

Speculation in Commodity Futures Markets, Inventories and the Price of Crude Oil *

Sung Je Byun †

March 1, 2016

Abstract

Refiners have an incentive to hold inventories even if they anticipate falling crude oil prices. This paper develops a model of the convenience yield arising from holding inventories. Using data on inventories and futures contract prices, I show that the convenience yield is inversely related to inventories, and positively related to oil prices. In addition to exhibiting seasonal and procyclical behaviors, I show that the historical convenience yield averages about 19% of the oil price from March 1989 to November 2014. Although some have argued that a breakdown of the relationship between crude oil inventories and prices following increased participation by financial investors after 2003 was evidence of an effect of speculation, I find that the proposed price-inventory relationship is stable over time. In light of this evidence, I conclude that the contribution of financial investors' activities to the crude oil market is weak.

*JEL classification: G13, G15, G17, G23

*Keywords: Speculation, Convenience yield, Forecasting oil prices, Stable oil price-inventory relationship

*I am very grateful to James Hamilton, Allan Timmermann, Valerie Ramey, Alexis Toda, Thomas Baranga, Christiane Baumeister, Soojin Jo and anonymous referees for their helpful comments and suggestions on earlier drafts of this paper.

†Financial Industry Studies Department, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX, USA. The views expressed are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

1 Introduction

The recent volatility of crude oil prices has renewed interest in the behavior of crude oil inventories. This paper examines the role of inventories in refiners' gasoline production and develops a structural model of the relation between crude oil prices and inventories.

In a competitive commodity market, a producer makes the optimal storage decision by equating his expected benefit with the relevant cost of holding inventories. The positive benefit could motivate a producer to hold inventories even if the producer anticipates falling prices in the future. Given this motivation, the importance of inventories for storable commodities has been widely recognized in the theory of storage literature. Among early researchers, Brennan (1958) defines this benefit as the convenience yield, which is the flow of services that accrues to an owner of the physical commodity but not to an owner of a contract for future delivery of the commodity.

In this paper, I propose an equilibrium storage model of the global crude oil market to study the role of crude oil inventories in refiners' gasoline production and to explain fluctuations of crude oil prices in terms of risk-averse refiners' benefits to hold inventories. First, I model determinants of crude oil inventories, building upon earlier research by Eckstein and Eichenbaum (1985) and Considine (1997). Specifically, I model the role of crude oil inventories directly as enhancing refiners' gasoline production by treating inventories as an essential factor of production following Kydland and Prescott (1982, 1988) and Ramey (1989). I derive an inverse relationship between convenience yield and the level of crude oil inventories, which has never been done in the literature. Instead, the common approach as in Pindyck (1994, 2001), Alquist and Kilian (2010) and Knittel and Pindyck (2015) is to assume such an inverse relationship when modeling determinants of crude oil prices and to infer convenience yield from the term spread observed in prices of crude oil futures contracts at two different maturities. Second, I provide empirical evidence that the convenience yield is inversely related to crude oil inventories, yet positively related to crude oil prices. Moreover, it exhibits a procyclical behavior along the business cycle with the convenience yield being increased during the expansionary period. In general, these findings are consistent with the earlier conjecture in the theory of storage. Third, I find from model estimates that the convenience yield accounts to about 19% of oil price on average from March 1989 to November 2014, holding other factors of production constant. Regarding the

observed seasonality in oil inventories, I find a higher convenience yield for spring and summer than for other seasons. These complement earlier researchers' empirical findings. Fourth, I identify a structural change in crude oil market fundamentals since 2004 from a rising permanent component of crude oil prices. Such a structural break is also evidenced in changing model parameter estimates for before and after June 2004, which is consistent with Tang and Xiong (2012), Büyükkşahin and Robe (2014) and Hamilton and Wu (2014). Lastly, I show that the proposed model provides more accurate 3-, 6-, 9- and 12-month-ahead crude oil spot price forecasts than the random walk model, indicating the stable price-inventory relation despite the potential structural break.

I illustrate one use of the proposed model by re-examining the possible contribution of financial investors on the stable price-inventory relation in the crude oil market. Given rapidly increasing financial investors' activities since 2004 (Irwin and Sanders 2011), a public debate has centered on the potential contribution of financial investors' participation in the commodity futures markets (Masters 2008, Kennedy 2012). One piece of evidence that may lead to such perception was provided by Merino and Ortiz (2005), who noted that the traditional model of the relation between crude oil prices and inventories, such as that developed by Ye et al. (2002), broke down after financial investor participation in crude oil markets increased. Merino and Ortiz extended this model to include a contribution of the long positions of non-commercial traders, and they documented that such a specification could better explain the relation between prices and inventories over 2001-2004.

To investigate the potential contribution of financial investors' activities on the crude oil market, I consider previous forecasting models provided by Ye et al. (2002) and Merino and Ortiz (2005) for two reasons. First, these models are grounded in the fundamental relationship between the two variables. Second, such simple models provide a good framework for examining popular political perceptions on the crude oil market. In this paper, I reproduce the good in-sample fit of the Ye specification, its breakdown after 2001 and the improvement provided by non-commercial positions over 2001-2004. However, I also find that even the latter relation falls apart after 2004. I show that although the traditional model seems to do well over the period 1992-2001, it is in fact misspecified even over that period and simply appears to do well because there is relatively little change in the randomly driven component of crude oil prices over that period. Despite the structural change after 2004, I show that the model of the

inventory-price relation proposed here is stable over time and furthermore, can account for both the apparent success of the Ye et al. (2002) specification on earlier data as well as its breakdown on subsequent data. This finding is consistent with Kilian and Murphy (2014) who document the stable forecasting relationship between the oil market variables using a structural vector autoregressive (VAR) model. Although the approach taken here is indirect by focusing on the refineries' storage decision, this paper adds further support to the recent academic literature documenting empirical evidence for no contribution of financial investors (or limited at most) from various angles (Fattouh et al. 2013, Kilian and Murphy 2014, Kilian and Lee 2014, Juvenal and Petrella 2015, Knittel and Pindyck 2015, Hamilton and Wu 2015).

The remainder of the paper proceeds as follows. Section 2 introduces a storage model where the convenience yield arises from the refiners' expected benefits to hold crude oil inventories. After introducing ways to deal with empirical issues, I provide a convenience yield measure at the end of section 2. In section 3, I begin by introducing data and an empirical procedure for estimating the proposed storage model. Next, I document estimation results together with explanations for maximum likelihood estimates of structural model parameters as well as empirical properties of the proposed convenience yield. After providing evidence on the structural break since 2004, I show the stable price-inventory relation proposed in this paper is robust to the structural break in the crude oil market. In section 4, I illustrate one use of the proposed model by investigating earlier forecasting models and show the strong forecasting relationship of inventories and prices, indicating a weak contribution of financial investors' activities in the crude oil market. Lastly, I conclude with some remarks.

2 Theoretical Model

In this section, I propose an equilibrium storage model of the global crude oil market to study the role of crude oil inventories in refiners' gasoline production. I study the optimal storage decision of refiners who have an incentive to hold inventories despite anticipated falling crude oil prices. Note that the first-order conditions for refiners have to hold regardless of the actions of financial investors in futures markets, because a refiner can always respond to arbitrage opportunities from storing oil and selling a futures contract. This is separate from the behavior of financial investors, who roll over maturing futures contracts forward rather than taking a

physical delivery of crude oil (Smith 2009, Hamilton and Wu 2015). However, financial investors' activities could be reflected through risk premiums that the refinery pays for its hedging activities in the crude oil market (Hamilton and Wu 2014). Although financialization in commodity markets affects risk-premiums, my model differs from Basak and Pavlova (2015). There is no role except consumption smoothing for storable commodities in their model, which neglects the role of storable commodities in the theory of storage literature. Furthermore, the objective in this paper is to build a structural relationship between relevant oil market variables and to explain fluctuations of crude oil prices in terms of risk-averse refiners' benefits to hold inventory, i.e., convenience yield. From this perspective, the adopted approach differs from Knittel and Pindyck (2015), in which the authors seek to isolate the speculative component of changes in prices and in the convenience yield. Lastly, the proposed model provides a fully specified structural relationship among oil market variables following the evidence established in the literature. This differs from the approach in the recent structural VAR literature such as Kilian and Murphy (2014), Kilian and Lee (2014) and Juvenal and Petrella (2015), whose objectives are to disentangle various sources of price changes using residuals obtained from an unrestricted VAR model for oil market variables.

2.1 Cost Function

Consider a representative producer who purchases a quantity q_t^P of crude oil at price S_t per barrel, of which q_t^U is used up in current production of a consumption good such as gasoline and the remainder goes to increase inventories i_t :¹

$$i_t = i_{t-1} + q_t^P - q_t^U, \quad (1)$$

where $q_t^P - q_t^U$ corresponds to net additional crude oil inventories during the time period t . If the quantity of oil consumed by the representative producer (q_t^U) is smaller than the quantity purchased (q_t^P), inventories accumulate and vice versa.

¹The proposed model abstracts a refiner's decision for gasoline production using crude oil from its complicated production decisions for multiple intermediary goods. It is because crude oil itself has accounted for more than 40% of the wholesale gasoline price. Excluding sales taxes, costs for the crude oil resource dominate the refinery's remaining operation costs such as manufacturing and marketing. Given its importance in gasoline production, I focus on crude oil inventories and explain the benefit from holding raw material inventories. From this perspective, the approach in this paper is different from most earlier literature that studies roles of finished good inventories in smoothing a production schedule and reducing marketing costs.

The cost function summarizes both the production and storage technology for the representative producer. Given the current crude oil spot price (S_t), the firm's costs come from two components: costs for purchasing resources and those for storing inventories over one period. The former is the product of crude oil spot price and amounts purchased (q_t^P), and the latter is assumed to be proportional to the current crude oil spot price.² For reflecting the idea of limited storage facilities in the short-run, storage costs are further assumed to take the form of a convex quadratic function. Given the current crude oil spot price (S_t) and previous level of inventories (i_{t-1}), the representative producer's cost function becomes

$$c(q_t^U, i_t; i_{t-1}, S_t) = S_t \{q_t^U + (i_t - i_{t-1})\} + S_t \left(c_0 + \frac{c_1}{2} i_t^2\right),$$

where I plug in the accounting identity of inventories following (1) after rearranging it for a quantity purchased (q_t^P) and $c_0, c_1 > 0$.

2.2 Production Function

The importance of inventories has been widely recognized in the theory of storage developed by Kaldor (1939), Working (1949) and Brennan (1958). In particular, Brennan (1958) defines the convenience yield as the flow of services that accrues to an owner of the physical commodity but not to an owner of a contract for future delivery of the commodity. These services can arise from reducing the probability of a stock-out of inventories, from inventories' role in production smoothing, and from future production cost saving. For example, Pindyck (1994) describes the role of inventories in production is "to reduce costs of adjusting production and to reduce marketing costs by facilitating schedule and avoiding stockouts". With his proposed measure for the intangible convenience yield using available futures contract prices, he finds evidence for such roles for copper, heating oil, and lumber. Further empirical tests for the convenience yield in commodities include Fama and French (1987, 1988), Deaton and Laroque (1992), Ng and Pirrong (1994), Routledge et al. (2000), Bryant et al. (2006), Gorton et al. (2013), Acharya et al. (2013), and Pirrong (2011).

