

What Do We Learn from the Price of Oil Futures?

Ron Alquist*
University of Michigan

Lutz Kilian
University of Michigan and CEPR

Preliminary and Incomplete! Do Not Quote without Permission!

First draft: February 28, 2007. This version: August 13, 2007

Abstract

It is common to rely on prices of oil futures to predict the path of the spot price of oil. Based on a two-country, two-period general equilibrium model of the spot and futures markets for crude oil, we clarify the conditions under which oil futures prices can predict the spot price. We document that these conditions are violated in the data. Using a newly constructed data set of monthly oil futures prices and oil spot prices that takes careful account of the exact dating of the underlying daily time series, we find that futures-based forecasts of oil prices tend to be less accurate than forecasts from alternative easy-to-use methods such as the no-change forecast. This result does not mean, however, that there is no useful information in oil futures prices. Under plausible conditions our model implies a positive correlation between the risk premium embodied in the price of oil futures and the component of the global spot price of oil driven by precautionary demand for oil. Under the random walk benchmark, this risk premium is observable. Its evolution over time suggests major shifts in precautionary demand during the Persian Gulf War and following the Asian crisis, for example. Our analysis provides independent evidence on how shifts in market expectations about future oil supply shortfalls affect the spot price of crude oil.

Key words: Crude oil; futures market; spot market; risk premium; basis; expectations; forecasting ability; precautionary demand.

JEL classification: C53, D51, G13, G15, Q31, Q43.

Acknowledgements: We thank Saul Hymans for assisting us in locating some of the data for this study. We also thank Ben Chabot, Steve Salant, and Ken West for helpful comments and suggestions. The first author acknowledges the generous support of the Center for International Business Education at the Stephen M. Ross School of Business, University of Michigan.

Correspondence: Department of Economics, University of Michigan, 611 Tappan St., Ann Arbor, MI 48109. Email: ralquist@umich.edu.

1. Introduction

Prices of crude oil futures are thought to convey useful real-time information regarding shifts in market expectations about future demand and supply conditions in the crude oil market. It is common to rely on futures prices of crude oil to infer the future path of the spot price of crude oil, in particular. Such forecasts are a crucial input for the macroeconomic models used at many central banks and at the International Monetary Fund. When macroeconomic forecasts fail, their failure can be often traced to inaccurate forecasts of the price of crude oil. Futures-based forecasts of the price of crude oil also play an important role in monetary policy decisions and affect financial markets' perceptions of the risks to price stability and sustainable growth.

This paper examines the common practice of predicting the spot price of crude oil based on the price of oil futures. Based on a theoretical model of the spot and futures markets for crude oil, we clarify the conditions under which prices of oil futures can predict future spot prices. We document that these conditions are violated in the data. Using a newly constructed data set of monthly oil futures prices and oil spot prices that takes careful account of the exact dating of the underlying daily time series, we show that futures-based forecasts in practice tend to be less accurate than forecasts from alternative easy-to-use models such as the random walk model. This result does not mean, however, that there is no useful information in oil futures prices. Our two-country, two-period general equilibrium model of the futures and spot markets for crude oil, implies that oil futures prices are informative about shifts in precautionary demand for crude oil, as reflected in the current spot price of crude oil. Specifically, in the presence of capacity constraints on crude oil production, the risk premium implied by a random walk or no-change forecast of the spot price of oil is shown to be positively correlated with the component of the spot price that is driven by precautionary demand for crude oil.

Our results confirm that the sharp spike in oil prices during the Persian Gulf War was mainly driven by an expectations shift reflected in higher precautionary demand, corroborating earlier results in Kilian (2007c) based on regression dummies as well as results in Kilian (2007a,b) based on historical decompositions. They also confirm that the temporary decline in oil prices following the Asian crisis and its reversal after 1999 reflected fluctuations in precautionary demand.

An independent empirical estimate of the precautionary demand component of the spot price of crude oil has recently been proposed by Kilian (2007a,b) for the period 1973-2006.

That estimate was based on a structural VAR model of the global crude oil market. We show that the VAR-based measure and the futures-based measure have a correlation of between 62 percent and 80 percent during 1989.1-2001.12, depending on the maturity of the futures data used. The correlation monotonically increases in the horizon.

