DALLASFED

Occasional Paper

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

Sung Je Byun

Federal Reserve Bank of Dallas
Financial Industry Studies Department

Occasional Paper 16-03
September 2016

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

Sung Je Byun *

September, 2016

Abstract

This paper examines the role of inventories in refiners' gasoline production and develops a structural model of the relationship between crude oil prices and inventories. Using data on inventories and prices of oil futures, I show that convenience yields decrease at a diminishing rate as inventories increase, consistent with the theory of storage. In addition to exhibiting seasonal and procyclical behaviors, I show that the historical convenience yield averages about 18 percent 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 financial investors' participation after 2004 was evidence of an effect of speculation, I find that the proposed price-inventory relationship is stable over time. The empirical evidence indicates that crude oil prices remained tied to oil-market fundamentals such as inventories, suggesting that the contribution of financial investors' activities was weak.

*Research Economist in Financial Industry Studies Department, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX, USA. E-mail:SungJe.Byun@dal.frb.org. 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.

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

1 Introduction

Recent volatility in 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 relationship between crude oil prices and inventories.

In a competitive commodity market, a producer makes the optimal storage decision by equating its expected benefit with the relevant cost of holding inventories. A positive expected 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 traditional theory of storage literature (Kaldor, 1939; Working, 1949; Brennan, 1958; Telser, 1958). For example, Brennan (1958) defines this benefit as the convenience yield, which is the flow of services that accrues to an owner of a physical commodity, but not to an owner of a derivative contract for future delivery of that commodity.

In this paper, I propose an equilibrium storage model of the global crude oil market to explain fluctuations of crude oil prices in terms of risk-averse refiners' benefits to hold crude oil inventories. I focus on the role of crude oil inventories since the refiners' gasoline production depends heavily on this raw materials inventory, which generally has a lengthier delivery time relative to other factors of production. Building upon earlier research by Eckstein and Eichenbaum (1985) and Considine (1997), I model the role of crude oil inventories directly by treating inventories as an essential factor of production (Kydland and Prescott, 1982; Kydland and Prescott, 1988; Ramey, 1989). Unlike earlier inventory studies employing linear quadratic models to study the role of finished goods inventory (e.g., Pindyck, 1994), I introduce a production function of the form that captures nonlinear effects of crude oil inventories on refiners' gasoline production. Hence, the crude oil inventories improve refiners' gasoline production efficiency without needing to make costly production adjustments (Brennan, 1958). The proposed storage model provides a nonlinear expression of the marginal convenience yield (henceforth, convenience yield) as a function of crude oil inventories. While the convenience yield could arise with risk-neutral refiners, I assume refiners are risk averse because risk premiums, the costs refiners pay to hedge their activities in storage markets, are channels through which refiners interact with financial investors (Hamilton and Wu, 2014). In such a case, the financial investors' activities could be reflected through risk premiums, although they generally do not carry crude oil inventories. In the proposed model, a shift in the convenience yield results from a change in refiners' expectations on future crude oil prices and associated changes in risk premiums. The equilibrium storage demand curve, defined as the difference between the convenience yield and the marginal storage cost, can be viewed as a mirror image of the storage supply curve in Working (1949), except that the risk-premiums affect the convenience yields as well as the marginal storage costs in the proposed model.

Using data on crude oil inventories and futures contract prices, I provide empirical evidence that the proposed storage model is consistent with the theory of storage, helping to explain fluctuations of crude oil prices in terms of the convenience yield that accrues to refiners only. First, the convenience yield is a convex and decreasing function of inventories, indicating that it decreases at a diminishing rate as the inventory increases. Second, I find that the convenience yield accounts for about 18 percent of the oil price on average from March 1989 to November 2014, holding other factors of production constant. Third, I find procyclical and seasonal behaviors in the convenience yield. While increasing during the expansionary period, the convenience yields for spring and summer are higher 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), Büyükşahin and Robe (2014) and Hamilton and Wu (2014). Lastly, I show that the proposed model provides more accurate one-, three-, six- and 12-month-ahead crude oil spot price forecasts than a random walk model, supporting a stable price-inventory relationship despite the structural break. 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 since it focuses 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) using different approaches (Fattouh et al., 2013; Kilian and Murphy, 2014; Kilian and Lee, 2014; Knittel and Pindyck, 2016; 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 refinery's expected benefits to hold crude oil inventories. After providing empirical predictions from the theory of storage, I illustrate the refinery's equilibrium storage decision, which is subject to adjustments as the refinery updates its forecasts of oil market conditions. In section 3, 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. 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 a refinery's gasoline production decisions. Given a large share of wholesale gasoline prices accrued to costs for crude oil resources, I abstract a refinery's decision for gasoline production from its complicated production decisions for multiple intermediary goods. I study the optimal storage decision of a refinery who has an incentive to hold inventories despite anticipated falling crude oil prices. Note that a refinery's first-order conditions 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 risk-averse 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 (2016). 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.

2.1 Cost Function

Consider a refinery 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:

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

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 refinery (q_t^U) is smaller than the quantity purchased (q_t^P), inventories accumulate and vice versa.

The cost function summarizes both the production and storage technology for the refinery. Given the current crude oil spot price (S_t), the total costs come from two components: costs for purchasing resources and costs 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.[1] 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 refinery's cost function becomes

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

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), Brennan (1958), and Telser (1958). For example, 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), Thurman (1988), Deaton and Laroque (1992), Ng and Pirrong (1994), Routledge et al. (2000), Bryant et al.

[1] The storage costs are defined in a broad sense; they include costs for insurance and transportation beside those necessary for using storage facilities. In general, costs for insurance and transportation are small, yet are proportional to the stored value.

(2006), Pirrong (2011), Gorton et al. (2013), Acharya et al. (2013), and Alquist et al. (2014).