The role of inventories as a factor of production has been widely adopted in earlier litera-

²The cost for storage is defined in a broad sense: it includes the cost for insurance and transportation beside that for using storage facilities. In general, these costs are determined as being proportional to the value that is stored.

ture. For example, Kydland and Prescott (1982, 1988) and Ramey (1989) treat inventories as essential factors of production for studies of aggregate fluctuations in the U.S. economy. Focusing on the petroleum refining industry, Considine (1997) examines determinants of inventory investment under joint production by treating inventories as quasi-fixed factors of production. In this paper, I model crude oil inventories as directly increasing the refinery's production capabilities. To motivate this role of inventories further, consider the refinery's decision problem for a periodic production schedule under resource constraints. Compared to firms in other industries, the refinery's production heavily depends on the raw material, in this case, crude oil. It generally takes more than a few weeks to receive additional crude oil delivery after placing an order in the spot market or making a transaction in the derivative market. For its current gasoline production, the refinery can use, at most, crude oil resources that are either carried over from the previous period or are delivered currently from the past transaction. In response to uncertainties in the aggregate gasoline demand and petroleum prices, the available inventories in the refinery's storage determine the attainable level of production and can improve production efficiency without needing to make costly production adjustments. Hence, the refinery has a motivation for holding inventories despite anticipated falling crude oil prices in the future. Given the realized technology process (z_t) , I propose a production function of the form,

$$f(q_t^U; i_{t-1}, z_t) = (e^{z_t} q_t^U)^\alpha \left\{ 1 - e^{-\theta_t \cdot i_{t-1}} \right\},$$

where $\alpha \in (0, 1)$ is the output share of the crude oil resource and $\theta_t > 0$ is the utilization parameter governing the production function's dependence on the previous level of inventory. The subscript t on utilization parameter (θ_t) is used to allow the possibility for taking different values depending on the season of the year - for example, taking a smaller value when the refinery tries to produce more gasoline for the summer driving season. The utilization function denoted by the term in curly brackets is bounded between 0 and 1, which is a concave function of the previous level of inventory. This captures the nonlinearity of a refinery's production arising from inventory stock-outs, approaching 0 at an increasing rate for i_{t-1} being close to 0 and approaching 1 at a reducing rate for sufficiently large i_{t-1} . The resource-augmenting technology process (z_t) follows the random walk process, that is, $z_{t+1} = z_t + \epsilon_{1,t+1}$ with $\epsilon_{1,t+1} \sim N(0, \sigma_1^2)$.

2.3 Theoretical Model

With the cost function and the production function as introduced earlier, the representative producer faces a dynamic programming problem. At the beginning of each period t , the representative producer faces the crude oil spot price of S_t , the realized technology process z_t and the exogenously determined real interest rate r_t . Given previous inventory (i_{t-1}), he makes decisions on the resource demand (q_t^U) and inventories (i_t) jointly in order to maximize his lifetime profits. Suppose he is a price taker in the crude oil market and he is risk averse with discount factor (Λ_t),³ which governs his risk aversion. Suppressing the superscript in the resource demand (q_t), the representative producer's objective is thus to choose $\{q_t, i_t\}_{t=1}^{\infty}$ so as to maximize

$$\Pi = \max_{\{q_t, i_t\}} E_0 \left[\sum_{t=1}^{\infty} \prod_{\tau=1}^t \Lambda_{\tau} e^{-r_{\tau}} \{f(q_t; i_{t-1}, z_t) - c(q_t, i_t; i_{t-1}, S_t)\} \right],$$

where i_{-1} is given. Note, I pose this as the representative agent's problem of producing consumption goods rather than gasoline, for the purposes of studying the role of inventories in a broad perspective - facing global phenomena of accumulating crude oil inventories, or petroleum as a whole. In spite of this subtle difference, I use the term "refinery" interchangeably for the facile understanding of our theoretical model and its implication.

The optimality conditions for the representative producer are summarized as:

$$\alpha \exp(\alpha z_t) q_t^{\alpha-1} \{1 - \exp(-\theta_t i_{t-1})\} = S_t, \quad (2)$$

$$E_t [\Lambda_{t+1} \theta_{t+1} \exp(-r_{t+1} - \theta_{t+1} i_t + \alpha z_{t+1}) q_{t+1}^{\alpha}] = S_t (1 + c_1 i_t) - E_t [\Lambda_{t+1} \exp(-r_{t+1}) S_{t+1}]. \quad (3)$$

Equation (2) is the optimality condition associated with the refinery buying one more barrel of crude oil and using the crude oil immediately for its current production. The resource demand is determined where the marginal product equals the marginal cost. Equation (3) is the optimality condition for the storage decision, in which the refinery equates expected marginal benefits with marginal costs of holding an additional barrel of crude oil inventory. A positive value of the expected marginal benefits, left-hand side in (3), introduces the convenience yield as directly increasing the refinery's future production capabilities. The representative producer

³In a standard macro model, we have $\Lambda_{\tau} = \frac{\beta \cdot U'(c_{\tau+1})}{U'(c_{\tau})} e^{r_{\tau}}$, where $U(\cdot)$ represents the agent's utility function, β is the discount factor and r_{τ} is the interest rate.

has an incentive of holding a positive level of inventory despite an expected capital loss in the future.

2.4 Aggregation, Seasonality and the Crude Oil Spot Price

The theory presented so far applies to one representative producer; however, available observations are aggregated at the global level. Given the trend growth observed from the global crude oil production, consumption and inventories, we assume that the number of representative producers increases over time at a fixed rate. Let N_t be the number of representative producers at period t and g be the exogenously determined average growth rate of the representative producer, i.e. $N_t = (1 + g) N_{t-1}$. The aggregate resource demands (Q_t) and inventories (I_t) become $Q_t = N_t \cdot q_t$ and $I_t = N_t \cdot i_t$, where Q_t and I_t will be associated with globally observed aggregate quantities of the crude oil consumed and inventories. Assuming competitiveness in the global refinery industry following Eckstein and Eichenbaum (1985) and Pindyck (2001), the aggregate production F and cost functions C become $F(\cdot) = N_t \cdot f(\cdot)$ and $C(\cdot) = N_t \cdot c(\cdot)$ where the representative producer's production $f(\cdot)$ and the cost function $c(\cdot)$ are as defined earlier.

Next, I introduce seasonally varying utilization parameters to deal with the strong seasonality observed both in the crude oil consumption and inventories. Specifically, I conjecture that seasonal variations exist in the representative producer's benefit, accordingly, the seasonally varying convenience yield. In the northern hemisphere, for example, the aggregate demands for gasoline increase during the summer for traveling. In most cases, refiners are able to predict these seasonally varying aggregate demands and tend to accumulate inventories in advance when it is profitable to meet increasing demands for gasoline by utilizing their production facilities efficiently. I propose to capture seasonal variations of the utilization parameters as, θ_t equals θ_1 if the month of the time period t corresponds to March, April, May, θ_2 for June, July, August, θ_3 for September, October, November, and θ_4 for December, January, February.

Lastly, I model the crude oil spot price process following the long-term/short-term model of Schwartz and Smith (2000) because of the model's flexibility to capture observed behaviors in the crude oil prices. In this empirical application, I adapt their model to a discrete time stochastic process to obtain the closed-form forecasts for the future crude oil spot price.

Suppose that the logarithm of crude oil spot price ($\log S_t$) consists of the long-term trend

(ξ_t) and the short-term deviation (χ_t) .⁴ The long-term trend follows the random walk process, reflecting the idea that oil prices themselves behave like a random walk at each time period. On the other hand, the short-term deviation from the long-term trend follows the mean-reverting AR(1) process:

$$\begin{aligned}\ln S_{t+1} &\equiv \xi_{t+1} + \chi_{t+1}, \\ \xi_{t+1} &= \xi_t + \epsilon_{2,t+1}, \\ \chi_{t+1} &= k \cdot \chi_t + \epsilon_{3,t+1},\end{aligned}$$

where $k \in (-1, 1)$ is the mean-reversion parameter. $\epsilon_{2,t+1}$ and $\epsilon_{3,t+1}$ are normally distributed innovation processes with $E[\epsilon_{2,t+1}] = E[\epsilon_{3,t+1}] = 0$, $Var[\epsilon_{2,t+1}] = \sigma_2^2$, $Var[\epsilon_{3,t+1}] = \sigma_3^2$ for $\forall t$. Though not serially uncorrelated with their own processes, innovations are correlated to each other only at the same time period, i.e., $corr[\epsilon_{2,t+1}, \epsilon_{3,s+1}] = \rho$ for $\forall t = s$.

Given the crude oil spot price process, I adopt the risk-neutral valuation framework to solve the representative producer's forecasting problem in the storage decision (3) as if he is risk-neutral.⁵ Specifically, I specify risk-neutral processes by subtracting risk premiums from underlying stochastic processes of state variables (z_t, ξ_t, χ_t) . While risk premiums are equilibrium prices for risks the producer pays for his hedging activities in the crude oil market, this form of risk adjustment is frequently adopted in the literature (Schwartz and Smith 2000, Pirrong 2011, and Hamilton and Wu 2014). Denoting by $\Upsilon \equiv [\tilde{\lambda}_z, \tilde{\lambda}_\xi, \tilde{\lambda}_\chi]'$ risk premiums, I assume zero, constant, and linear state-dependent risk premiums on the technology (z_t) , long-term trend (ξ_t) and short-term deviation processes (χ_t) respectively: $\tilde{\lambda}_z \equiv 0$, $\tilde{\lambda}_\xi = \lambda_\xi$, $\tilde{\lambda}_\chi \equiv \lambda_\chi + \varpi\chi_t$ with

⁴Another way to understand the adopted spot price process is a two-factor model. While the long-term factor evolves according to a random walk process, the short-term factor reflects a temporal deviation from the first factor. In this paper, the level of crude oil inventories, accordingly the convenience yield, helps disentangle the long- and short-term factors from the observed crude oil futures prices of which risk premiums are adjusted by the risk-neutral valuation framework.

⁵Implicit assumptions are the deterministic interest rate and the redundancy of the futures contract. The former guarantees the price equivalence between futures and forward contracts. More importantly, the latter validates the proposed approach of specifying risk-neutral processes by subtracting the risk premiums from underlying processes. See more details for the application of the risk-neutral valuation framework in Duffie (Dynamic Asset Pricing Theory, 2001, pp. 167-174)

$\lambda_\xi > 0$, $\lambda_\chi > 0$ and ϖ being constants. Hence, risk-neutral stochastic processes are of the form:

$$\begin{bmatrix} z_{t+1} \\ \xi_{t+1} \\ \chi_{t+1} \end{bmatrix} = \begin{bmatrix} 0 \\ -\lambda_\xi \\ -\lambda_\chi \end{bmatrix} + \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & k^Q \end{bmatrix} \begin{bmatrix} z_t \\ \xi_t \\ \chi_t \end{bmatrix} + \begin{bmatrix} \epsilon_{1,t+1}^Q \\ \epsilon_{2,t+1}^Q \\ \epsilon_{3,t+1}^Q \end{bmatrix},$$

where $k^Q \equiv k - \varpi \in (-1, 1)$. Assuming that $\epsilon_{1,t+1}^Q$ is independent of both $\epsilon_{2,t+1}^Q$ and $\epsilon_{3,t+1}^Q$, innovations for risk-neutral stochastic processes ($\epsilon_{1,t+1}^Q, \epsilon_{2,t+1}^Q, \epsilon_{3,t+1}^Q$) have identical properties as described earlier.

Since the crude oil spot price is conditionally log-normally distributed, the closed-form forecasting for the one-period-ahead crude oil spot price is provided as⁶

$$E_t^Q [S_{t+1}] = \exp \left\{ \xi_t + k^Q \chi_t - \lambda_\xi - \lambda_\chi + \frac{1}{2} s_1^Q \right\}, \quad (4)$$

where $s_1^Q \equiv Var_t^Q [\log S_{t+1}] = \sigma_2^2 + \sigma_3^2 + 2\rho\sigma_2\sigma_3$. Here the superscript Q indicates the calculation under the risk-neutral processes as opposed to those under the physical measure such as $E_t[\cdot]$ and $Var_t[\cdot]$.