Starting in 2002.1, however, the correlation between the risk premium and the VAR-based estimate of the precautionary demand component of the spot price vanishes. Whereas the VAR-based measure suggests an increase in concerns about future oil supply shortfalls, especially in 2005 and 2006 (consistent with an increase in geopolitical uncertainty, strong expected demand for crude oil and expectations of tight oil supplies), the futures-based estimate suggests a steady decline in precautionary demand after 2003. We show that this divergence is consistent with evidence of an unprecedented increase in speculative activities in oil futures markets starting in 2002. As increased demand for futures from speculators drives up futures prices, the risk premium falls, creating the mistaken impression that precautionary demand has fallen. We conclude that oil futures prices, although they are not useful for predicting future spot prices and in fact are dominated by simple no-change forecasts, may contain useful information about the determinants of the current spot price. They will not be reliable, however, unless major structural shifts in the composition of traders are accounted for.

The remainder of the paper is organized as follows. In section 2, we document the use of prices of oil futures as predictors of spot prices at central banks and international organizations. In section 3, we introduce a two-period, two-country general equilibrium model of the spot and futures markets for crude oil. In the model, an oil-producing country exports oil to an oil-consuming country that uses oil in producing a final good to be traded for oil or consumed domestically. Both oil producers and oil consumers may insure against uncertainty about stochastic oil endowments by holding oil inventories or trading oil futures. The spot and futures prices of oil are determined endogenously and simultaneously. Within this theoretical framework, we establish the conditions under which the current futures price is the expected spot price of oil, and we discuss the implications of our analysis for the specification of forecasting models.

Whether the conditions required for futures prices to be accurate predictors are satisfied in practice, is an empirical question that we take up in section 4. We provide a systematic evaluation of the predictive accuracy of the leading forecasting models based on the forecast

evaluation period 1991.1-2007.2. This period includes several major upheavals in the crude oil market. A robust finding across all horizons from 1 month to 12 months is that the no-change forecast is the most accurate forecast based on the mean-squared forecast error and the mean absolute forecast error. Simple no-change forecasts tend to be more accurate than forecasts based on futures prices or the futures spread. We also show that professional forecasts of the price of crude oil such as the 3-month and 12-month survey forecasts provided by *Consensus Economics Inc.* are dominated by both the random walk model and futures-based forecasts.

The poor forecasting accuracy of oil futures prices is consistent with the presence of a large and time-varying risk premium embodied in the futures price. In section 5 we study the determinants of this risk premium. In our theoretical model, fluctuations in the risk premium are driven by shifts in uncertainty about future oil supply shortfalls. Our theoretical analysis implies a positive correlation between the risk premium embodied in the futures price and the component of the spot price of oil driven by precautionary demand for crude oil. Thus, the risk premium may be viewed as an index of fluctuations in the price of crude oil driven by precautionary demand for oil. Given our evidence that simple no-change forecasts are most accurate in forecasting the spot price of crude oil, we can express the empirical counterpart of this risk premium in terms of observables. We discuss the evidence for expectations shifts in oil markets since 1989 and we compare the evolution of this risk premium to the VAR-based estimate of the precautionary demand component presented in Kilian (2007a,b). We also propose a measure of the term structure of risk premia that conveys the market's assessment of the preponderance of short-term versus long-term risks in the crude oil market. The concluding remarks are in section 6.

2. Oil Price Futures Are Widely Used as Predictors of the Future Spot Price

Futures contracts are financial instruments that allow traders to lock in today a price at which to buy or sell a fixed quantity of the commodity in a predetermined date in the future. In contrast to forward contracts, futures contracts can be retraded between inception and maturity on a futures exchange such as the New York Mercantile Exchange (NYMEX). A further difference between futures and forward contracts is that the holder of a futures contract realizes the gains or losses from holding that contract continuously rather than at maturity. The NYMEX offers institutional features that allow traders to transact anonymously. These features reduce individual default risk and ensure homogeneity of the traded commodity, making the futures market a low-cost and

liquid mechanism for hedging against and for speculating on oil price risks. The NYMEX light sweet crude contract is the most liquid and largest volume market for crude oil trading (NYMEX 2007a).