The role of inventories as a factor of production has been widely adopted in earlier literature. 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. It is because a refinery's gasoline production heavily depends on the raw material such as crude oil, yet faces a long delivery time in general.[2] In response to uncertainties in gasoline demand and crude oil 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 even if it anticipates falling crude oil prices. Given the realized technology process (z_t), I propose a production function of the form,

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

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 possiblity 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 the second 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

[2]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 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.

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 $\epsilon_{1,t+1} \sim N\left(0, \sigma_1^2\right)$.

2.3 Optimal Storage Decision Problem

With the cost function and the production function as introduced earlier, the refinery faces a dynamic programming problem. At the beginning of each period t, the refinery 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}), the refinery makes decisions on the resource demand (q_t^U) and inventories (i_t) jointly to maximize its lifetime profits. Suppose it is a price taker in the crude oil market and is risk averse with discount factor (Λ_t),[3] which governs its risk aversion. Suppressing the superscript in the resource demand (q_t), the refinery's objective is, thus, to choose $\{q_t, i_t\}_{t=1}^{\infty}$ so as to maximize

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

where i_{-1} is given.

The optimality conditions for the refinery are summarized as:

$$\alpha e^{\alpha z_t} q_t^{\alpha-1} \left(1 - e^{-\theta_t i_{t-1}}\right) = S_t, \tag{2}$$

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

[3] In a standard macro model, we have $\Lambda_\tau = \frac{\beta \cdot U'\left(c_{\tau+1}\right)}{U'\left(c_\tau\right)} e^{r_\tau}$, where $U(\cdot)$ represents the agent's utility function, β is the subjective discount factor and r_τ is the interest rate. The role of crude oil inventories such as the convenience yield could arise in the storage model under a risk-neutrality assumption ($\Lambda_\tau = 1$), which is widely adopted in most empirical storage literature for analytic tractability. However, the risk-averseness assumption is essential for the purpose of this paper since risk premiums paid by refiners for their hedging activities are mediums through which financial investors' activities could be reflected. In the online appendix (available at https://sites.google.com/site/sungjebyun/), I present parameter estimates of the proposed storage model for risk-neutral refineries, of which results are comparable to those in Table 1.

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 products equal the marginal costs. 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. Hence, the refinery has an incentive to hold a positive level of inventory despite an expected capital loss in the future.

2.4 Equilibrium Predictions

The theory of storage predicts decreasing and convex convenience yield; that is, the convenience yield decreases at a lessening rate as inventory increases. Assuming that the refinery internalizes its future resource demand in its storage decision, I rewrite the left-hand side in (3) by plugging the expected resource demand (q_{t+1}) obtained from (2) into the storage decision. Holding other variables including the refinery's spot price forecasts constant, the optimality condition for the storage decision becomes

$$e^{-\theta_{t+1} i_t} \cdot \left(1 - e^{-\theta_{t+1} i_t}\right)^{\phi} \cdot \Phi_{0,t+1} = S_t \left(1 + c_1 i_t\right) - E_t \left[\Lambda_{t+1} e^{-r_{t+1}} S_{t+1}\right], \tag{3'}$$

where $\phi = \frac{\alpha}{1-\alpha}$ and $\Phi_{0,t+1} = \alpha^{\phi} \cdot E_t \left[\Lambda_{t+1} \theta_{t+1} e^{-r_{t+1} + \phi z_{t+1}} S_{t+1}^{-\phi}\right]$.

Let MCY_t be the left-hand side of the equation (3'). The theory of storage provides two equilibrium predictions;

$$\frac{\partial MCY_t}{\partial i_t} = \left\{(1 + \phi) \cdot e^{-\theta_{t+1} i_t} - 1\right\} \cdot \Phi_{1,t+1} < 0, \tag{4}$$

indicating the convenience yield is inversely related to inventories if the term in curly brackets is negative, provided by $\Phi_{1,t+1} = \Phi_{0,t+1} \cdot \theta_{t+1} e^{-\theta_{t+1} i_t} \left(1 - e^{-\theta_{t+1} i_t}\right)^{\phi - 1} > 0$. And,

$$\frac{\partial^2 MCY_t}{\partial i_t^2} = \left\{ \left(\phi - e^{\theta_{t+1} i_t}\right) \cdot \left(\frac{\phi e^{-\theta_{t+1} i_t}}{1 - e^{-\theta_{t+1} i_t}} - 1\right) - (1 + \phi) \right\} \cdot \Phi_{2,t+1} > 0, \tag{5}$$

indicating the convenience yield is a convex function of inventories if the term in curly brackets is positive, provided by $\Phi_{2,t+1} = \Phi_{0,t+1} \cdot \theta_{t+1}^2 e^{-2\theta_{t+1} i_t} \left(1 - e^{-\theta_{t+1} i_t}\right)^{\phi - 1} > 0$. Estimates for parameters such as α, θ_{t+1}, which will be in Table 1 of section 3.2, are consistent with the above two equilibrium predictions for all sample periods, providing empirical evidence supporting the theory of storage.

The equilibrium storage decision provided by (3') differs from the earlier theory of storage literature in two ways. First, it describes storage demand for raw material, in this case, crude oil, whereas most earlier literature except Alquist and Kilian (2010) studies the storage supply of finished goods. Hence, the storage demand curve, defined as the difference between the left-hand side and the right-hand side of the equation (3'), is a convex and decreasing function of inventories; this is a mirror image of the storage supply curve in Working (1949). More importantly, the equilibrium storage decision reflects risk premiums that the risk-averse refinery pays for its hedging activities, for which costs vary depending on the state of the economy and the level of the crude oil inventory. Here the risk premiums provide a channel from which the refinery adjusts its storage decision in response to changes in the financial market's belief on economic conditions (Brennan 1958).