Similarly, the τ -periods-ahead crude oil spot price forecast is inferred recursively by

$$\begin{aligned} E_t^Q [S_{t+\tau}] &= \exp \left\{ E_t^Q [\log S_{t+\tau}] + \frac{1}{2} Var_t^Q [\log S_{t+\tau}] \right\} \\ &= \exp \left\{ \xi_t + (k^Q)^\tau \cdot \chi_t - \tau \cdot \lambda_\xi - \lambda_\chi \left(1 + k^Q + \dots + (k^Q)^{\tau-1} \right) + \frac{1}{2} s_\tau^Q \right\}, \end{aligned} \quad (5)$$

where $s_\tau^Q \equiv \tau \cdot \sigma_2^2 + \frac{1-(k^Q)^{2\tau}}{1-(k^Q)^2} \cdot \sigma_3^2 + 2\rho\sigma_2\sigma_3 \cdot \frac{1-(k^Q)^\tau}{1-k^Q}$ is the conditional variance associated with τ -periods-ahead forecast of the crude oil spot price.

⁶The conditional expectation and the variance of $\log S_{t+1}$ under the risk-neutral processes are respectively given as, $E_t^Q [\log S_{t+1}] = \xi_t + k^Q \chi_t - \lambda_\xi - \lambda_\chi$ and $Var_t^Q [\log S_{t+1}] = \sigma_2^2 + \sigma_3^2 + 2\rho\sigma_2\sigma_3$. Hence, one-month-ahead crude oil spot price forecast becomes $E_t^Q [S_{t+1}] = \exp \left(E_t^Q [\log S_{t+1}] + \frac{1}{2} Var_t^Q [\log S_{t+1}] \right)$ due to Jensen inequality. For arbitrary $\phi > 0$, it can be further generalized as $E_t^Q [(S_{t+1})^{-\phi}] = \exp \left\{ -\phi \cdot (\xi_t + k^Q \cdot \chi_t - \lambda_\xi - \delta) + \frac{1}{2} \phi^2 s_1^Q \right\}$.

2.5 Equilibrium Prediction

With approaches introduced in the previous subsection, the producer's optimality conditions in (2) and (3) become

$$\alpha \exp(\alpha z_t) (Q_t/N_t)^{-(1-\alpha)} \{1 - \exp(-\theta_t I_{t-1}/N_{t-1})\} = \exp(\xi_t + \chi_t), \quad (6)$$

$$\begin{aligned} & \alpha^\phi \theta_{t+1} (1+g) \{1 - \exp(-\theta_{t+1} I_t/N_t)\}^\phi \exp(-\theta_{t+1} I_t/N_t) \\ & \times \exp \left\{ -r_{t+1} + \phi (z_t - \xi_t - k^Q \chi_t + \lambda_\xi + \lambda_\chi) + \frac{1}{2} \phi^2 (s_1^Q + \sigma_1^2) \right\} \\ = & \exp(\xi_t + \chi_t) \left[1 + c_1 I_t/N_t - \exp \left\{ -r_{t+1} - (1 - k^Q) \chi_t - \lambda_\xi - \lambda_\chi + \frac{1}{2} s_1^Q \right\} \right], \end{aligned} \quad (7)$$

where $\phi \equiv \frac{\alpha}{1-\alpha} > 0$ and $\theta_t, \theta_{t+1} > 0$ are utilization parameters associated with the production in the current period and the storage decision for the next period, respectively.⁷ In equilibrium, the aggregate resource and inventory demands are provided as a function of exogenous (r_{t+1}, N_t),⁸ and endogenous state variables ($I_t, Q_t, z_t, \xi_t, \chi_t$).

Notice that the left-hand side of the inventory demand equation (7) is the convenience yield, which is the refinery's expected benefit from holding an additional barrel of crude oil. The proposed convenience yield is time-varying (Gibson and Schwartz 1990) in line with state variables; it is inversely related to the current inventories (Pindyck 1994, Alquist and Kilian 2010, and Knittel and Pindyck 2015), indicating that inventory holders have larger benefits when inventories are relatively low. Furthermore, it exhibits asymmetric responses to the long-term dynamics and short-term deviation components of the crude oil spot process (Schwartz and Smith 2000). In general, the proposed convenience yield is consistent with the theory of storage.

However, it is worth noting that the inverse relationship in this paper is derived within the optimal storage decision of the refinery who naturally has an incentive to hold crude oil

⁷For example, θ_1 are used for both θ_t and θ_{t+1} when the period t corresponds to March. However, we use θ_2 in place of θ_{t+1} when the period t corresponds to May. It is because it becomes summer in the next period while it is currently spring.

⁸In the aggregate resource demand equation (6), the number of the representative producers in the previous period (N_{t-1}) works as "positive externality," yet the number at current period (N_t) works as "negative externality." This confirms our intuition under the limited storage facilities; an increasing competition tends to raise marginal storage costs in the short run due to congestion, but it also motivates growths in storage technologies in the long run.

inventories, which has never been done in the literature. Instead, researchers commonly assume such a negative relationship when studying the behavior of crude oil price and its relationship with inventories via convenience yield. When measuring the convenience yield, the proposed storage model also differs from earlier researchers' approaches in two ways. First, the proposed convenience yield reflects convex storage costs that are essential considerations for refiners (or inventory holders in general) besides their expectations on future price movements. Second, data on crude oil inventories are used along with estimates for parameters and unobservable state variables. These contrast with the widely adopted approach of inferring the net convenience yield (abstracting marginal storage costs) by using the term spread calculated from prices for crude oil futures contracts (Pindyck 1994, Alquist and Kilian 2010, and Knittel and Pindyck 2015). In the following section, I provide a historical convenience yield measure calculated from the inventory demand equation (7), using parameter estimates as well as state variables that include data on crude oil inventory.

3 Empirical Results

3.1 Data and Empirical Implementation

The monthly global crude oil inventories, consumptions and prices for crude oil futures contracts are used for estimating the model. Given the lack of data on global crude oil inventories, I construct the monthly series of global crude oil inventories following Kilian and Murphy (2014), where total U.S. crude oil inventories are scaled by the ratio of Organisation for Economic Co-operation and Development (OECD) petroleum stocks over U.S. petroleum stocks. Next, I construct the monthly series of global crude oil consumptions by subtracting monthly changes in global crude oil inventories from monthly global crude oil productions. While OECD petroleum stocks is provided by the International Energy Agency, U.S crude oil inventories, petroleum stocks and global crude oil productions are provided by the U.S. Energy Information Administration (EIA). Then, I construct monthly prices for the crude oil futures contracts from daily prices for the light sweet crude oil (aka WTI crude oil) futures contracts traded in the New York Mercantile Exchange. More specifically, after obtaining daily prices for futures contracts having maturities up to 13 months from datastream, I construct monthly prices of a given maturity

using the first calendar date price of corresponding futures contracts.⁹ The monthly U.S. LIBOR is used for representing the effective historical risk-free interest rate. Lastly, the U.S.CPI, provided by OECD, is used for deflating nominal prices of the crude oil futures contracts into real prices as well as adjusting nominal interest rates for realized inflation. The time period of the analysis is from March 1989 to November 2014, where the beginning period of the analysis is determined from the availability of prices for a 12-month crude oil futures contract, and the ending period is determined from the availability of the OECD Inventory.

Assuming an inelastic supply of crude oil following Kilian (2008) and Alquist and Kilian (2010),¹⁰ I estimate model parameters along the aggregate demands provided by (6) and (7). As noted earlier, I treat crude oil spot price as being unobservable since the spot price is not available in general.¹¹ Instead, I use prices of crude oil futures contracts to estimate unobservable components in the crude oil spot price following Schwartz and Smith (2000). Furthermore, I allow for measurement errors in the observables such as crude oil demands, inventories and prices for crude oil futures contracts during the estimation procedure. This addresses the widely recognized measurement issues in oil quantity variables (Baumeister and Kilian 2012) and mispriced oil futures contracts at times (Kilian and Murphy 2014)

Now, I provide three sets of observation equations as linear functions of state variables (z_t, ξ_t, χ_t) for the estimation of the proposed storage model. The first two observation equations are readily available from equilibrium decisions for the aggregate resource demands and inventory demands.

Consider the aggregate resource demands from (6). Taking the log and subtracting the

⁹In general, the construction of the monthly prices is consistent with the procedure described in Alquist and Kilian (2010). One exception is to use the price of τ -month futures contract at the first calendar date in this paper, compared to using the end-of-month price of the same futures contract in Alquist and Kilian (2010). Both approaches are intended to avoid the potential liquidity issue in the futures market since trading in the current delivery month ceases on the third business day prior to the 25th calendar day of the month preceding the delivery month.

¹⁰Given capacity constraints in world crude oil production, the supply side of the crude oil market has been commonly modeled by an endowment structure, where unexpected disruption/discovery in the oil production shifts the supply curve. While estimating parameters along the aggregate demands, I confirm that forecasting errors in resource and inventory demand equations are not correlated with the exogenous crude oil supply shock suggested in Kilian (2009).

¹¹Strictly speaking, the spot price is not available. While using the cash price for proxying the spot price, Knittel and Pindyck (2015) note the difference between the two; whereas the spot price is a price for immediate delivery at a specific location of a specific grade of oil, the cash price refers to an average transaction price, which reflects discounts or premiums resulting from buyer-seller relationships.

previous aggregate resource demands ($\log Q_{t-1}$) from the resulting equation yields

$$\begin{aligned} \log(Q_t/Q_{t-1}) &= \phi z_t - (1 + \phi) \xi_t - (1 + \phi) \chi_t \\ &+ (1 + \phi) \log \alpha (1 - \exp(-\theta_t I_{t-1}/N_{t-1})) + \log N_t - \log Q_{t-1}. \end{aligned} \quad (8)$$

Similarly, consider the log of the aggregate inventory demands (7) and linear approximations of non-linear terms in the resulting equation. Linear approximations are taken for $\log(I_t)$ being around $\log(I_{t-1})$, and for χ_t being close to $\frac{-\lambda_\chi}{1-k^Q}$ since χ_t is mean-reverting around $\frac{-\lambda_\chi}{1-k^Q}$. Rearranging this, we obtain a linear relationship between a change in the aggregate demands for inventories (I_t) and state variables (z_t, ξ_t, χ_t) as

$$\log(I_t/I_{t-1}) = \Phi_1(x_t) \cdot z_t + \Phi_2(x_t) \cdot \xi_t + \Phi_3(x_t) \cdot \chi_t + \Phi_4(x_t), \quad (9)$$

where coefficients $\Phi_1(x_t)$, $\Phi_2(x_t)$, $\Phi_3(x_t)$, $\Phi_4(x_t)$ are functions of parameters and a vector of observations $x_t \equiv (r_{t+1}, I_{t-1})$ that are either exogenous or predetermined at period t . Appendix A.1. provides detailed derivations as well as coefficients of the above inventory observation equation (9).