It is commonplace in policy institutions, including many central banks and the International Monetary Fund (IMF), to use the price of oil futures as a proxy for the market's expectation of the spot price of crude oil. A widespread view is that prices of NYMEX futures contracts are good proxies for the expected spot price of oil and better predictors of oil prices than econometric forecasts in particular. Forecasts of the spot price of oil are important inputs in the macroeconomic forecasting exercises that these institutions produce. For example, the European Central Bank (ECB) employs the future oil price in constructing the inflation and output-gap forecasts that guide monetary policy (see Svensson 2005). Likewise the IMF relies on futures prices as a predictor of future spot prices (see International Monetary Fund (2007), p. 42). Futures-based forecasts of the price of oil also play a role in policy discussions at the Federal Reserve Board (see, e.g., Greenspan 2004a,b; Bernanke 2004; Gramlich 2004). This is not to say that forecasters do not recognize the potential limitations of futures-based forecasts of the price of oil. Nevertheless, the perception is that oil futures prices, imperfect as they may be, are the best available forecasts of the spot price of oil.

Given the prominence of futures-based forecasts of the price of oil in practice, it is important not only to assess the empirical evidence for the forecasting ability of oil futures prices (as we do in section 4), but to use economic theory to understand the link between the spot market for crude oil and the futures market for crude oil. In section 3, we propose a theoretical framework for thinking about the relationship between the price of crude oil futures and the expected price of crude oil. Our analysis is based on a two-country, two-period general equilibrium model. We explicitly model the interaction of oil producers in the Middle East and oil consumers in the United States. The insights provided by this model will also be used to guide and inform our empirical analysis in sections 4 through 5.

3. A Two-Country General Equilibrium Model of the Oil Futures and Oil Spot Markets

The model in this section provides explicit microeconomic foundations for Pindyck's (1994, 2001) analysis. Pindyck discusses how equilibrium prices and quantities in generic spot and futures markets are endogenously determined by the interaction of the spot market, the market for storage, and the futures market. Our analysis builds on Townsend (1978). We extend the

framework used by Townsend in three directions. First, we add production of a final good. The model postulates that one representative trader in the economy, the United States, uses oil as an input in producing a final consumption good. Second, the model allows the two representative traders, the United States and Saudi Arabia, to carry inventories of oil from the first period to the second period. Third, the two representative traders have access to a bond market that allows them to save at a risk-free rate.¹

3.1. Model Description

There are two countries, the United States and Saudi Arabia. In the first period, Saudi Arabia trades its known oil endowment ω with the United States in exchange for a consumption good that the United States produces from oil to be delivered at the end of the period. The United States consumes some of the final consumption good and sells the rest to Saudi Arabia. Saudi Arabia's second period oil endowment is uncertain. There are two states of nature. With probability θ , Saudi Arabia receives the random endowment $\omega + \varepsilon$ in state 1. With probability $1 - \theta$, it receives the endowment $\omega - \hat{\varepsilon}$ in state 2, where $\hat{\varepsilon} = \varepsilon \theta / (1 - \theta)$ to ensure that the endowment shock has expectation zero. The variance of the second period oil endowment is σ^2 .

In the first period, both the United States and Saudi Arabia choose: (1) inventory holdings to carry forward to the second period; (2) the number of oil futures contracts that deliver one barrel of oil in the second period; and (3) the number of risk-free bonds that pay R in the second period. In addition, the United States chooses the quantity of oil to be used as the input into the production of the consumption good.

The optimization problem for the United States can be solved by backward induction. In the second period, the United States' problem is

$$\begin{aligned} & \max_{\{Z_{2s}\}} U(Y_{2s}) \\ & s.t. \\ & Y_{2s} = F(Z_{2s}) - \frac{S_{2s}}{P_{2s}}(Z_{2s} - I_{US}) + N_{US} \left(\frac{S_{2s}}{P_{2s}} - \frac{F_1}{P_{2s}} \right) + \frac{(1+R)}{P_{2s}} B_{US}, \end{aligned}$$

where Z_{2s} is the quantity of oil the United States chooses in period and state $s = 1, 2$; $U(\cdot)$ is the

¹ In related work, Britto (1984) proposes a general equilibrium model of a closed economy in which futures and spot prices of a generic commodity are determined endogenously. His model allows for production risk, but does not include inventories or savings.

United States' utility function defined over the consumption good; Y_{2s} is the quantity of the consumption good available to the United States in period 2 and state s ; $F(\cdot)$ is the strictly concave production function that generates the consumption good from oil; I_{US} is the quantity of oil the United States hold as inventory; N_{US} is the number of oil futures contracts the United States buy or sell in the first period for delivery in the second period; S_{2s} is the spot price of oil in state s ; P_{2s} is the price of the consumption good in state s ; F_1 is the price of a futures contract that delivers one barrel of oil in the second period; R is the risk-free interest rate; and B_{US} is the number of bonds the United States hold. In the second period I_{US} , N_{US} , and B_{US} are state variables.