Figure 1 illustrates the refinery's equilibrium storage decision and its dynamic adjustment process in response to anticipated falling crude oil prices. For illustrative purposes, I calibrate the convenience yields and the marginal costs using parameter estimates in Table 1 and estimates of state variables, which are going to be discussed in the following section. Two solid lines represent the convenience yield curve (MCY_1) and marginal cost curve (MSC_1) for holding an additional barrel of crude oil inventory in January 2003, during which prices of West Texas Intermediate (WTI) crude oil futures contracts for all maturities were below crude oil spot prices ("normal backwardation").[4] In January 2003, the equilibrium storage demand (I_1) is determined at the intersection of the two solid lines. Given falling oil prices anticipated from prices of crude oil futures contracts in February 2003, the refinery's marginal cost curve (MSC_2) shifts upward since stored crude oil becomes less valuable in the future. Holding the convenience yield curve constant, the refinery's storage demand would decrease because of insufficient expected benefits from holding inventories. Given the role of crude oil inventories in enhancing the refinery's future gasoline production, the convenience yield curve (MCY_2) also shifts upward in February 2003, of which the expected benefits dominate effects from the increased marginal costs. Hence, the equilibrium storage demand increases to I_2, although prices of crude oil futures contracts in February 2003 remain to indicate falling crude oil prices as they appeared in January 2003.

2.5 Aggregation, Seasonality and the Crude Oil Spot Price

The theory presented so far applies to one refinery; however, available observations are aggregated at the global level.[5] Given the trend growth observed from the global crude oil production, consumption and inventories, I assume that the number of refineries increases over time at a fixed rate. Let N_t be the number of refineries in period t and g be the exogenously determined average

[4]In the online appendix, I provide supplemental figures to highlight the proposed storage model's flexibility in explaining oil market conditions, which are observed from crude oil futures spreads. I also compare equilibrium storage decisions and dynamic adjustment processes during a few selected historical episodes under the risk-neutrality assumption with those under the risk-averseness assumption, which highlight the relevance of the risk-averseness assumption adopted in the proposed storage model.

[5]Brennan et al. (1997) argues that the convenience yield is an illusionary phenomenon caused by spatial aggregation of stocks, provided by empirical evidence that no stocks are held at a monetary loss in terms of local prices. However, the local wheat marketing system of Western Australia studied by Brennan et al. is qualitatively different from the global crude oil market in this paper. While wheat is a perishable good and farmers do not have storage technologies, grain storage and rail transport in the local wheat market have operated as monopolies. Furthermore, the local and regional supply of wheat each has small variation within the same crop year, leaving no incentive to hold wheat inventories besides monetary profits.

growth rate of the refinery, i.e., $N_t = (1+g)\, N_{t-1}$.[6] 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 consumption 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 refinery'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 (Ye et al., 2002). Specifically, I conjecture that seasonal variations exist in the refinery'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 and May, θ_2 for June, July and August, θ_3 for September, October and November, and θ_4 for December, January and 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 ($\log S_{t+1}$) consists of the long-term trend

[6] The Monthly Refinery Capacity Report of the U.S. Energy Information Administration (EIA) provides annual estimates for operable atmospheric crude oil distillation capacity in U.S. since 1982. This series exhibits linear trend growth, which is consistent with the assumption for dealing with the linear trend observed in the aggregate data.

(ξ_{t+1}) and the short-term deviation (χ_{t+1}); that is $\ln S_{t+1} \equiv \xi_{t+1} + \chi_{t+1}$.[7] 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, $\xi_{t+1} = \xi_t + \epsilon_{2,t+1}$. On the other hand, the short-term deviation from the long-term trend follows the mean-reverting AR(1) process, $\chi_{t+1} = k \cdot \chi_t + \epsilon_{3,t+1}$, 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 refinery's forecasting problem in the storage decision (3) as if it is risk-neutral.[8] 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 that the refinery pays for its hedging activities in the crude oil market, this form of risk adjustment is frequently adopted in the literature (Schwartz and Smith, 2000; Duffee, 2002; Casassus and Collin-Dufresne, 2005; Pirrong, 2011; Hamilton and Wu, 2014). Denoting by $\Upsilon \equiv \left[\widetilde{\lambda}_z, \widetilde{\lambda}_\xi, \widetilde{\lambda}_\chi\right]'$ 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: $\widetilde{\lambda}_z \equiv 0$, $\widetilde{\lambda}_\xi = \lambda_\xi$, $\widetilde{\lambda}_\chi \equiv$

[7] Another 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 prices of crude oil futures contracts, in which risk premiums vary depending on the level of crude oil inventory.

[8] Implicit assumptions are the deterministic interest rate and the redundancy of the futures contract. The former guarantees the price equivalance between futures and forward contracts, whereas small differences are found in the empirical literature (Chow et al., 2000). 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).

$\lambda_\chi + \varpi\chi_t$ with ϖ being constants. Hence, risk-neutral stochastic processes are of the form:

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

where $k^Q \equiv k - \varpi \in (-1, 1)$. Assuming that $\epsilon^Q_{1,t+1}$ is independent of both $\epsilon^Q_{2,t+1}$ and $\epsilon^Q_{3,t+1}$, innovations for risk-neutral stochastic processes ($\epsilon^Q_{1,t+1}, \epsilon^Q_{2,t+1}, \epsilon^Q_{3,t+1}$) 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 as[9]

$$\log E^Q_t [S_{t+1}] = \xi_t + k^Q\chi_t - \lambda_\xi - \lambda_\chi + \frac{1}{2}s^Q_1,$$

where $s^Q_1 \equiv Var^Q_t [\log S_{t+1}] = \sigma^2_2 + \sigma^2_3 + 2\rho\sigma_2\sigma_3$. 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

$$\log E^Q_t [S_{t+\tau}] = \xi_t + \left(k^Q\right)^\tau \cdot \chi_t - \tau \cdot \lambda_\xi - \frac{1 - \left(k^Q\right)^\tau}{1 - k^Q} \cdot \lambda_\chi + \frac{1}{2}s^Q_\tau, \tag{6}$$

where $s^Q_\tau \equiv \tau \cdot \sigma^2_2 + \frac{1-\left(k^Q\right)^{2\tau}}{1-\left(k^Q\right)^2} \cdot \sigma^2_3 + 2\rho\sigma_2\sigma_3 \cdot \frac{1-\left(k^Q\right)^\tau}{1-k^Q}$ is the conditional variance associated with τ-periods-ahead forecast of the crude oil spot price.