Lastly, I use closed-form spot price forecasts provided in (5) for evaluating prices of crude oil futures contracts with various maturities. To see this, consider the risk-neutral stochastic processes in Section 2.4. Under the risk-neutral valuation paradigm, the “no arbitrage” price of the contingent claims coincides with the expected future cash flows under the risk-neutral stochastic process. Since there is no cash payment when a futures contract is traded, the price of the crude oil futures contract with one month to maturity ($F_{t,1}$) coincides with the one-period-ahead spot price forecast under the risk-neutral measure. Similarly, the price of the crude oil futures contract with τ periods to maturity ($F_{\tau,t}$) is the τ -period-ahead spot price forecast under the risk neutral measure, i.e., $F_{t,\tau} = E_t^Q[S_{t+\tau}]$.¹² After taking log on both sides, I subtract the logarithm of the same crude oil futures contract in the previous period ($\log F_{t-1,\tau+1}$). This

¹²Hamilton and Wu (2014) shed new light on the time-varying risk premium in oil futures markets, relying on an affine factor structure to futures prices. In this paper, risk-premium is also time-varying and depends on a short-term deviation process, which is a factor reflecting the interaction between the level of inventories and temporal behaviors of spot prices. To see this, the risk premium associated with τ -month to maturity crude oil futures contract is $(E_t^Q[S_{t+\tau}] - E_t[S_{t+\tau}]) / E_t[S_{t+\tau}] \approx \log E_t^Q[S_{t+\tau}] - \log E_t[S_{t+\tau}]$. Using spot price forecasts under the risk-neutral measure (5) and under the physical measure without risk adjustments, this further reduces to $A(k, k^Q; \tau) \cdot \chi_t + B(k, k^Q, \lambda_\xi, \lambda_\chi, \sigma_2, \sigma_3, \rho; \tau)$, a time-varying function of the short-term deviation process (χ_t), as well as parameters through $A(\cdot)$ and $B(\cdot)$, of which signs and magnitudes depend on model parameters.

provides a linear relationship between a periodic return of the crude oil futures contracts with τ periods to maturity and state variables (ξ_t, χ_t) as

$$\log(F_{t,\tau}/F_{t-1,\tau+1}) = \xi_t + (k^Q)^\tau \cdot \chi_t - \tau \cdot \lambda_\xi - \lambda_\chi \cdot \frac{1 - (k^Q)^\tau}{1 - k^Q} + \frac{1}{2}s_\tau^Q - \log F_{t-1,\tau+1}, \quad (10)$$

where $s_\tau^Q \equiv \tau \cdot \sigma_2^2 + \frac{1 - (k^Q)^{2\tau}}{1 - (k^Q)^2} \cdot \sigma_3^2 + 2\rho\sigma_2\sigma_3 \cdot \frac{1 - (k^Q)^\tau}{1 - k^Q}$ is the conditional variance associated with τ -periods-ahead forecast of the crude oil spot price.

Given three sets of observation equations, namely, (8) for crude oil consumption, (9) for inventories and (10) for futures contract prices for $1, \dots, 12$ months to maturity, I estimate parameters by maximum likelihood using the Kalman filter. Let $\zeta_t = (z_t, \xi_t, \chi_t)'$ denote a 3×1 vector of unobservable state variables and y_t denote a 14×1 vector of observed variables at each period t , i.e., $y_t = (\ln Q_t/Q_{t-1}, \log(I_t/I_{t-1}), \ln F_{t,1}/F_{t-1,2}, \dots, \ln F_{t,12}/F_{t-1,13})'$. Then y_t can be described in terms of a linear function of unobservable state vector ζ_t . The state-space representation is given by the following system of equations:

$$\xi_{t+1} = C + G\xi_t + v_{t+1},$$

$$y_t = a(x_t) + H(x_t)'\xi_t + w_t,$$

where $v_{t+1} = (\epsilon_{1,t+1}^Q, \epsilon_{2,t+1}^Q, \epsilon_{3,t+1}^Q)'$ is a 3×1 vector of risk-neutral innovation processes and w_t is a 14×1 vector of measurement errors. In the transition equation, $C = (0, -\lambda_\xi, -\lambda_\chi)'$ and $G = \text{diag}(1, 1, k^Q)$ with diag being a diagonal matrix operator whose diagonal elements are provided in the parentheses. In the observation equation, $H(x_t)'$ and $a(x_t)$ are specified following (8), (9) and (10).

Given observations (y_t, x_t) for $t = 1, \dots, T$, the sample log-likelihood can be constructed using the Kalman filter

$$\mathcal{L}(\Theta) = \sum_{t=1}^T \log f_{Y_t|X_t, \mathbf{Y}_{t-1}}(y_t|x_t, \mathbf{y}_{t-1}),$$

where $f_{Y_t|X_t, \mathbf{Y}_{t-1}}(y_t|x_t, \mathbf{y}_{t-1})$ is the conditional likelihood for a 14×1 vector of observables conditional on their lagged values as well as available observations that are either exogenous or predetermined. Appendix A.2. provides a detailed explanation for calculating the above sample

log likelihood using the Kalman filter as well as parametric restrictions for $H(x_t)'$ and $a(x_t)$ in the observation equation.¹³

3.2 Results

The second column in Table 1 reports maximum likelihood estimates and asymptotic standard errors.¹⁴ Estimates of the model parameters are consistent with the theoretical restrictions and historical observations. As for the estimates themselves, several points stand out. First, the output share (α) estimate indicates that crude oil resources account for approximately 39% of consumption good production globally. Second, when the real price of crude oil was \$99.21 per barrel and the inventory level was 4,138.71 million barrels per day in January 2008, estimates for storage cost functions indicate that the marginal storage cost necessary for increasing an additional barrel of crude oil inventory is approximately \$0.91 per barrel, holding other things being equal. Third, I find smaller estimates of the utilization parameters for spring (θ_1) and summer (θ_2) than other seasons. These estimates indicate that utilizations for spring and summer are higher than those for fall and winter since the smaller θ implies the refinery's higher dependence on available inventories. While all seasonal utilization parameters ($\theta_1, \theta_2, \theta_3, \theta_4$) are statistically significant, the smallest difference among utilization parameters is 0.0009 between fall (θ_3) and winter (θ_4).¹⁵ When crude oil inventories correspond to the median level of 2,455.15 millions barrels per day during July 2003, a 0.0009 decrease in θ yields a \$0.10 per barrel increase in the convenience yield holding other variables constant. Fourth, positive mean reversion (k^Q) indicates that the short-term deviation of the crude oil spot price is highly persistent under the risk-neutral measure. Given the slowly moving long-term trend and the positive correlation (ρ^Q) between the two underlying processes, this further implies the highly persistent crude oil spot process under the risk-neutral measure. Fifth, the long-term risk premium (λ_ξ) is positive and statistically significant, indicating a risky long-term trend component in crude oil price movement. Given the highly persistent oil price, producers need to provide counterparties

¹³See more details regarding the state-space representation and the Kalman filter for maximum likelihood estimation in Hamilton (Time Series Analysis, 1994).

¹⁴To estimate the model, the likelihood function is reparametrized while preserving restrictions on parameters. The reported MLE are re-calculated from the estimated MLE, and asymptotic standard errors are calculated based on initial parametric specifications.

¹⁵In general, seasonal utilization parameters are statistically different from each other at 10% significance level except two consecutive pairs of summer (θ_2) and fall (θ_3), and fall (θ_3) and winter (θ_4). Wald statistic for testing the null hypothesis of $H_0 : \theta_1 = \theta_2 = \theta_3 = \theta_4$ is 5.99 with the corresponding p-value being 0.11.

with sufficient compensations for hedging their price risks, especially being associated with the long-term risk factor. On the other hand, the short-term risk premium (λ_χ) is negative, yet statistically insignificant. Hence, arbitrageurs do not necessarily require a high risk premium when providing liquidities for hedgers when facing the short-term risk factor.

Figure 1 plots the historical decomposition of the crude oil spot price with the long-term trend (Panel 1) and the short-term deviation (Panel 2). The model fit is provided in Panel 3 by overlaying the model forecasts from (4) with the real price of the nearest maturity crude oil futures contract (henceforth, crude oil price) on a logarithm scale. In the historical decomposition, previous episodes in the crude oil market are shown through the lens of the long-term and short-term processes. As a random walk process, the long-term trend represents a slowly moving trend in the crude oil spot price. While stable until 2003, the long-term trend has been increasing from the beginning of 2004 to the first half of 2008, reaching at the higher equilibrium level compared to the previous period. On the other hand, the short-term deviation process is a mean-reverting process around its long-run average ($\frac{-\lambda_\chi}{1-kQ} = 0.44$), of which historical observations are consistent with previous episodes documented in the earlier literature (Hamilton 2009, Kilian 2009, and Kilian and Murphy 2014).

Panel 1 in Figure 2 plots the equilibrium convenience yield as a percentage of the crude oil spot price. The convenience yield fluctuates considerably along the business cycle, repeating the same patterns around the past recession and recovery periods. When the global economy starts to recover from the early 2000s' recession, for example, the convenience yield sharply increases from 21.72% in February 2002 to 24.03% in February 2003. On the other hand, it declines from 18.43% to 16.67% between December 2007 and June 2009 (Great Recession). The procyclical convenience yield can be explained by its negative relationship with crude oil inventories (Pindyck 1994). During the early 2000 recovery period, the crude oil inventories declined by 0.12%, which reflects increasing crude oil demands for immediate gasoline productions as well as for future uses. On the other hand, crude oil inventories had increased by 3.87% during the Great Recession. The correlation coefficient estimates are -0.46 with the level of crude oil inventories and -0.21 with the monthly changes in crude oil inventories. This finding is consistent with the theory of storage literature. Moreover, I find asymmetric responses of the convenience yield toward two components in the crude oil spot price. The convenience yield is positively (negatively) related to the short-term deviation (long-term trend) components with the correla-

tion coefficient being 0.56 (-0.45) over the sample period. When the convenience yield increased during the early 2000 recovery period, the short-term deviation increased by 95.12%, whereas the long-term trend decreased by 3.00%. During the Great Recession, the short-term deviation had decreased by 33.24%, whereas the long-term component had increased by 3.17%. Lastly, I find that the historical average convenience yield is 19.10% over the sample period. Holding other factors of production constant, historical average gains in the refinery's productivity can be translated into a significant benefit such as 19.10% relative to the spot price at which the refinery purchases an additional barrel of crude oil for future production. This again highlights the importance of using data on inventories and taking associated storage costs into account when measuring convenience yield.

For comparison, Panel 2 plots the negative of crude oil futures spreads, which are widely adopted convenience yield measures among researchers. Here I extend the spread series in Figure 4 of Alquist and Kilian (2010) using monthly real prices for futures contracts with three, six, nine and 12-months to maturity and real crude oil spot price provided by the EIA. In general, observations from negative spreads are analogous to those from the proposed convenience yield in Panel 1. One exception is the Persian Gulf War period, where the negative spreads exhibit the positive spike provided by the sharp spikes in prices for futures contracts. Besides the fact that the size of spikes are larger for the spreads from longer-term futures contracts,¹⁶ it can be explained by the hypothesized role of crude oil inventories, leaving a precautionary role as unmodelled dynamics. Despite such a difference, one can visually confirm that the proposed convenience yield exhibits qualitatively similar patterns during historical events, documented in Alquist and Kilian (2010).

3.3 Stable Forecasting Relationship

Given increased financial investors' participation since 2004, it has been a controversial issue that these new participants have changed the crude oil market structure. For example, Tang and Xiong (2012) find an increased correlation between oil prices and prices of non energy

¹⁶The proposed convenience yield in Panel 1 measures the expected benefit from one-month storage of an additional barrel of crude oil. From this perspective, the spread to be compared with Panel 1 should be the shortest spread defined as a (log) real price difference between spot and futures contract with one month to maturity. When calculating this, I find the small magnitude of the one-month spread to be compared with the convenience yield. Moreover, for a few periods, it exhibits erroneous behaviors that are qualitatively different from spreads provided by longer maturity futures contracts.

commodities, raising doubts about potential spillover effects on commodity markets from other financial markets. On the other hand, Fattouh (2010) shows that the increased liquidity provided by financial investors facilitates arbitrage activities, enhancing oil market integration. In this section, I begin by providing evidence of a structural change by estimating the model under two sub-sample periods split after June 2004. Then, I show that recursive forecasts provided by the proposed model are more accurate than the no-change forecast at three, six, nine and 12-month horizons, highlighting that the price-inventory relationship is stable despite a structural change in the crude oil market.