In the first period, the United States chooses the quantity of oil, inventories, the number of oil futures contracts, and the number of bonds to maximize its utility:

$$\begin{aligned} & \max_{\{Z_1, I_{US}, N_{US}, B_{US}\}} U(Y_1) + g(I_{US}, \sigma) + \beta \left[\theta J_1(I_{US}, N_{US}, B_{US}) + (1 - \theta) J_2(I_{US}, N_{US}, B_{US}) \right] \\ & s.t. \\ & Y_1 = F(Z_1) - \frac{S_1}{P_1} (Z_1 + I_{US}) - \frac{B_{US}}{P_1}, \end{aligned}$$

where Z_1 is the quantity of oil the United States chooses in the first period; Y_1 is quantity of the consumption good available to the United States in period 1; $g(\cdot)$ is an increasing, strictly concave function in both inventories and the variance of the oil endowment that measures the convenience yield accruing to the United States from holding inventory between period one and period two; $\beta \in (0, 1)$ is the discount factor; S_1 is the spot price of oil period 1; P_1 is the price of the consumption good in period one; and J_s is the value function in state s .

Saudi Arabia's decision problem can be solved in the first period, because it faces no decision in the second period. Saudi Arabia's problem is

$$\begin{aligned} & \max_{\{I_{SA}, N_{SA}, B_{SA}\}} V(C_1) + h(I_{SA}, \sigma) + \beta \left[\theta V(C_{21}) + (1 - \theta) V(C_{22}) \right] \\ & s.t. \end{aligned}$$

$$C_1 = \frac{S_1}{P_1}(\omega - I_{SA}) - \frac{B_{SA}}{P_1}$$

$$C_{21} = \frac{S_{21}}{P_{21}}(\omega + \varepsilon + I_{SA}) + N_{SA}\left(\frac{S_{21}}{P_{21}} - \frac{F_1}{P_{21}}\right) + \frac{(1+R)}{P_{21}}B_{SA}$$

$$C_{22} = \frac{S_{22}}{P_{22}}(\omega - \hat{\varepsilon} + I_{SA}) + N_{SA}\left(\frac{S_{22}}{P_{22}} - \frac{F_1}{P_{22}}\right) + \frac{(1+R)}{P_{22}}B_{SA}$$

where $V(\cdot)$ is Saudi Arabia's utility function defined over the consumption good; C_1 is the quantity of the consumption good available to Saudi Arabia in the first period; C_{2s} is the quantity of the consumption good available to Saudi Arabia in period two and state $s = 1, 2$; $h(\cdot)$ is a strictly concave function in both its arguments that measures the convenience yield that accrues to Saudi Arabia from holding oil inventories between period one and period two; I_{SA} is the quantity of oil Saudi Arabia holds as inventory; and B_{SA} is the number of bonds Saudi Arabia holds.

These two optimization problems jointly yield 15 equations. There are six first-order conditions and three budget constraints for the United States; and there are three first-order conditions and three budget constraints for Saudi Arabia. In addition, there are eight market-clearing conditions:

$$C_1 + Y_1 = F(Z_1)$$

$$C_{21} + Y_{21} = F(Z_{21})$$

$$C_{22} + Y_{22} = F(Z_{22})$$

$$Z_1 = \omega - I_{US} - I_{SA}$$

$$Z_{21} = \omega + \varepsilon + I_{US} + I_{SA}$$

$$Z_{22} = \omega + \hat{\varepsilon} + I_{US} + I_{SA}$$

$$N_{US} + N_{SA} = 0$$

$$B_{US} + B_{SA} = 0$$

These 23 equations allow us in principle to solve for the 22 endogenous variables, the six consumption variables ($Y_1, Y_{21}, Y_{22}, C_1, C_{21}, C_{22}$); the two inventory holdings (I_{US}, I_{SA}); the two futures positions (N_{US}, N_{SA}); the two bond holdings (B_{US}, B_{SA}); the three input choices

(Z_1, Z_{21}, Z_{22}); and the seven relative prices ($\frac{S_1}{P_1}, \frac{F_1}{P_1}, \frac{S_{21}}{P_{21}}, \frac{S_{22}}{P_{22}}, \frac{F_1}{P_{21}}, \frac{F_1}{P_{22}}, R$). These conditions pin

down the general equilibrium in the economy.