Note that I use the closed-form spot price forecast provided in (6) for evaluating prices of

[9]The conditional expectation and the variance of $\log S_{t+1}$ under the risk-neutral processes are respectively given as, $E^Q_t [\log S_{t+1}] = \xi_t + k^Q\chi_t - \lambda_\xi - \lambda_\chi$ and $Var^Q_t [\log S_{t+1}] = \sigma^2_2 + \sigma^2_3 + 2\rho\sigma_2\sigma_3$. Hence, one-month-ahead crude oil spot price forecast is provided by $logE^Q_t [S_{t+1}] = E^Q_t [\log S_{t+1}] + \frac{1}{2}Var^Q_t [\log S_{t+1}]$ due to Jensen inequality. For arbitrary $\phi > 0$, it can be further generalized as $logE^Q_t \left[(S_{t+1})^{-\phi}\right] = -\phi \cdot \left(\xi_t + k^Q \cdot \chi_t - \lambda_\xi - \delta\right) + \frac{1}{2}\phi^2 s^Q_1$.

crude oil futures contracts with various maturities. 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 contracts with τ periods to maturity ($F_{t,\tau}$) is the τ-period-ahead spot price forecast under the risk-neutral measure; that is, $F_{t,\tau} = E_t^Q [S_{t+\tau}]$.

3 Empirical Results

3.1 Data and Empirical Implementation

The data set includes monthly observations from March 1989 to November 2014, where the beginning period is determined by the availability of prices for a 12-month crude oil futures contract. The monthly series of global crude oil inventories, consumption and real prices for crude oil futures contracts are used for estimating the proposed storage model in section 2. 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). Using the constructed series of global crude oil inventories, I construct the monthly series of global crude oil consumption by subtracting monthly changes in global crude oil inventories from monthly global crude oil production. Given prices of crude oil futures contracts having maturities up to 12 months, I calculate real prices of those futures contracts using the U.S. Consumer Price Index (2010.5=100). Lastly, I calculate the monthly real risk-free interest rate by adjusting the U.S. LIBOR with the realized inflation observed in the U.S. Consumer Price Index. The appendix provides a detailed discussion of the data set construction procedure as well as the data source.

Assuming an inelastic supply of crude oil following Kilian (2008) and Alquist and Kilian

(2010), I estimate model parameters along the aggregate resource and storage demands. While treating crude oil spot price as being unobservable, I use prices of crude oil futures contracts having maturities up to 12 months to estimate unobservable components in the crude oil spot price following Schwartz and Smith (2000). During the estimation procedure, I allow for measurement errors in the observables such as crude oil consumption, inventories and prices for crude oil futures contracts. This addresses the widely recognized measurement issues in oil quantity variables (Baumeister and Kilian, 2012) and mispriced oil futures contracts (Kilian and Murphy, 2014). The online appendix provides detailed explanations for the estimation procedure, which includes derivations of three sets of observation equations obtained from the equilibrium conditions (2), (3') and (6), and a construction of the log likelihood using a Kalman filter.

3.2 Results

The first column in Table 1 reports maximum likelihood estimates and asymptotic standard errors. Estimates of the model parameters are consistent with historical observations and the empirical predictions from the theory of storage. In regard to the latter observation, I find that the terms in the curly brackets in equations (4) and (5) are consistent with the theoretical restrictions during the entire sample period, confirming that the convenience yield is a decreasing and convex function of the crude oil inventory. The online appendix provides a supplemental figure that plots estimates for two theoretical restrictions during the sample period. As for the estimates themselves, several points stand out. First, the output share (α) estimate indicates that crude oil resources account for approximately 39 percent of gasoline production globally, close to the estimate obtained from U.S. EIA methodology.[10] Second, when the real price of crude oil was \$99.21 per barrel and the inventory level was $2,455.15$ million barrels per day

[10]Following the methodology introduced in weekly "Gasoline and Diesel Fuel Updates", I calculate shares of crude oil resource costs in a gallon of gasoline price by using refiners' acquisition cost for imported crude oil and regular conventional gasoline price provided by U.S. EIA. The costs for crude oil resource accounts for approximately 49% of the U.S. retail gasoline price on average during the total sample period.

in January 2008, estimates for storage cost functions indicate that costs for storage facility necessary for increasing an additional barrel of crude oil inventory are approximately \$18.39 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. Since the smaller θ implies the refinery's higher dependence on available inventories, these estimates indicate that utilizations for spring and summer are higher than those for fall and winter (Milonas and Henker, 2001). While all seasonal utilization parameters ($\theta_1, \theta_2, \theta_3, \theta_4$) are statistically significant, the smallest difference among utilization parameters is 0.0010 between fall (θ_3) and winter (θ_4).[11] When crude oil inventories correspond to the median level of $2,206.16$ millions barrels per day during February 2003, a 0.0010 decrease in θ yields a \$0.13 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, refineries need to provide counterparties 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 high risk-premiums when providing liquidities for hedgers when facing the short-term risk factor.

Figure 2 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

[11] 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), fall (θ_3) and winter (θ_4). For a null hypothesis of $H_0 : \theta_1 = \theta_2 = \theta_3 = \theta_4$, Wald test statistic is 6.05 with the corresponding p-value being 0.11.

by overlaying the model forecasts from (6) with the real price of the nearest maturity crude oil futures contract 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 ($\frac{-\lambda_\chi}{1-k^Q} = 0.45$), of which historical observations are consistent with previous episodes documented in the earlier literature (Hamilton, 2009; Kilian, 2009; Kilian and Murphy, 2014).