I estimate the proposed model again using data for two sub-sample periods split after June 2004. The first sub-sample spans from March 1989 to June 2004 (184 observations), and the second sub-sample spans from July 2004 to November 2014 (125 observations). The third and fourth columns in Table 1 report maximum likelihood estimates and asymptotic standard errors. In general, estimates are qualitatively similar to each other and those from the total sample. Turning to the sizes of parameter estimates, however, there is a clear sign of a structural change in the oil market. Here, I find a reduced role of crude oil inventory relative to timely purchased resources from an increased convenience yield parameter (θ_t) while preserving the same seasonal pattern. The convenience yields calculated from parameter estimates across two sub-sample periods exhibit the same piecewise historical movement, yet at different levels. Holding other factors of production constant across two periods, the average convenience yields are 33.15% and 15.00%, respectively, for the first and second periods. However, this does not indicate the lower convenience yield during the latter period due to the omitted factors in the refinery's gasoline production.¹⁷ Instead, when combined with the increased crude oil inventories, such a difference is reflected in the model through higher trend (g) and lower marginal storage cost (c_1) estimates for the second period. From parameter estimates describing the crude oil spot price, I find that the statistical properties for the crude oil spot price has also changed. More specifically, the crude oil spot price becomes more persistent, which is evidenced by the increased autocorrelation (k^Q) for the mean-reverting short-term process and increased cross-correlation (ρ^Q) between short- and long-term processes. The increased persistence of the crude oil spot price results in higher risk premiums that are necessary for refiners to hedge heightened risks

¹⁷Similarly, the smaller output share (α) from the second sub-period does not imply the reduced importance of the crude oil resource in gasoline production.

in the crude oil market. This can be confirmed by the increased prices of both long- (λ_ξ) and short-term (λ_χ) risks during the latter period. To sum up, the sub-sample exercise in this section provides empirical evidence of the structural change in the crude oil market, which are consistent with Tang and Xiong (2012), Büyüksahin and Robe (2014) and Hamilton and Wu (2014) among others.

Next, I perform an out-of-sample forecasting exercise to investigate whether the structural change affects the price-inventory relation proposed in this paper. Throughout the forecasting exercise, I forecast the log price of oil because the model specifies the relationship between the log price and the state variables. To reduce computational burden during the recursive estimation procedure, I use a smaller set of observations consisting of futures contract prices with one, three, six and 12 months to maturity, oil consumptions and inventories. Given evidence for the structural break, I adopt a rolling fixed estimation window method based on its robustness in the presence of nonstationarity (Giacomini and White 2006). With 184 months of estimation window, I fit the model to a sample of 184 months, generate one-step-ahead forecasts for prices of futures contracts with one, three, six and 12 months to maturity and drop the oldest observation from the sample when adding the new data. I repeat this process and evaluate the performance of 125 monthly out-of-sample forecasts from July 2004 to November 2014. I also generate one, three, six and 12-month forecasts for the crude oil spot price by using fitted unobservable state variables and parameter estimates.¹⁸ To evaluate the performance, I set the benchmark as the random walk model forecast, of which a one-month-ahead forecast for τ -month to maturity futures contract is the current price for the same futures contract and the τ -month-ahead spot price forecasts coincides with the current crude oil spot price.

Table 2 reports recursive mean squared prediction error (MSPE) and mean absolute error (MAE) ratios relative to the no-change forecasts. Panel 1 evaluates performance of one-month-ahead price forecasts for crude oil futures contracts with one, three, six and 12 months to maturity. Compared to the benchmark, the proposed model provides less accurate price forecasts for all futures contracts using MSPE criteria (column 2). Although more accurate forecasts

¹⁸Let $\xi_{t+1|t}$ and $\chi_{t+1|t}$ be one-month-ahead forecasts for long-term trend and short-term deviation processes, respectively. Given parameter estimates, price forecasts for τ -month maturity futures contract are calculated from (10) with unobservable state variables (ξ_t, χ_t) being replaced by their forecasted values ($\xi_{t+1|t}, \chi_{t+1|t}$). Similarly, denoting by $\xi_{t|t}$ and $\chi_{t|t}$ the fitted state variables, the τ -month-ahead forecast for crude oil spot price is calculated from (5) with unobservable state variables (ξ_t, χ_t) being replaced by their fitted values ($\xi_{t|t}, \chi_{t|t}$). Appendix A.2. provides a detailed explanation for the recursive calculation necessary for generating fitted values (A4, A5) and forecasts (A6, A7) of unobservable state variables.

for short-maturity futures contracts are found from the MAE criteria (column 6), I also find that such improvements are not statistically significant based on tests for equal forecasting accuracy suggested by Diebold and Mariano (1995) and West (1996) (column 7). Given abnormal behaviors in the financial markets during the Great Recession period (Gürkaynak and Wright 2012), I recalculated the relative MSPE and MAE ratios after ruling out price forecasts during this period. When excluding forecasts from November 2008 to March 2009, I find that the model provides more accurate one-month-ahead price forecasts for contracts with one and three months to maturity under MSPE (column 4) and one, three, and six months to maturity under MAE (column 8). Yet, the statistical significance for such improvements is weak in general. These results indicate that prices for futures contracts are harder to predict than spot prices. A better forecast of the futures price could come from explicitly modeling the risk premium in these contracts as in Baumeister and Kilian (2014).

Panel 2 reports recursive MSPE and MAE of one, three, six and 12-month-ahead crude oil spot price forecasts relative to the no-change forecast for crude oil spot prices. Compared to the no-change forecast, I find that the model provides more accurate spot price forecasts for all horizons considered under both MSPE and MAE criteria. Such improvements are statistically significant at the 10% level for three and 12-month-ahead forecasts using MSPE (columns 2 and 3), and for three, six and 12-month-ahead forecasts using MAE (columns 6 and 7). When ruling out a few forecasts for the evaluation, I find larger and more statistically significant improvements across forecasting horizons in general under both MSPE and MAE criteria. These findings indicate that a price-inventory relationship proposed in this paper is stable despite potential structural changes in the oil market, consistent with Kilian and Murphy (2014) who develop a four-variable structural VAR model. The improved forecasting accuracies in predicting future crude oil spot price are consistent with Baumeister and Kilian (2012) who document the improved accuracies from widely adopted forecasting models when using real-time data.

4 The Effect of Financial Investors in the Crude Oil Market

So far, I investigate the role of crude oil inventories for the refinery's gasoline production and document empirical evidence for the convenience yield around which crude oil inventories and prices center. Although failing to predict prices of crude oil futures contract, I show that the

proposed storage model provides more accurate crude oil spot price forecasts than the no-change model.

In this section, I illustrate one use of the framework by re-examining the possible contribution of financial investors using the proposed model. To do so, I begin by introducing simple forecasting models that also use oil inventories as an explanatory variable. Then I investigate a breakdown in their forecasting relationships since 2004, which contrasts with the stable price-inventory relationship documented in the previous section.

Traditionally, traders would react to releases of inventory data in a predictable way. Unexpectedly high inventories were taken as a sign of weak demand and strong supplies, and the price would decline. Unexpectedly low inventories would predict an increase in the price. The fact that this relationship seemed less reliable in the period of heavy financial participation led some to conclude that financial investors were changing the fundamental behavior of crude oil prices.

One description of the traditional relation between prices and inventories was developed by Ye et al. (2002), who defined the normal inventory, NI_t as the fitted value from a regression of the observed OECD inventory OI_t on a constant, deterministic time trend and monthly dummies:

$$OI_t = \beta_0 + \beta_1 \cdot t + \sum_{p=2}^{12} \beta_p \cdot D_p, \quad (11)$$

where ($D_p, p = 2, \dots, 12$) are monthly dummies; i.e., $D_p = 1$ if the period t corresponds to the month p , but $D_p = 0$ otherwise. They further defined $RI_t \equiv OI_t - NI_t$ as the residual from this regression. Allowing for the possibility that positive and negative residuals had different effects as well as an effect of the year-over-year change $AI_t = OI_t - OI_{t-12}$, Ye et al. (2002) proposed a forecasting model (hereafter, traditional forecasting model) as follows:

$$F_{1,t} = a + \sum_{i=0}^5 b_i \cdot RI_{t-i} + \sum_{i=0}^5 c_i \cdot LI_{t-i} + d \cdot AI_t + e \cdot F_{1,t-1} + \varepsilon_t, \quad (12)$$

where $F_{1,t}$ is a real price of the nearest-maturity crude oil futures contract in month t and ε_t is a regression error.

Figure 3 plots in-sample and out-of-sample forecasts following the traditional forecasting model. The period for in-sample forecasts coincides with the earlier literature from January

1992 to February 2001. Using estimated coefficients from the in-sample forecasting period, I calculate out-of-sample forecasts from March 2001 to November 2014. The solid line in March 2001 denotes the beginning of out-of-sample forecasts. The long-dashed line in Figure 3 confirms the strong in-sample forecasting ability documented in the earlier literature. However, its strong forecasting ability disappears over the out-of-sample forecasting period.

Given a poor performance of the traditional forecasting model after 2004, Merino and Ortiz (2005) proposed the extended forecasting model by building upon the former model. After expressing RI_t in a percentage of total level ($RI\%_t$) and including U.S. Commodity Futures Trading Commission (CFTC) non-commercial long position (X_t), the extended forecasting model is provided as follows:

$$F_{1,t} = \beta_0 + \beta_1 \cdot t + \beta_2 \cdot \Delta RI\%_t + \beta_3 \cdot RI\%_{t-1} + \beta_4 \cdot \Delta X_t + \beta_5 \cdot X_{t-1} + \eta_t, \quad (13)$$

where $\Delta RI\%_t = RI\%_t - RI\%_{t-1}$, $\Delta X_t = X_t - X_{t-1}$ and η_t is a regression error. One thing to note is that Merino and Ortiz (2005) included the deterministic time trend (t) in the forecasting equation as opposed to the lagged price ($F_{1,t-1}$) considered in Ye et al. (2002). Figure 4 plots in-sample and out-of-sample forecasts following the extended forecasting model. The period for in-sample forecasts coincides with Merino and Ortiz (2005) from January 1992 to June 2004. Using estimated coefficients from the in-sample forecasting period, I calculate out-of-sample forecasts from July 2004 to November 2014. The solid line in July 2004 denotes the beginning of the out-of-sample forecasts. Along with the extended model forecasts, I also calculate the basic model forecasts from the literature, where the latter model excludes X_t in the former forecasting equation (13). Figure 4 confirms the improved in-sample forecasting ability of the extended model as documented in Merino and Ortiz (2005): The extended model exhibits the better in-sample forecasts than the basic model from 2001 to 2004, while both generate similar in-sample forecasts from 1992 to 2001. However, when I extend the sample to include data since publication of Merino and Ortiz's study, even their model with non-commercial traders does quite badly. This finding contrasts with the stable price-inventory relationship evidenced by the improved forecasting accuracies in the previous section.