The model has direct implications for the relationship between the futures price and the expected spot price of crude oil. Taking the first-order condition of $U(Y_{2s})$ with respect to U.S. holdings of futures contracts (N_{US}), we obtain

$$\beta \left[\theta U'(Y_{21}) \left(\frac{S_{21}}{P_{21}} - \frac{F_1}{P_{21}} \right) + (1-\theta) U'(Y_{22}) \left(\frac{S_{22}}{P_{22}} - \frac{F_1}{P_{22}} \right) \right] = 0.$$

Rearranging and using the definition of $\text{cov}(\cdot)$, this expression can be solved for today's futures price:

$$(1) \quad F_1 = \frac{\text{cov}\left(U'(Y_2), \frac{S_2}{P_2}\right)}{E\left[\frac{U'(Y_2)}{P_2}\right]} + \frac{E[U'(Y_2)]E\left[\frac{S_2}{P_2}\right]}{E\left[\frac{U'(Y_2)}{P_2}\right]}.$$

From the perspective of modern asset pricing theory, the first term in expression (1) may be interpreted as a *risk premium*. Under risk neutrality, the covariance term is zero, and using a first-order Taylor series approximation about the mean, we obtain:

$$(1') \quad F_1 \approx E[S_2]$$

If the United States are risk averse, in contrast, then expression (1) suggests that there is no simple relationship between the futures price of crude oil and the expected spot price. Thus, the question of whether futures prices for crude oil can forecast the future spot price depends on the unknown preference structure and is ultimately an empirical question. Section 4 will explore this empirical question in detail. In section 5, we will return to the model proposed here and will develop more fully its implications for the futures market, the spot market, and the evolution of the risk premium over time. Before we evaluate the forecast accuracy of futures prices in practice, however, it is useful to consider the implications of the possible presence of a risk premium for forecasting.

3.2. The Role of the Risk Premium in Forecasting Oil Prices

The theoretical model in section 3.1 suggests that the presence of a risk premium in futures prices of oil may undermine their forecasting ability. This observation has important

implications for the relative performance of alternative forecasting models. Let $F_t^{(h)}$ denote the current price of the futures contract that matures in h periods, S_t the current spot price of oil, and $E_t[S_{t+h}]$ the expected future spot price at date $t+h$ conditional on information available at t . Expression (1) implies that $F_t^{(h)}$ will equal $E_t[S_{t+h}]$ only under risk neutrality. More generally, $F_t^{(h)}$ and $E_t[S_{t+h}]$ will differ. The difference by construction is the *risk premium*. Short of taking a stand on traders' preferences, modeling the risk premium requires that we specify the expected future spot price and make assumptions about how traders form expectations. For any horizon h , the h -period ahead risk premium is given by

$$(2) \quad \rho_t^{(h)} \equiv \frac{E_t[S_{t+h}] - F_t^{(h)}}{S_t}.$$

The risk premium associated with holding the futures contract is expressed relative to the current spot price because the value of the futures contract is zero at inception.² By solving equation (2) for $F_t^{(h)}$, we obtain:

$$(3) \quad F_t^{(h)} = E_t[S_{t+h}] - \rho_t^{(h)} S_t$$

which links the futures price to the risk premium. Equation (3) illustrates that the futures price differs from the expected spot price if and only if there is a risk premium.

Accommodating the risk premium poses a problem in empirical work because expectations of market participants are not in general observable. One approach that sidesteps this problem is to assume that markets are informationally efficient and that ex-post realized returns are a good proxy for ex-ante expected returns. This assumption is implicit in forecast efficiency regressions that evaluate whether futures prices are rational expectations of future

² If the risk premium is positive, the expected return from holding a long futures position is positive – that is, the expected return associated with buying the futures contract is positive – and oil users have an incentive to purchase a futures contract at the observed futures price, while oil suppliers have an incentive to sell it; if the risk premium is negative, the expected return from holding a short futures position is positive – that is, the expected returns associated with selling the futures contract is positive – and oil suppliers have an incentive to buy a futures contract and oil users to sell it.

spot prices. A rejection of the restrictions implied by this model is interpreted as evidence of the existence of a time-varying risk premium (see, e.g., Fama and French 1987, 1988; Chernenko, Schwarz, and Wright 2004). Such tests also postulate that the trader's loss function coincides with the econometrician's quadratic loss function (see Elliott, Komunjer, and Timmermann 2005). Neither assumption is necessarily plausible in practice.