Panel 1 in Figure 3 plots the equilibrium convenience yield as a percentage of the crude oil spot price.[12] The convenience yield fluctuates considerably along the business cycle, repeating the same patterns around the past recession and recovery periods (Gibson and Schwartz, 1990). When the global economy starts to recover from the early 2000s recession, for example, the convenience yield sharply increases from 20.76 percent in February 2002 to 22.94 percent in February 2003. On the other hand, it declines from 17.58 percent to 15.93 percent 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; Alquist and Kilian, 2010; Knittel and Pindyck, 2016). During the early 2000 recovery period, the crude oil inventories declined by 0.12 percent, which reflects increasing crude oil demands for immediate gasoline production as well as for future uses. On the other hand, crude oil inventories had increased by 3.87 percent during the Great Recession. The correlation coefficient estimates are -0.45 with the level of crude oil inventories and -0.21 with the monthly changes in crude oil

[12]The convenience yield is calculated by plugging estimates of parameters and state variables into the left-hand side of the refinery's optimal storage decision in equation (3'). The online appendix provides detailed explanations for the convenience yield calculation, including a closed-form convenience yield measure provided by the proposed storage model.

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 correlation coefficient being 0.57 (-0.45) over the sample period. When the convenience yield increased during the early 2000 recovery period, the short-term deviation increased by 91.67 percent, whereas the long-term trend increased by 2.86 percent. During the Great Recession, the short-term deviation had decreased by 33.46 percent, whereas the long-term component had decreased by 3.04 percent. Lastly, I find that the convenience yield averages about 18.24 percent of the oil price over the sample period.[13] Holding other factors of production constant, historical average gains in the refinery's productivity can be translated into a significant benefit such as 18.24 percent relative to the crude oil 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 quantifying the 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 proxied by U.S. refiners' acquisition cost for imported crude oil. In general, observations from negative spreads are analogous to those from the 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 spikes are larger for the spreads from longer-term

[13] The average convenience yield is similar in magnitude to the estimate in Considine (1997), who finds 22% of the average annual convenience yield in an earlier sample period from April 1981 to October 1994 by estimating a set of equations that includes optimal demands for inventories. Though interesting, two estimates are not directly comparable to each other mainly because of two sources of market incompleteness in the theoretical models. Considine studies the role of inventories in the the risk-neutral refinery's production of multiple petroleum intermediary goods, where the convenience yield arises from reducing adjustment costs for production, investment, and stocks of crude oil supplies and work in process.

futures contracts,[14] it can be explained by the hypothesized role of crude oil inventories, leaving a precautionary role unmodeled. Despite such a difference, one can visually confirm that the equilibrium convenience yield exhibits qualitatively similar patterns during historical events documented in the empirical literature, which infers net convenience yield (abstracting marginal storage costs) by using the term spread calculated from prices of crude oil futures contracts (Pindyck, 1994; Alquist and Kilian, 2010; Alquist et al., 2014; Knittel and Pindyck, 2016).

3.3 Stable Forecasting Relationship

Given increased participation of financial investors 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 nonenergy commodities, raising doubts about potential spillover effects on commodity markets from other financial markets. On the other hand, Fattouh (2010) shows that 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 proposed storage model under two subsample periods split after June 2004. Then, I show that out-of-sample forecasts are more accurate than a no-change forecast at horizons of three, six, nine and 12 months, highlighting that the price-inventory relationship is stable despite a structural change in the crude oil market.

I estimate the model again using data for two subsample periods. The first subsample spans from March 1989 to June 2004 (184 observations), and the second spans from July 2004 to November 2014 (125 observations). The second and third columns in Table 1 report maximum likelihood estimates and asymptotic standard errors. In general, estimates are qualitatively sim-

[14]The proposed convenience yield measures benefits expected from one month's storage of an additional barrel of crude oil. From this perspective, a one-month spread, defined as a (log) price difference between spot and the nearest maturity futures contract, provides a fair comparison to the convenience yield in Panel 1 of Figure 3. When calculating the one-month spread, I find that it is generally smaller in magnitude than the proposed convenience yield. Furthermore, it exhibits erroneous behaviors for a few periods, during which its behaviors are qualitatively different from the spreads in Panel 2.

ilar 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 for 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 subsample 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 percent and 14.98 percent, respectively, for the first and second periods, during which the crude oil inventory increased from $2,157.76$ million barrels to $2,480.25$ million on average. I find higher trend (g) and lower marginal storage cost (c_1) estimates for the second period that are also consistent with the observed fast increase in crude oil inventory during the second period. Here, the smaller output share (α) from the second period does not indicate a reduced share of the crude oil resource in gasoline production due to omitted factors in the refinery's gasoline production.[15] From parameter estimates describing the crude oil spot price, I find that the statistical properties for the crude oil spot price also changes. 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 in the crude oil market. This can be confirmed by the increased price of long-term risks (λ_ξ) during the latter period. In summary, the subsample exercise in this section provides empirical evidence of the structural change in the crude oil market, which is consistent with Tang and Xiong (2012), Büyükşahin and Robe (2014) and Hamilton and Wu (2014), among others.

[15] The average shares of crude oil resource costs calculated following the U.S. EIA methodology are 39 percent and 62 percent, respectively, for the first and second subsample periods. A potential explanation for a sizable difference during the second subsample period can be found from an omitted labor demand decision in the refinery's production. From Bureau of Labor Statistics estimates, I find changes in the labor productivity growth rate across the two subsample periods. While labor productivity in the refining industry increased 2.3 percent annually for the 1989-2014 period, a faster growth rate occurred for the 1989-2004 period (3.4 percent) than for the 2005-2014 period (0.5 percent).