On the other hand, the optimizing model developed in Section (2) suggests that both equations (12) and (13) are misspecified. To see whether the proposed model exhibits any of the

instabilities of either (12) and (13), I perform an analogous out-of-sample evaluation. In order to highlight the stability of the relationship, I use the same parameter estimates from pre-June 2004 in Table 1 for calculating out-of-sample forecasts recursively from July 2004 to November 2014. Given fixed parameter estimates, I iterate the following three steps recursively for calculating one-month-ahead oil price forecasts while updating available observations. First, given available observations, I update the state vector and the variance matrix using the formula for updating a linear projection. Second, given updated vectors and matrices, I calculate one-month-ahead forecasts with linear projections. Third, given one-month-ahead forecasts, I calculate one-month-ahead out-of-sample forecasts. Figure 5 displays forecasts of the nearest-maturity crude oil futures contract price. The solid vertical line in July 2004 denotes the beginning of the out-of-sample forecasts. I find the stable crude oil price forecasting relationship both in-sample and out-of-sample. In this forecasting exercise, out-of-sample forecasts are based on information that is either known in advance or predetermined at each forecasting period. Hence, I find that currently available information is sufficient to produce a reliable forecast for the next month's crude oil futures price given the stable price-inventory relationship proposed in this paper.

If equation (12) was misspecified, why did it seem to have such a good fit prior to 2001? To answer this question, suppose that inventories had been exactly as predicted by the proposed model; that is, replace OI_t on the left-hand side of (11) with \widehat{OI}_t , the predicted equilibrium level of inventories from equation (9). I then reproduce the Ye et al. exercise on this generated data, the results of which are shown in Figure 3 (short-dashed line). Even though equation (9) is the correct relationship that exactly describes these generated data, and even though (11) and (12) are misspecified, they appear to do a reasonable job until 2004, after which they completely break down.

The explanation for the apparent success of the misspecified model in the earlier period can be found in the fact that the permanent component of crude oil prices, ξ_t , was relatively stable until 2003 but then began to increase substantially (see the top panel of Figure 1). During the period when ξ_t was relatively stable, simple relationships like (11) did not do a bad job of fitting the data, even though they fundamentally mischaracterize the true relationship between inventories and crude oil prices. With the large changes in ξ_t observed since 2004, the need for a better model of the relationship between oil prices and inventories, such as that proposed in this paper, becomes quite evident.

5 Conclusion

This paper proposes an equilibrium storage model of the global crude oil market to study the role of crude oil inventories in refiners' gasoline production. Given a heavy dependence of refinery's production on resources on hand, I model the role of crude oil inventories directly as enhancing gasoline production by treating inventories as an essential factor of production, following Kydland and Prescott (1982, 1988) and Ramey (1989). The model is simple, yet it provides a measure for convenience yield that refiners expect to receive when storing a barrel of crude oil.

Using data on crude oil inventories, consumption and prices for futures contracts, I find empirical evidence on convenience yield that is consistent with the theory of storage literature. I show that the proposed convenience yield is inversely related to inventories, yet positively related to oil prices. While exhibiting seasonal and procyclical behaviors, I show that the historical convenience yield is about 19% of the oil price on average from March 1989 to November 2014. Moreover, I identify a structural change in crude oil market fundamentals since 2004 that coincides with the period of rapidly increasing financial investors' activities. I document empirical evidence on the structural change from a rising long-term component in the spot price process and from changing parameter estimates for before and after June 2004. Despite the structural break, I show that the price-inventory relationship proposed in this paper is stable over time, providing more accurate forecasts for future crude oil prices compared with the random walk model as well as earlier forecasting models. In light of this evidence, I conclude that the contribution of financial investors' activities is weak in the crude oil market.

References

- Acharya, V. V., Lochstoer, L. A., and Ramadorai, T. (2013). Limits to arbitrage and hedging: evidence from commodity markets. *Journal of Financial Economics*, 109(2):441–465.
- Alquist, R. and Kilian, L. (2010). What do we learn from the price of crude oil futures? *Journal of Applied Econometrics*, 25(4):539–573.
- Basak, S. and Pavlova, A. (2015). A model of financialization of commodities. *Journal of Finance*. Forthcoming.
- Baumeister, C. and Kilian, L. (2012). Real-time forecasts of the real price of oil. *Journal of Business & Economic Statistics*, 30(2):326–336.
- Baumeister, C. and Kilian, L. (2014). A general approach to recovering market expectations from futures prices with an application to crude oil.
- Brennan, M. J. (1958). The supply of storage. *American Economic Review*, 48(1):50–72.
- Bryant, H. L., Bessler, D. A., and Haigh, M. S. (2006). Causality in futures markets. *Journal of Futures Markets*, 26(11):1039–1057.
- Büyükkahin, B. and Robe, M. A. (2014). Speculators, commodities and cross-market linkages. *Journal of International Money and Finance*, 42:38–70.
- Considine, T. J. (1997). Inventories under joint production: an empirical analysis of petroleum refining. *Review of Economics and Statistics*, 79(3):493–502.
- Deaton, A. and Laroque, G. (1992). On the behaviour of commodity prices. *Review of Economic Studies*, 59(1):1–23.
- Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3):253–265.
- Duffie, D. (2001). *Dynamic asset pricing theory*. Princeton University Press.
- Eckstein, Z. and Eichenbaum, M. S. (1985). Inventories and quantity-constrained equilibria in regulated markets: the us petroleum industry, 1947–1972. *Energy, Foresight, and Strategy*.
- Fama, E. F. and French, K. R. (1987). Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage. *Journal of Business*, 60(1):55–73.
- Fama, E. F. and French, K. R. (1988). Business cycles and the behavior of metals prices. *Journal of Finance*, 43(5):1075–1093.
- Fattouh, B. (2010). The dynamics of crude oil price differentials. *Energy Economics*, 32(2):334–342.
- Fattouh, B., Kilian, L., and Mahadeva, L. (2013). The role of speculation in oil markets: what have we learned so far? *Energy Journal*, 34(3):7–33.
- Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6):1545–1578.
- Gibson, R. and Schwartz, E. S. (1990). Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance*, 45(3):959–976.
- Gorton, G. B., Hayashi, F., and Rouwenhorst, K. G. (2013). The fundamentals of commodity futures returns. *Review of Finance*, 17(1):35–105.
- Gürkaynak, R. S. and Wright, J. H. (2012). Macroeconomics and the term structure. *Journal of Economic Literature*, 50(2):331–367.

- Hamilton, J. D. (1994). *Time series analysis*. Princeton University Press.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity*, 1(Spring):215–284.
- Hamilton, J. D. and Wu, J. C. (2014). Risk premia in crude oil futures prices. *Journal of International Money and Finance*, 42:9–37.
- Hamilton, J. D. and Wu, J. C. (2015). Effects of index-fund investing on commodity futures prices. *International Economic Review*, 56(1):187–205.
- Irwin, S. H. and Sanders, D. R. (2011). Index funds, financialization, and commodity futures markets. *Applied Economic Perspectives and Policy*, 33:1–31.
- Juvenal, L. and Petrella, I. (2015). Speculation in the oil market. *Journal of Applied Econometrics*, 30:621–649.
- Kaldor, N. (1939). Speculation and economic stability. *Review of Economic Studies*, 7(1):1–27.
- Kennedy, J. P. (2012). The high cost of gambling on oil. *New York Times*, April 10.
- Kilian, L. (2008). Exogenous oil supply shocks: how big are they and how much do they matter for the us economy? *Review of Economics and Statistics*, 90(2):216–240.
- Kilian, L. (2009). Not all oil price shocks are alike: disentangling demand and supply shocks in the crude oil market. *American Economic Review*, 99(3):1053–1069.
- Kilian, L. and Lee, T. K. (2014). Quantifying the speculative component in the real price of oil: The role of global oil inventories. *Journal of International Money and Finance*, 42:71–87.
- Kilian, L. and Murphy, D. P. (2014). The role of inventories and speculative trading in the global market for crude oil. *Journal of Applied Econometrics*, 29(3):454–478.
- Knittel, C. R. and Pindyck, R. S. (2015). The simple economics of commodity price speculation. *American Economic Journal*. forthcoming.
- Kydland, F. E. and Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 50(6):1345–1370.
- Kydland, F. E. and Prescott, E. C. (1988). The workweek of capital and its cyclical implications. *Journal of Monetary Economics*, 21(2-3):343–360.
- Masters, M. W. (2008). Testimony before committee on homeland security and governmental affairs. *US Senate, Washington, May 20*.
- Merino, A. and Ortiz, A. (2005). Explaining the so-called ‘price premium’ in oil markets. *OPEC review*, 29(2):133–152.
- Ng, V. K. and Pirrong, S. C. (1994). Fundamentals and volatility: storage, spreads, and the dynamics of metals prices. *Journal of Business*, 67(2):203–230.
- Pindyck, R. S. (1994). Inventories and the short-run dynamics of commodity prices. *Rand Journal of Economics*, 25(1):141–159.
- Pindyck, R. S. (2001). The dynamics of commodity spot and futures markets: a primer. *Energy Journal*, 22(3):1–30.
- Pirrong, C. (2011). *Commodity price dynamics: a structural approach*. Cambridge University Press.
- Ramey, V. A. (1989). Inventories as factors of production and economic fluctuations. *American Economic Review*, 79(3):338–354.

- Routledge, B. R., Seppi, D. J., and Spatt, C. S. (2000). Equilibrium forward curves for commodities. *Journal of Finance*, 55(3):1297–1338.
- Schwartz, E. and Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7):893–911.
- Smith, J. L. (2009). World oil: market or mayhem? *Journal of Economic Perspectives*, 23(3):145–164.
- Tang, K. and Xiong, W. (2012). Index investment and financialization of commodities. *Financial Analysts Journal*, 68(6):54–74.
- West, K. D. (1996). Asymptotic inference about predictive ability. *Econometrica*, 64(5):1067–1084.
- Working, H. (1949). The theory of price of storage. *American Economic Review*, 39(6):1254–1262.
- Ye, M., Zyren, J., and Shore, J. (2002). Forecasting crude oil spot price using oecd petroleum inventory levels. *International Advances in Economic Research*, 8(4):324–333.

Figure 1: Historical Decomposition of the Crude Oil Spot Price

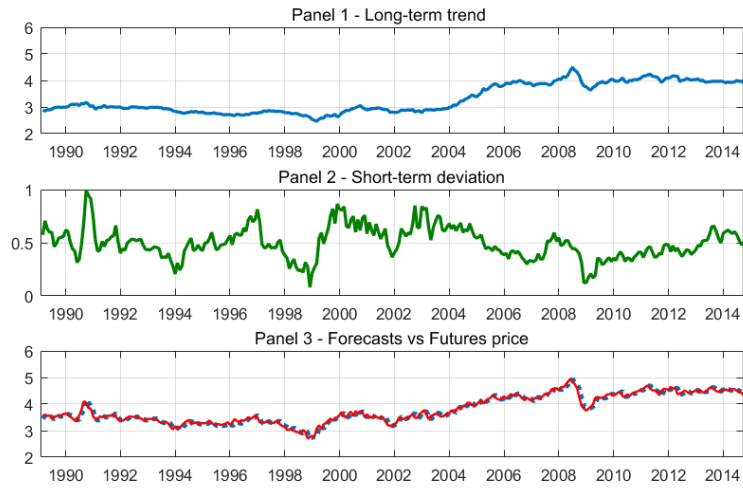


Figure 1 plots the historical decomposition of the crude oil spot price with the long-term trend process (Panel 1) and the short-term deviation process (Panel 2). Panel 3 plots the crude oil spot price (short-dashed line) together with the price for nearest-maturity crude oil futures contract (solid line) on a log scale. The unit is dollars per barrel (2010.5=100) on a log scale.