An alternative approach is to use proxies for market expectations of the spot price such as survey expectations. While economists have used survey data extensively in measuring the risk premium embedded in foreign exchange futures (see Chinn and Frankel 1995), this approach has not been applied to oil futures, with the exception of recent work by Wu and McCallum (2005). The survey approach allows the researcher to identify the ex-ante risk premium, dispensing with the assumption that expectations are necessarily model-consistent.

3.3. Alternative Measures of the Expected Spot Price of Oil

3.3.1. Expected Future Spot Prices from Surveys

Given the significance of crude oil to the international economy, it is surprising that there are few organizations that produce monthly forecasts of future spot prices. In the oil industry, where the spot price of oil is critical to investment decisions, oil firms tend to make annual forecasts of future spot prices for horizons as long as 15-20 years, but these are not publicly available. The U.S. Department of Energy's International Energy Agency (IEA) uses a structural econometric model of crude oil supply and demand to produce quarterly forecasts of the spot price of oil, but these forecasts are available only beginning in late 2004. The Economist Intelligence Unit has produced annual forecasts since the 1990s for horizons of up to 5 years. None of these sources provides monthly forecasts.

A widely used source of monthly forecasts of the price of crude oil is Consensus Economics Inc., a U.K.-based company that compiles private sector forecasts in a variety of countries. Initially, the sample consisted of more than 100 private firms, but it now contains about 70 firms. Of interest to us are the survey expectations for the 3- and 12-month ahead spot price of West Texas Intermediate crude oil, which corresponds to the type and grade usually delivered under the NYMEX futures contract. The survey provides the arithmetic average, the minimum, the maximum, and the standard deviation for each survey month beginning in October 1989 and ending in February 2007. We use the arithmetic mean at the relevant horizon:

$$(4) \quad \hat{S}_{t+h|t} = S_{t,h}^{CF}$$

3.3.2. Expected Future Spot Prices from Econometric Models

An alternative to modeling expectations of spot prices for crude oil is based on econometric models. In this section, we consider three types of econometric models: the random walk without drift, the random walk with drift, and the Hotelling model. Treating the conditional expectation from any one of these econometric models as a proxy for the market expectations, equation (2) allows us to construct the risk premium.

The random walk forecast of the future spot price of crude oil is simply the current spot price:

$$(5) \quad \hat{S}_{t+h|t} = S_t$$

Thus changes in the spot price are unpredictable. This parsimonious forecasting model is consistent with the view of Peter Davies, chief economist of British Petroleum, that “we cannot forecast oil prices with any degree of accuracy over any period whether short or long” (see Davies 2007). The random walk model has worked well in many other asset pricing contexts. Even if the spot price of crude oil does not truly follow a random walk, random walk forecasts may be attractive in terms of their mean-squared prediction error (MSPE) since the reduction in variance from excluding other predictors in small samples may more than offset the omitted variable bias.

Given that oil prices have been persistently trending upward (or downward) at times, we also consider a random walk model with drift. One possibility is to estimate this drift recursively, resulting in the forecasting model:

$$(6) \quad \hat{S}_{t+h|t} = S_t (1 + \alpha).$$

Alternatively, a local drift term may be estimated using rolling regressions:

$$(7) \quad \hat{S}_{t+h|t} = S_t (1 + \Delta \bar{s}_t^{(l)})$$

where $\hat{S}_{t+h|t}$ is the forecast of the spot price at $t+h$; and $1 + \Delta \bar{s}_t^{(l)}$ is the geometric average of the monthly percent change for the preceding l months, i.e., the percent change in the spot price between t and $t-l+1$. This model postulates that traders extrapolate from the spot price's recent behavior when they form expectations about the future spot price.