Next, I perform an out-of-sample forecasting exercise to investigate whether the structural change affects the price-inventory relationship proposed in this paper. Throughout the forecasting exercise, I forecast the log real price of oil because the model specifies the relationship between the log real price and the state variables. To reduce computational burden during the out-of-sample forecasting exercise, I use a smaller set of observations consisting of futures contract prices with one, three, six and 12 months to maturity, oil consumption 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 an estimation window of 184 months, 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. To evaluate the performance, I set a benchmark as a random walk model forecast, for which a one-month-ahead forecast for the τ-month-to-maturity futures contract is the current price for the same futures contract. The benchmark forecasts for the τ-month-ahead spot price is the current crude oil spot price, which is proxied by U.S. refiners' acquisition cost for crude oil imports (Baumeister and Kilian, 2012). Notice that the storage model forecast nests the random walk forecast as a special case when $k^Q = 1$. However, empirical evidence indicates the contrary; estimates of k^Q are statistically different from 1 in sample for total and two subsample periods (Table 1), and out-of-sample estimates are also statistically different from 1 during the entire rolling fixed estimation window.[16] Given this empirical evidence, I compare their forecasting accuracies using tests for equal forecasting accuracy suggested by Diebold and Mariano (1995) and West (1996)

[16]The online appendix provides a series of k^Q estimates obtained during the out-of-sample forecasting exercise.

(henceforth, DMW).[17]

Table 2 reports mean squared prediction error (MSPE) and mean absolute error (MAE) ratios relative to the no-change forecasts. Panel 1 evaluates the performance of one-month-ahead price forecasts for crude oil futures contracts with one, three, six and 12 months to maturity. Compared with the benchmark, the proposed model provides less-accurate price forecasts for all futures contracts using MSPE criteria (column 1). Although a more accurate forecast for short-maturity futures contract is found from the MAE criteria (column 5), I also find that such improvement is not statistically significant based on the DMW tests for equal forecasting accuracy (column 6). Given abnormal behaviors in the financial markets during the Great Recession period (Gürkaynak and Wright 2012), I recalculate 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, three and six months to maturity under MSPE (column 3) and under MAE (column 7). 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 premiums in these contracts, as in Baumeister and Kilian (2014).

Panel 2 reports MSPE and MAE for one-, three-, six- and 12-month-ahead crude oil spot price forecasts relative to the no-change forecast for crude oil spot prices. Compared with the no-change forecast, I find that the equilibrium storage 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 percent level for three- and 12-month-ahead forecasts using MSPE (columns 1 and 2), and for three-, six- and 12-month-ahead forecasts using MAE (columns

[17] When comparing forecasting accuracies of two nested models, size distortions of the DMW test statistics are widely recognized in the forecasting literature (Clark and West, 2007; McCracken, 2007). Using asymptotically valid critical values provided in McCracken (2007), I find stronger statistical significance for out-of-sample forecasts than those reported in Table 2.

5 and 6). 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 the price-inventory relationship proposed in this paper is stable despite the structural change in the oil market, consistent with Kilian and Murphy (2014), who develop a four-variable structural VAR model. The improved accuracies in forecasting future crude oil spot prices are consistent with Baumeister and Kilian (2012), who document the improved accuracies from widely adopted forecasting models when using real-time data.[18]

4 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 the refiners expect to receive when storing a barrel of crude oil.

Using data on crude oil inventories, consumption and prices of crude oil futures contracts, I find empirical evidence on the convenience yield that is consistent with the theory of storage literature. I show that the proposed convenience yield decreases at a diminishing rate as the inventory increases. While finding seasonal and procyclical behaviors, I show that the historical convenience yield is about 18 percent 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

[18] The proposed storage model also provides more accurate out-of-sample spot price forecasts for all horizons using an expanding window method adopted in Baumeister and Kilian (2012). However, a widely adopted bootstrap method for their statistical significance is computationally burdensome given highly nonlinear structural relationships among oil market variables.

coincides with the period of rapidly increasing financial investor activity. I document empirical evidence of 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 spot prices than the random walk model. The empirical evidence indicates that crude oil prices remained tied to oil market fundamentals such as inventories, suggesting that the contribution of financial investor activity to the crude oil market after 2004 was weak.

Acknowledgements

I am grateful to Thomas Baranga, Christiane Baumeister, James Hamilton, Soojin Jo, Allan Timmermann, Alexis Toda and Valerie Ramey and five anonymous referees for their helpful comments and suggestions on earlier drafts of this paper. I thank the participants of the 2016 Society for Nonlinear Dynamics and Econometrics conference, 2016 Energy and Commodity Finance conference, 92nd annual Western Economic Association International conference, and the seminar series at the University of California, San Diego and University of California, Riverside.

Appendix: Data Construction Procedure and Data Source

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 Organization for Economic Cooperation and Development (OECD) petroleum stocks over U.S. petroleum stocks. Given highly correlated growth rates of U.S. and OECD petroleum inventories, this procedure provides reasonable estimates of global crude oil inventories. The monthly series of OECD petroleum stocks is provided by the International Energy Agency (IEA), and U.S. crude oil inventories and petroleum stocks are provided by the U.S. Energy Information Administration (EIA).

With the inventory dynamics introduced in 2.1, I construct the monthly series on global crude oil consumption by subtracting monthly changes in global crude oil inventories from monthly global crude oil production provided by the U.S. EIA. To check the validity of the constructed series on global crude oil consumption, I calculate annual estimates of global crude oil consumption by scaling annual estimates of U.S. crude oil production by the ratio of annual global petroleum consumption estimates over annual petroleum production estimates. The annual estimates of U.S. crude oil production are provided by the U.S. EIA, and those for global petroleum consumption and production are provided by the IEA. I confirm that the (annual) estimated growth rates of consumption and those calculated from the constructed monthly series on global crude oil consumption are highly correlated, with a correlation coefficient of 0.84 over an extended sample period from 1981 to 2013.