Figure 2: Convenience yield

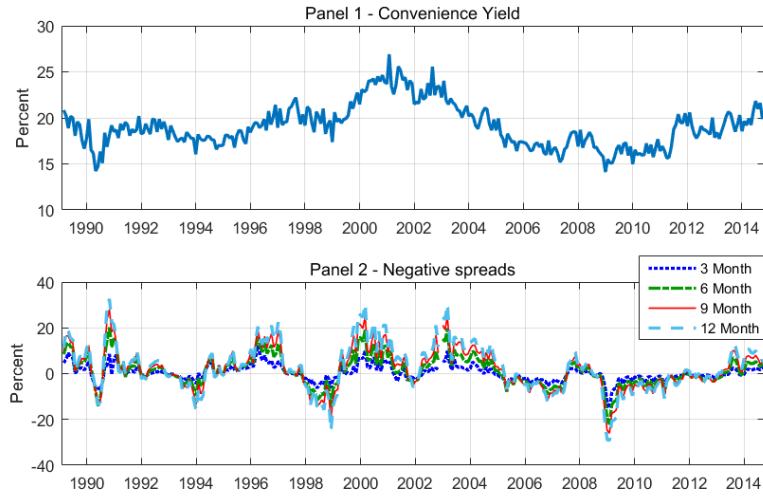


Figure 2 plots the convenience yield measure. Panel 1 displays the equilibrium convenience yield as a percentage of the crude oil spot price. Panel 2 displays negative spreads calculated from prices of crude oil futures contracts with 3-, 6-, 9- and 12-months to maturity, and the real crude oil spot price provided by the U.S. Energy Information Administration.

Figure 3: Forecasts of the Ye et al. (2002) Model

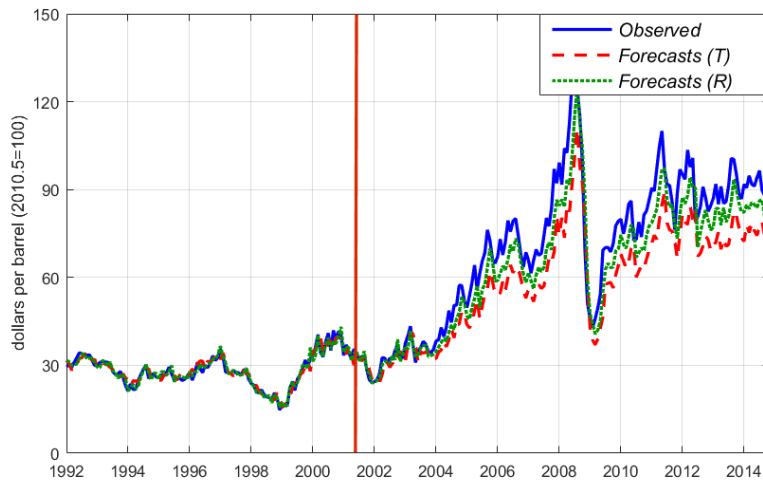


Figure 3 plots in-sample and out-of-sample forecasts of the real crude oil futures price following Ye et al. (2002). In-sample forecasts are obtained by fitting the traditional forecasting model of Ye et al. (2002) from January 1992 to February 2001. Out-of-sample forecasts are calculated recursively from March 2001 to November 2014, using coefficient estimates from the in-sample forecasting period. The traditional model's forecasts, Forecasts (T), are marked by the long-dashed line. The short-dashed line indicates in-sample and out-of-sample forecasts from the same model, with observed crude oil inventories being replaced by those predicted from the proposed storage model, which is denoted by Forecasts (R). For comparison, the solid line plots observed real crude oil futures prices. The vertical line denotes the beginning of out-of-sample forecasts in March 2001.

Figure 4: Forecasts of the Merino and Ortiz (2005) Model




Figure 4 plots in-sample and out-of-sample forecasts of the real crude oil futures price following Merino and Ortiz (2005). In-sample forecasts are obtained by fitting the forecasting models of Merino and Ortiz (2005) from January 1992 to June 2004. Out-of-sample forecasts are calculated recursively from July 2004 to November 2014, using coefficient estimates from in-sample forecasting period. The extended (short-dashed line) and the basic (long-dashed line) refer to the forecasting models with and without a role for non-commercial traders. For comparison, the solid line plots observed real crude oil futures prices. The vertical line in July 2004 denotes the beginning of the out-of-sample forecasts.

Figure 5: Forecasts of the Proposed Storage Model




Figure 5 plots in-sample and out-of-sample forecasts of the real crude oil futures price from the proposed storage model, Forecasts (S). In-sample forecasts are calculated from the proposed storage model by using observations from March 1989 to June 2004, and out-of-sample forecasts are calculated recursively from forecasts for state variables by using parameter estimates from the in-sample estimation period.

Table 1: Parameter Estimates

<i>Period</i>	1989.3-2014.11		1989.3-2004.6		2004.7-2014.11	
<i>Output share</i>						
α	0.3933	(0.0122)	0.4347	(0.0154)	0.3470	(0.0118)
<i>Convenience yield</i>						
θ_1 (<i>spring</i>)	0.1854	(0.0059)	0.1397	(0.0078)	0.1761	(0.0074)
θ_2 (<i>summer</i>)	0.1878	(0.0058)	0.1415	(0.0077)	0.1788	(0.0073)
θ_3 (<i>fall</i>)	0.1896	(0.0058)	0.1434	(0.0077)	0.1811	(0.0072)
θ_4 (<i>winter</i>)	0.1905	(0.0059)	0.1436	(0.0077)	0.1822	(0.0074)
<i>Storage cost</i>						
c_1	0.0051	(0.0003)	0.0152	(0.0017)	0.0063	(0.0008)
<i>Trend</i>						
g	0.0003	(< 0.0000)	0.0002	(< 0.0000)	0.0004	(< 0.0000)
<i>Risk premium</i>						
λ_ξ (<i>long</i>)	0.0026	(0.0002)	0.0001	(0.0004)	0.0063	(0.0004)
λ_χ (<i>short</i>)	-0.0443	(0.0697)	-0.0386	(0.0655)	-0.0351	(0.0486)
<i>Spot process</i>						
k^Q	0.9014	(0.0018)	0.9074	(0.0028)	0.9313	(0.0021)
ρ^Q	0.0968	(0.0596)	-0.0173	(0.0843)	0.1206	(0.0966)
σ_1	0.2470	(0.0126)	0.2246	(0.0143)	0.2638	(0.0190)
σ_2	0.0518	(0.0021)	0.0377	(0.0020)	0.0652	(0.0042)
σ_3	0.0789	(0.0033)	0.0910	(0.0050)	0.0623	(0.0045)
Likelihood	14, 191.74		8, 174.74		6, 852.38	

This table reports maximum likelihood estimates (MLE) for the model parameters and their asymptotic standard errors in parentheses. In the estimation procedure, the likelihood function is reparametrized while preserving restrictions on parameters. The reported MLE are recalculated from the estimated MLE, and asymptotic standard errors are estimated by approximating the second derivative of the log-likelihood functions at MLE estimates. The second column provides parameter estimates for the entire sample period, while the third and last columns report estimates from two sub-samples. In order to preserve space, estimates for forecasting errors are not reported, although they are jointly estimated with other model parameters. The last row reports the maximum value achieved for the log of the likelihood function.

Table 2: Recursive MSPE and MAE Ratios Relative to the No-Change Forecast

	MSPE				MAE			
	Total		Excluded		Total		Excluded	
Panel 1: Prices of futures contracts with τ -months to maturity								
<i>Horizon</i>								
1	1.0028	(<i>n/a</i>)	0.9830	(0.33)	0.9875	(0.25)	0.9755	(0.12)
3	1.0044	(<i>n/a</i>)	0.9881	(0.31)	0.9954	(0.36)	0.9847	(0.14)
6	1.0241	(<i>n/a</i>)	1.0076	(<i>n/a</i>)	1.0063	(<i>n/a</i>)	0.9976	(0.43)
12	1.1120	(<i>n/a</i>)	1.1329	(<i>n/a</i>)	1.0421	(<i>n/a</i>)	1.0419	(<i>n/a</i>)
Panel 2: τ -month-ahead spot price forecast								
<i>Horizon</i>								
1	0.9715	(0.17)	0.9734	(0.23)	0.9806	(0.18)	0.9767	(0.16)
3	0.9479	(0.09)	0.8826	(0.07)	0.9555	(0.03)	0.9395	(0.02)
6	0.9253	(0.11)	0.8394	(0.10)	0.9310	(0.05)	0.9126	(0.05)
12	0.8140	(0.03)	0.7936	(0.05)	0.8776	(0.02)	0.8705	(0.04)

Panel 1 reports the Mean Squared Prediction Error (MSPE) and Mean Absolute Error (MAE) of one-month-ahead forecasts for prices of futures contracts with 3-, 6-, 9- and 12-months maturities. Panel 2 reports those of 3-, 6-, 9- and 12-month-ahead spot price forecasts recursively provided by the model forecasts. For each horizon, columns 2 and 6 report recursive MSPE and MAE ratios relative to those from no-change forecasts. While boldface indicates improvements relative to the no-change forecast, adjacent columns in parentheses report p -values from tests for equal forecasting accuracy suggested by Diebold and Mariano (1995) and West (1996). Given abnormal financial market behaviors after the Great Recession, columns 4 and 8 provide recursive MSPE and MAE ratios calculated after excluding forecasts for 5 months from 2008.11 to 2009.3.

Appendix A: Detailed Explanations

A.1. Equilibrium Prediction and Aggregate Inventory Demands

The equilibrium prediction for the aggregate inventory demands is obtained by combining the representative producer's optimal decisions with his forecasts for the future crude oil spot price. Recall the optimality conditions from the model

$$e^{z_t} Q_t = \left(\alpha e^{z_t} \left(1 - \exp \left(-\theta_t \frac{I_{t-1}}{N_{t-1}} \right) \right) \right)^{\frac{1}{1-\alpha}} (S_t)^{-\frac{1}{1-\alpha}} \cdot N_t,$$

$$e^{-r_{t+1}} E_t^Q \left[(e^{z_{t+1}} Q_{t+1})^\alpha (N_{t+1})^{1-\alpha} \frac{\theta_{t+1}}{N_t} \exp \left(-\theta_{t+1} \frac{I_t}{N_t} \right) \right] = S_t \left(1 + c_1 \frac{I_t}{N_t} \right) - e^{-r_{t+1}} E_t^Q [S_{t+1}],$$

where we rearrange the equation for aggregate resource demands in (6).

