigma}_{ej}$ and so contributes 0 to the weighted objective function.

Unrestricted seasonal model. For this model the specifications for γ_j and Γ_j are the same as in the baseline model, while the expressions for the other parameters become

$$\Omega_1 = \left(\rho^{Q}\right)^3 \Sigma_4 \Sigma_4' \left(\rho^{Q\prime}\right)^3$$

$$\Omega_j = \left(\rho^{Q}\right)^{4-j} \Sigma_{j-1} \Sigma_{j-1}' \left(\rho^{Q\prime}\right)^{4-j} \quad \text{for } j = 2,3,4$$

$$\Phi_1 = \left(\rho^{Q}\right)^3 \rho_4$$

$$\Phi_j = \left(\rho^{Q}\right)^{4-j} \rho_{j-1} \left[\left(\rho^{Q}\right)^{-1}\right]^{4-j+1} \quad \text{for } j = 2,3,4 \tag{50}$$

$$\phi_1 = H_1 \begin{bmatrix} \alpha_3 \\ \alpha_7 \\ \alpha_{11} \end{bmatrix} + H_1 \begin{bmatrix} \beta_3' \\ \beta_7' \\ \beta_{11}' \end{bmatrix} c_4 - \Phi_1 H_1 \begin{bmatrix} \alpha_0 \\ \alpha_4 \\ \alpha_8 \end{bmatrix}$$

$$\phi_j = H_1 \begin{bmatrix} \alpha_{4-j} \\ \alpha_{8-j} \\ \alpha_{12-j} \end{bmatrix} + H_1 \begin{bmatrix} \beta_{4-j}' \\ \beta_{8-j}' \\ \beta_{12-j}' \end{bmatrix} c_{j-1} - \Phi_j H_1 \begin{bmatrix} \alpha_{5-j} \\ \alpha_{9-j} \\ \alpha_{13-j} \end{bmatrix} \quad \text{for } j = 2,3,4.$$

Minimum-chi-square estimation in this case is achieved by first choosing $\{\xi, \alpha_0, \Sigma_1, \Sigma_2, \Sigma_3, \Sigma_4\}$ so as to minimize the distance from the OLS estimates $\{\widehat{\gamma}_j, \widehat{\Gamma}_j, \widehat{\Omega}_j\}_{j=1}^{4}$. From ξ we can then calculate ρ^{Q}, with which we can obtain ρ_j analytically from (50) in order to fit these OLS coefficients perfectly. The values c_{j-1} are likewise obtained analytically from $\widehat{\phi}_j$.

References

Acharya, Viral V., Lochstoer, Lars A., Ramadorai, Tarun, 2013. Limits to arbitrage and hedging: evidence from commodity markets (forthcoming). Journal of Financial Economics.