Lastly, I construct monthly prices for crude oil futures contracts from daily prices for the light sweet crude oil futures contracts (a.k.a. WTI crude oil futures contracts) obtained from Datastream. Given daily prices for futures contracts having maturities up to 12 months, I construct monthly prices of a given maturity using the first calendar date price of corresponding futures contracts. This is 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. Given prices of crude oil futures contracts having maturities up to 12 months, I calculate real prices of those futures contracts using the U.S. Consumer Price Index (2010.5=100, seasonally unadjusted) provided by the OECD. While prices of Brent futures contracts are more relevant than those of WTI futures contracts for capturing global oil price dynamics, I choose the latter mainly because of the data availability for prices of continuous contracts. Prices of WTI futures contracts of up to 12 months to maturity are available back to March 1989, whereas those of Brent futures contracts are available after April 1994. Though deviating on a daily basis, a WTI-Brent spread, defined as the price difference between WTI and Brent futures contracts, maintains a stable range during most sample periods, which reflects the costs necessary to carry trades between the two delivery locations.

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., Bauer, G. H., and Diez de Los Rios, A. (2014). What does the convenience yield curve tell us about the crude oil market? *Bank of Canada*. Working paper 2014-42.

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. (2016). A model of financialization of commodities. *Journal of Finance*, 71(4):1511–1556.

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. *University of Notre Dame*. Working paper.

Brennan, D., Williams, J., and Wright, B. D. (1997). Convenience yield without the convenience: A spatial–temporal interpretation of storage under backwardation. *Economic Journal*, 107(443):1009–1022.

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.

Büyükşahin, B. and Robe, M. A. (2014). Speculators, commodities and cross-market linkages. *Journal of International Money and Finance*, 42:38–70.

Casassus, J. and Collin-Dufresne, P. (2005). Stochastic convenience yield implied from commodity futures and interest rates. *Journal of Finance*, 60(5):2283–2331.

Chow, Y.-F., McAleer, M., and Sequeira, J. (2000). Pricing of forward and futures contracts. *Journal of Economic Surveys*, 14(2):215–253.

Clark, T. E. and West, K. D. (2007). Approximately normal tests for equal predictive accuracy in nested models. *Journal of Econometrics*, 138(1):291–311.

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.

Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *Journal of Finance*, 57(1):405–443.

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 U.S. 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. (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.

Kaldor, N. (1939). Speculation and economic stability. *Review of Economic Studies*, 7(1):1–27.

Kilian, L. (2008). Exogenous oil supply shocks: How big are they and how much do they matter for the U.S. 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. (2016). The simple economics of commodity price speculation. *American Economic Journal: Macroeconomics*, 8(2):85–110.

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.

McCracken, M. W. (2007). Asymptotics for out of sample tests of granger causality. *Journal of Econometrics*, 140(2):719–752.

Milonas, N. T. and Henker, T. (2001). Price spread and convenience yield behaviour in the international oil market. *Applied Financial Economics*, 11(1):23–36.

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 Analyst Journal*, 68(6):54–74.

Telser, L. G. (1958). Futures trading and the storage of cotton and wheat. *Journal of Political Economy*, 66(3):233–255.

Thurman, W. N. (1988). Speculative carryover: An empirical examination of the U.S. refined copper market. *RAND Journal of Economics*, pages 420–437.

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: Equilibrium Storage Demand: January, February 2003

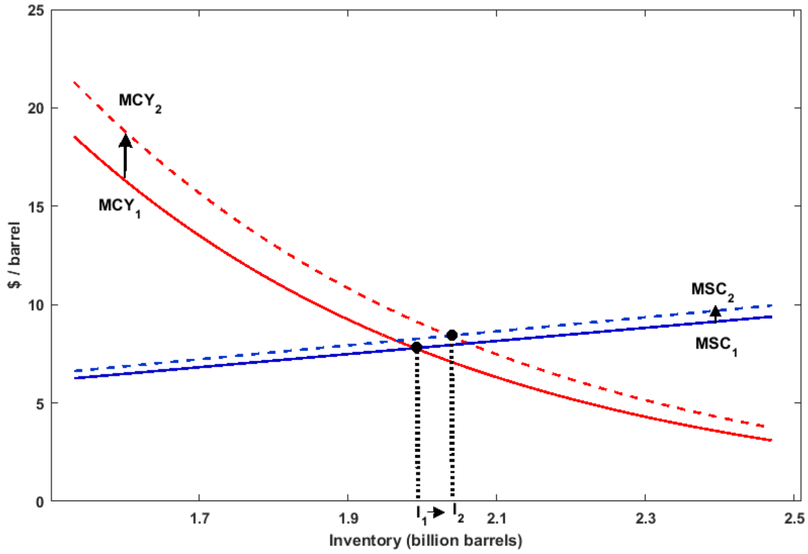

Note: This plots the refinery's storage decision and its dynamic adjustment process in response to an expected falling crude oil price during the period between January and February 2003. Two solid lines represent convenience yields and marginal costs in January 2003, while two dashed lines represent those in February 2003. For illustrative purposes, I calibrate the convenience yields and the marginal costs as functions of crude oil inventories using parameter estimates provided in Table 1 and estimates of state variables. The unit of the vertical axis is dollars per barrel (2010.5=100), with which the convenience yields and the marginal costs are measured.

Figure 2: Historical Decomposition of the Crude Oil Spot Price

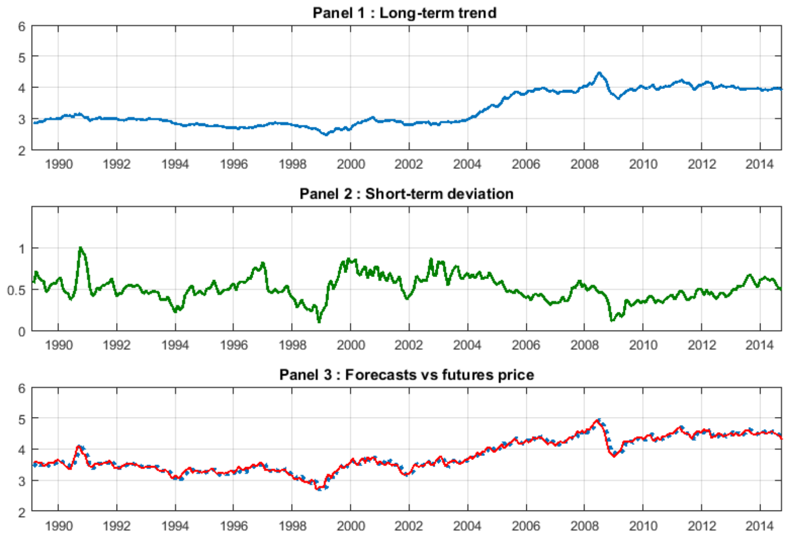

Note: This plots the historical decomposition of the crude oil spot price with the long-term trend process (Panel 1) and the short-term deviation process (Panel 2). Panel 3 plots the crude oil spot price (short-dashed line) together with the price for the nearest-maturity crude oil futures contract (solid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