After moving the time subscript one period forward for the rearranged aggregate resource demands, we plug the resulting equation into aggregate inventory demands whose left-hand side becomes:

$$L(I_t, N_t, r_{t+1}) \cdot E_t^Q \left[(e^{z_{t+1}})^{\frac{\alpha}{1-\alpha}} (S_{t+1})^{-\frac{\alpha}{1-\alpha}} \right] = S_t \left(1 + c_1 \frac{I_t}{N_t} \right) - e^{-r_{t+1}} E_t^Q [S_{t+1}],$$

where $L(I_t, N_t, r_{t+1}) \equiv e^{-r_{t+1}} \alpha^{\frac{\alpha}{1-\alpha}} \theta_{t+1} (1+g) \left(1 - \exp \left(-\theta_{t+1} \frac{I_t}{N_t} \right) \right)^{\frac{\alpha}{1-\alpha}} \exp \left(-\theta_{t+1} \frac{I_t}{N_t} \right)$ does not depend on the representative producer's forecasts on the future technology process z_{t+1} and the future spot price S_{t+1} . Assuming the independence of the technology process and the log-normal property such as $E_t^Q \left[(e^{z_{t+1}})^\phi \right] = \exp \left(\phi z_t + \frac{1}{2} \phi^2 \sigma_1^2 \right)$ for $\phi \equiv \frac{\alpha}{1-\alpha}$ under the risk-neutral measure, the equilibrium storage equation is further reduced to:

$$\begin{aligned} & L(I_t, N_t, r_{t+1}) \cdot \exp \left\{ -\phi \left(-z_t + \xi_t + k^Q \chi_t - \lambda_\xi - \lambda_\chi \right) + \frac{1}{2} \phi^2 \left(\sigma_1^2 + s_1^Q \right) - r_{t+1} \right\} \\ = & S_t \left[1 + \frac{c_1 I_t}{N_t} - \exp \left\{ -\left(1 - k^Q \right) \chi_t - \lambda_\xi - \lambda_\chi + \frac{1}{2} s_1^Q - r_{t+1} \right\} \right]. \end{aligned}$$

After taking the log of both sides, the measurement equation of the aggregate inventory demands is obtained by approximating non-linear terms in the resulting equation. Linear approximations are taken where χ_t and I_t are close to $\frac{-\lambda_\chi}{1-k^Q}$ and I_{t-1} , respectively, since χ_t is mean-reverting around $\frac{-\lambda_\chi}{1-k^Q}$ and I_t is close to I_{t-1} for most of the sample. Denoting by $x_t \equiv (r_{t+1}, I_{t-1})$ a vector of observations that are either exogenous or predetermined at period t , we have:

$$\log(I_t/I_{t-1}) = \Phi_1(x_t) \cdot z_t + \Phi_2(x_t) \cdot \xi_t + \Phi_3(x_t) \cdot \chi_t + \Phi_4(x_t),$$

where $\Phi_1(x_t)$, $\Phi_2(x_t)$, $\Phi_3(x_t)$ and $\Phi_4(x_t)$ are as specified as follows:

$$\begin{aligned}\Phi_0(x_t) &= \left[\frac{I_{t-1}}{N_t} \left\{ \frac{\phi \theta_{t+1} \exp\left(-\frac{\theta_{t+1} I_{t-1}}{N_t}\right)}{1 - \exp\left(-\frac{\theta_{t+1} I_{t-1}}{N_t}\right)} - \theta_{t+1} - \frac{c_1}{1 + \frac{c_1 I_{t-1}}{N_t} - \exp\left(-\lambda_\xi + \frac{1}{2} s_1^Q - r_{t+1}\right)} \right\} \right]^{-1} \\ \Phi_1(x_t) &= \Phi_0(x_t) \times (-\phi) \\ \Phi_2(x_t) &= \Phi_0(x_t) \times (1 + \phi) \\ \Phi_3(x_t) &= \Phi_0(x_t) \times \left\{ 1 + \phi k^Q + \frac{(1 - k^Q) \exp\left(-\lambda_\xi + \frac{1}{2} s_1^Q - r_{t+1}\right)}{1 + \frac{c_1 I_{t-1}}{N_t} - \exp\left(-\lambda_\xi + \frac{1}{2} s_1^Q - r_{t+1}\right)} \right\} \\ \Phi_4(x_t) &= \Phi_0(x_t) \times \left[\frac{\theta_{t+1} I_{t-1}}{N_t} - \phi \ln \left\{ 1 - \exp\left(-\frac{\theta_{t+1} I_{t-1}}{N_t}\right) \right\} - \ln \left\{ \alpha^\phi \theta_{t+1} (1 + g) \right\} \right. \\ &\quad \left. + \ln \left\{ 1 + \frac{c_1 I_{t-1}}{N_t} - \exp\left(-\lambda_\xi + \frac{1}{2} s_1^Q - r_{t+1}\right) \right\} + \frac{\lambda_\chi \exp\left(-\lambda_\xi + \frac{1}{2} s_1^Q - r_{t+1}\right)}{1 + \frac{c_1 I_{t-1}}{N_t} - \exp\left(-\lambda_\xi + \frac{1}{2} s_1^Q - r_{t+1}\right)} \right. \\ &\quad \left. - \phi (\lambda_\xi + \lambda_\chi) - \frac{1}{2} \phi^2 (s_1^Q + \sigma_1^2) + r_{t+1} \right].\end{aligned}$$

A.2. State-Space Representation and Kalman Filter

Given the set of observations that are linear functions of state variables, I estimate associated parameters by maximum likelihood using the Kalman filter. First, I represent the system of equations in the state-space form. Let $\zeta_t = (z_t, \xi_t, \chi_t)'$ denote a 3×1 vector of unobservable state variables and y_t denote a 14×1 vector of observed variables at each period t . Then y_t can be described in terms of a linear function of unobservable state vector ζ_t . Denoting by x_t a vector of observations that are either exogenous or predetermined, the state-space representation is given by the following system of equations:

$$\xi_{t+1} = C + G\xi_t + v_{t+1}, \quad (\text{A1})$$

$$y_t = a(x_t) + H(x_t)' \xi_t + w_t, \quad (\text{A2})$$

where $a(x_t)$ and $H(x_t)$ in the observation equation denote a 14×1 vector valued function and a 14×2 matrix valued function, which are specified below. In the transition equation, C and G denote respectively a 3×1 coefficient vector and a 3×3 coefficient matrix; that is, $C = (0, -\lambda_\xi, -\lambda_\chi)'$ and $G = \text{diag}(1, 1, k^Q)$ with diag being a diagonal matrix operator whose diagonal elements are provided in the parentheses

Let \mathbf{y}_{t-1} denote data observed through date $t - 1$; i.e., $\mathbf{y}_{t-1} \equiv (y'_{t-1}, y'_{t-2}, \dots, y'_1, x'_{t-1}, \dots, x'_1)'$.

Assuming that conditional on x_t and \mathbf{y}_{t-1} , the vector $(v'_{t+1}, w'_t)'$ has the normal distribution

$$\begin{bmatrix} v_{t+1} \\ w_t \end{bmatrix} | x_t, \mathbf{y}_{t-1} \sim N \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} Q & 0 \\ 0 & R \end{bmatrix} \right), \quad (\text{A3})$$

where $Q(i, i) = \sigma_i^2$ for $i = 1, 2, 3$ and $Q(2, 3) = Q(3, 2) = \rho^Q \sigma_2 \sigma_3$ with $\sigma_1, \sigma_2, \sigma_3 > 0$ and $\rho^Q \in (-1, 1)$ are unknown parameters in state variables. $R = \text{diag}(\sigma_Q^2, \sigma_I^2, \sigma_{F,1}^2, \dots, \sigma_{F,12}^2)$ is a 14×14 measurement error matrix. Here $\sigma_I^2, \sigma_Q^2 > 0$ are measurement errors for aggregate demands for inventories and resources, and $\sigma_{F,\tau}^2$ is the measurement error for the price of the crude oil futures contract with the τ -periods until maturity for $\tau = 1, \dots, 12$.

We are left with specifying one unknown matrix $H(x_t)'$ and one unknown vector $a(x_t)$, yet restrictions are obtained from three sets of observation equations derived earlier. Recall these observation equations following (8), (9) and (10). With these, we can specify $H(x_t)$ and $a(x_t)$ following

$$H(x_t)' = \begin{bmatrix} \phi & -1 - \phi & -1 - \phi \\ \Phi_1(x_t) & \Phi_2(x_t) & \Phi_3(x_t) \\ 0 & 1 & \Omega_{1,1}(x_t) \\ \vdots & \vdots & \vdots \\ 0 & 1 & \Omega_{1,12}(x_t) \end{bmatrix} \quad a(x_t) = \begin{bmatrix} \Psi(x_t) \\ \Phi_4(x_t) \\ \Omega_{2,1}(x_t) \\ \vdots \\ \Omega_{2,12}(x_t) \end{bmatrix},$$

which completes the state-space representation.

Next, consider the calculation of the conditional likelihood using the Kalman filter for the maximum likelihood estimation. Suppose $\xi_t | \mathbf{y}_{t-1} \sim N(\widehat{\xi}_{t|t-1}, P_{t|t-1})$. Since x_t contains only strictly exogenous variables or lagged values of \mathbf{y} , we also have $\xi_t | \mathbf{y}_{t-1}, x_t \sim N(\widehat{\xi}_{t|t-1}, P_{t|t-1})$.

With (A1), (A2) and (A3), it can be shown that

$$\begin{bmatrix} \xi_t \\ y_t \end{bmatrix} | \mathbf{y}_{t-1}, x_t \sim N \left(\begin{bmatrix} \widehat{\xi}_{t|t-1} \\ a(x_t) + [H(x_t)]' \widehat{\xi}_{t|t-1} \end{bmatrix}, \begin{bmatrix} P_{t|t-1} & P_{t|t-1} H(x_t) \\ H(x_t)' P_{t|t-1} & H(x_t)' P_{t|t-1} H(x_t) + R \end{bmatrix} \right),$$

where $a(x_t)$ and $H(x_t)$ are treated as deterministic conditional on x_t .

Since they are jointly normally distributed, the conditional distribution of ξ_t given $\mathbf{y}_{t-1}, x_t, y_t$ becomes $\xi_t | \mathbf{y}_{t-1}, x_t, y_t \sim N(\widehat{\xi}_{t|t}, P_{t|t})$ where we used the formula for updating a linear projection:

$$\widehat{\xi}_{t|t} = \widehat{\xi}_{t|t-1} + P_{t|t-1} H(x_t) [H(x_t)' P_{t|t-1} H(x_t) + R]^{-1} \times [y_t - a(x_t) - [H(x_t)]' \widehat{\xi}_{t|t-1}] \quad (\text{A4})$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1} H(x_t) [H(x_t)' P_{t|t-1} H(x_t) + R]^{-1} H(x_t)' P_{t|t-1}. \quad (\text{A5})$$

One can calculate sample likelihood using the Kalman filter as derived earlier given $\xi_{t+1}|\mathbf{y}_t \sim N\left(\widehat{\xi}_{t+1|t}, P_{t+1|t}\right)$ with

$$\widehat{\xi}_{t+1|t} = C + G\widehat{\xi}_{t|t} \quad (\text{A6})$$

$$P_{t+1|t} = GP_{t|t}G' + Q \quad (\text{A7})$$

Assume that the initial state ξ_1 is distributed as $N\left(\widehat{\xi}_{1|0}, P_{1|0}\right)$. Given observations $\{y_t, x_t\}$ for $t = 1, \dots, T$, the distribution of y_t conditional on $(y_{t-1}, \dots, y_1, x_t, x_{t-1}, \dots, x_1)$ is normal with mean $\left(a(x_t) + H(x_t)' \widehat{\xi}_{t|t-1}\right)$ and variance $\left(H(x_t)' P_{t|t-1} H(x_t) + R\right)$; that is, for $t = 1, 2, \dots, T$, we have the conditional likelihood following

$$\begin{aligned} f_{Y_t|X_t, Y_{t-1}}(y_t|x_t, \mathbf{y}_{t-1}) &= (2\pi)^{-n/2} |H(x_t)' P_{t|t-1} H(x_t) + R|^{-1/2} \\ &\times \exp\left[-\frac{1}{2} \left(y_t - a(x_t) - H(x_t)' \widehat{\xi}_{t|t-1}\right)'\right. \\ &\left. \times \left(H(x_t)' P_{t|t-1} H(x_t) + R\right)^{-1} \left(y_t - a(x_t) - H(x_t)' \widehat{\xi}_{t|t-1}\right)\right]. \end{aligned}$$

With this, the sample log likelihood becomes

$$\mathcal{L}(\Theta) = \sum_{t=1}^T \log f_{Y_t|X_t, Y_{t-1}}(y_t|x_t, \mathbf{y}_{t-1}),$$

where the unknown parameters are estimated numerically.