Alquist, Ron, Gervais, Olivier, 2011. The Role of Financial Speculation in Driving the Price of Crude Oil. Working paper. Bank of Canada.
Alquist, Ron, Kilian, Lutz, 2007. What Do We Learn from the Price of Crude Oil Futures? CEPR. Working paper No. 6548.
Alquist, Ron, Kilian, Lutz, 2010. What do we learn from the price of crude oil futures? Journal of Applied Econometrics 25, 539–573.
Ang, Andrew, Piazzesi, Monika, 2003. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. Journal of Monetary Economics 50, 745–787.
Bauer, Michael D., Rudebusch, Glenn D., Wu, Jing Cynthia, 2012. Correcting estimation bias in dynamic term structure models. Journal of Business & Economic Statistics 30, 454–467.
Bessembinder, Hendrik, 1992. Systemic risk, hedging pressure, and risk premiums in futures markets. Review of Financial Studies 5, 637–667.
Buyuksahin, Bahattin, Robe, Michel A., 2011. Does 'Paper Oil' Matter?. Working paper American University.
Carter, Colin A., Rausser, Gordon C., Schmitz, Andrew, 1983. Efficient asset portfolios and the theory of normal backwardation. Journal of Political Economy 91, 319–331.
Casassus, Jaime, Collin-Dufresne, Pierre, 2006. Stochastic convenience yield implied from commodity futures and interest rates. Journal of Finance 60, 2283–2331.
Chang, Eric C., 1985. Returns to speculators and the theory of normal backwardation. Journal of Finance 40, 193–208.
Chen, Ren-Raw, Scott, Louis, 1993. Maximum likelihood estimation for a multifactor equilibrium model of the term structure of interest rates. The Journal of Fixed Income 3, 14–31.
Cochrane, John H., Piazzesi, Monika, 2009. Decomposing the yield curve. In: AFA 2010 Atlanta Meetings Paper.
Cortazar, Gonzalo, Naranjo, Lorenzo, 2006. An N-factor Gaussian model of oil futures prices. Journal of Futures Markets 26, 243–268.
Creal, Drew D., Wu, Jing Cynthia, 2013. Estimation of Non-Gaussian Affine Term Structure Models. Chicago Booth Research Paper No. 13-72.
Dai, Qiang, Singleton, Kenneth J., 2002. Expectation puzzles, time-varying risk premia, and affine models of the term structure. Journal of Financial Economics 63, 415–441.
De Roon, Frans, Nijman, Theo, Veld, Chris, 2000. Hedging pressure effects in futures markets. Journal of Finance 55, 1437–1456.
Duffee, Gregory R., 2002. Term premia and interest rate forecasts in affine models. The Journal of Finance 57, 405–443.
Duffie, Darrell, 1992. Dynamic Asset Pricing Theory. Princeton University Press, Princeton, NJ.
Duffie, Darrell, Kan, Rui, 1996. A yield-factor model of interest rates. Mathematical Finance 6, 379–406.
Etula, Erkko, 2013. Broker-dealer risk appetite and commodity returns. Journal of Financial Econometrics 11, 486–521.
Fattouh, Bassam, Kilian, Lutz, Mahadeva, Lavan, 2013. The role of speculation in oil markets: what have we learned so far? Energy Journal 34, 7–33.
Hamilton, James D., 1994. Time Series Analysis. Princeton University Press, Princeton, New Jersey.
Hamilton, James D., 2009. Causes and Consequences of the Oil Shock of 2007–08. Brookings Papers on Economic Activity Spring, pp. 215–259.
Hamilton, James D., Wu, Jing Cynthia, 2012a. The effectiveness of alternative monetary policy tools in a zero lower bound environment. Journal of Money, Credit & Banking 44 (s1), 3–46.
Hamilton, James D., Wu, Jing Cynthia, 2012b. Identification and estimation of Gaussian affine term structure models. Journal of Econometrics 168, 315–331.
Hamilton, James D., Wu, Jing Cynthia, 2012c. Effects of Index-fund Investing on Commodity Futures Prices. Working paper. UCSD.
Hamilton, James D., Wu, Jing Cynthia, 2013. Testable implications of affine term structure models (forthcoming). Journal of Econometrics.
Irwin, Scott H., Sanders, Dwight R., 2012. Testing the masters hypothesis in commodity futures markets. Energy Economics 34, 256–269.
Joslin, Scott, Singleton, Kenneth J., Zhu, Haoxiang, 2011. A new perspective on Gaussian dynamic term structure models. Review of Financial Studies 24, 926–970.
Keynes, John M., 1930. A Treatise on Money, vol. 2. Macmillan, London.
Kilian, Lutz, Murphy, Daniel P., 2013. The role of inventories and speculative trading in the global market for crude oil (forthcoming). Journal of Applied Econometrics.
Mou, Yiqun, 2010. Limits to Arbitrage and Commodity Index Investment: Front-running the Goldman Roll. Working paper. Columbia University.
Schwartz, Eduardo S., 1997. The stochastic behavior of commodity prices: implications for valuation and hedging. Journal of Finance 52, 923–973.
Schwartz, Eduardo S., Smith, James E., 2000. Short-term variations and long-term dynamics in commodity prices. Management Science 46, 893–911.
Singleton, Kenneth J., 2011. Investor Flows and the 2008 Boom/Bust in Oil Prices. Working paper. Stanford University.
Tang, Ke, Xiong, Wei, 2012. Index investment and the financialization of commodities. Financial Analysts Journal 68, 54–74.
Vasicek, Oldrich, 1977. An equilibrium characterization of the term structure. Journal of Financial Economics 5, 177–188.
Vayanos, Dimitri, Vila, Jean-Luc, 2009. A Preferred-habitat Model of the Term Structure of Interest Rates. Working paper. London School of Economics.