Figure 3: Convenience yield

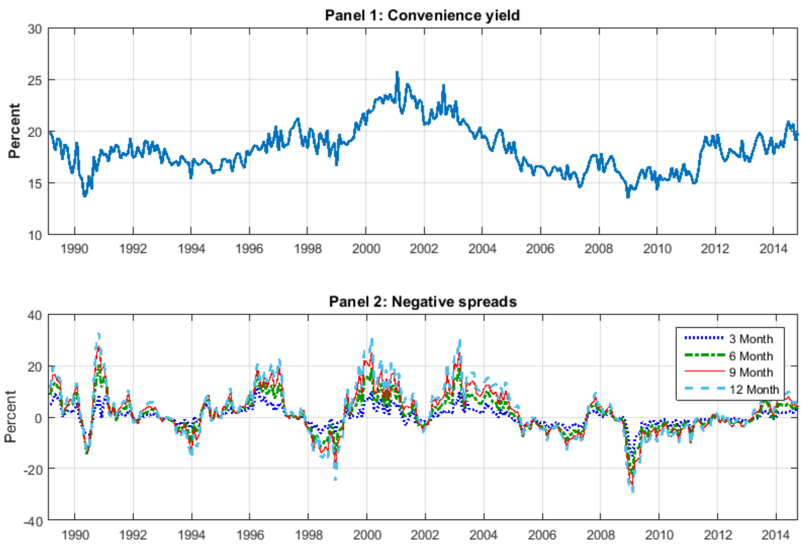

Note: Panel 1 displays the convenience yield as a percentage of the crude oil spot price. Panel 2 displays negative spreads calculated from real prices of crude oil futures contracts with three-, six-, nine- and 12-month maturities, and the real U.S refiners' acquisition cost for imported crude oil.

Table 1: Parameter Estimates

Period	(1) 1989.3-2014.11		(2) 1989.3-2004.6		(3) 2004.7-2014.11	
Output share						
α	0.3934	(0.0122)	0.4347	(0.0154)	0.3708	(0.0131)
Convenience yield						
θ_1 *(spring)*	0.1885	(0.0058)	0.1397	(0.0078)	0.1798	(0.0076)
θ_2 *(summer)*	0.1907	(0.0058)	0.1415	(0.0077)	0.1826	(0.0075)
θ_3 *(fall)*	0.1926	(0.0058)	0.1434	(0.0077)	0.1849	(0.0074)
θ_4 *(winter)*	0.1936	(0.0059)	0.1436	(0.0077)	0.1860	(0.0076)
Storage cost						
c_1	0.0086	(0.0008)	0.0152	(0.0017)	0.0066	(0.0008)
Trend						
g	0.0007	(< 0.0000)	0.0002	(< 0.0000)	0.0004	(< 0.0000)
Risk premium						
λ_ξ *(long)*	0.0024	(0.0002)	0.0001	(0.0004)	0.0063	(0.0004)
λ_χ *(short)*	-0.0427	(0.0664)	-0.0386	(0.0655)	-0.0387	(0.0486)
Spot process						
k^Q	0.9061	(0.0019)	0.9074	(0.0028)	0.9313	(0.0021)
ρ^Q	0.0966	(0.0601)	-0.0173	(0.0843)	0.1206	(0.0966)
σ_1	0.2438	(0.0124)	0.2246	(0.0143)	0.2469	(0.0180)
σ_2	0.0512	(0.0021)	0.0377	(0.0020)	0.0652	(0.0042)
σ_3	0.0778	(0.0033)	0.0910	(0.0050)	0.0623	(0.0045)
Likelihood	$14,211.81$		$8,174.74$		$6,845.20$	

Note: 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: MSPE and MAE Ratios Relative to the No-Change Forecast

	MSPE				MAE			
	Total		Excluded		Total		Excluded	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel 1: Prices of futures contracts with τ-months to maturity								
Horizon								
1	1.0038	(n/a)	**0.9842**	(0.33)	**0.9884**	(0.27)	**0.9765**	(0.13)
3	1.0088	(n/a)	**0.9937**	(0.39)	1.0020	(n/a)	**0.9920**	(0.23)
6	1.0176	(n/a)	**0.9994**	(0.49)	1.0176	(n/a)	**0.9927**	(0.30)
12	1.0147	(n/a)	1.0122	(n/a)	1.0147	(n/a)	1.0087	(n/a)
Panel 2: τ-month-ahead spot price forecast								
Horizon								
1	**0.9769**	(0.21)	**0.9804**	(0.28)	**0.9876**	(0.26)	**0.9844**	(0.24)
3	**0.9479**	(0.09)	**0.8825**	(0.07)	**0.9555**	(0.03)	**0.9394**	(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)

Note: 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 one-, three-, six- and 12-month maturities. Panel 2 reports those of one-, three-, six- and 12-month-ahead spot price forecasts recursively provided by the model forecasts. For each horizon, columns 1 and 5 report 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 3 and 7 provide recursive MSPE and MAE ratios calculated after excluding forecasts for five months from 2008.11 to 2009.3.