A third forecasting model is motivated by Hotelling's (1931) model, which predicts that the price of an exhaustible resource such as oil grows at the riskless rate of interest.

$$(8) \quad \hat{S}_{t+h|t} = S_t (1 + i_{t,h})$$

where $i_{t,h}$ refers to the interest rate at the relevant maturity h . Assuming perfect competition, no arbitrage, and no uncertainty, oil companies extract oil at a rate that equates: (1) the value today of selling the oil less the costs of extraction; (2) and the present value of owning the oil, which, given the model's assumptions, is discounted at the riskless rate. In competitive equilibrium, oil companies extract crude oil at the socially optimal rate. Although the Hotelling model seems too stylized to generate realistic predictions, we include this method given recent evidence that the Hotelling model does well in forecasting the future spot price of oil (see Wu and McCallum 2005).

3.3.3. Futures Prices as Expected Future Spot Prices

Perhaps the most common approach to constructing oil price forecasts is to rely on oil futures prices. A widespread view is that oil futures prices are good proxies for expected spot prices and better predictors than econometric forecasts in particular, despite the possible presence of a risk premium. In its simplest form, this forecasting model is:

$$(9) \quad \hat{S}_{t+h|t} = F_t^{(h)}$$

As discussed in section 3, in the presence of a risk premium, there is not necessarily a connection between the current crude oil futures price and the expected spot price. From a theoretical perspective, high futures prices relative to the current spot price are consistent with

increasing, constant, or falling spot prices (see Duffie 1989). Figure 1 shows a stylized example that illustrates the three cases. In each case, the futures price of \$50 exceeds the spot price of \$40, yet, depending on the movement of the risk premium, $E_t[S_{t+h}]$ may be higher than, equal to, or lower than the current spot price of \$40. Without further assumptions, the futures price is not a suitable predictor of the future spot price.

The common practice of basing oil price forecasts on futures prices may be defended, however, provided that (1) the risk premium is approximately constant; or (2) that the risk premium is small if it is time-varying. In that case, the futures price may still contain useful information about the future spot price and indeed may be the best available predictor.³

3.3.4. Forecasts Based on the Futures Spread

An alternative approach to forecasting the spot price of oil is to use the spread between the spot price and the futures price as an indicator of whether oil prices are likely to go up or down. This approach is referred to, for example, in a speech by Federal Reserve Board Governor Edward Gramlich (2004). The rationale for using the spread to forecast the future spot price is clear from rearranging (2):

$$\frac{F_t^{(h)}}{S_t} = \frac{E_t[S_{t+h}]}{S_t} - \rho_t^{(h)}$$

The spread, $F_t^{(h)}/S_t$, by construction contains information about both the expected percent change in the spot price and the risk premium. As with all forecasting models based on futures price data, for the spread to be a good predictor of the change in the spot price, the risk premium must be negligible.⁴ As Figure 1 shows, a large positive spread today is consistent with either rising, constant, or falling expected spot prices, depending the magnitude of the risk premium. We will explore the forecasting accuracy of the spread based on the forecasting models:

³ Prior empirical work provides no consensus about the existence of a risk premium in the crude oil market or commodity futures markets more broadly. Some authors have argued that commodity futures prices, and crude oil futures in particular, contain no risk premium (see Chernenko et al. 2004); others have found evidence of a significant time-varying risk premium in average commodity futures prices comparable in magnitude to the equity risk premium (see Gorton and Rouwenhorst 2006).

⁴ Alternatively, one could estimate the risk premium in real time and purge the effect of the risk premium on futures prices. This possibility is discussed in Chernenko et al. (2004) and Pagano and Pisani (2006).

$$(10) \quad \hat{S}_{t+h|t} = S_t \left(1 + \left(\alpha + \beta \ln(F_t^{(h)} / S_t) \right) \right),$$

$$(11) \quad \hat{S}_{t+h|t} = S_t \left(1 + \beta \ln(F_t^{(h)} / S_t) \right),$$

$$(12) \quad \hat{S}_{t+h|t} = S_t \left(1 + \ln(F_t^{(h)} / S_t) \right),$$

which differ in the extent to which the forecasting model is constrained to coincide with efficient market models (see, e.g., Chernenko et al. 2004).

4. Are Oil Prices Predictable? A Forecast Accuracy Comparison

4.1. Data Description and Timing Conventions

Our analysis is based on daily prices of crude oil futures traded on the NYMEX from the commercial provider *Price-Data.com*. The time series begins in March 30, 1983, when crude oil futures were first traded on the NYMEX, and extends through February 28, 2007. Crude oil futures can have maturities as long as 7 years. Trading ends on the last trading day four days prior to the 25th calendar day preceding the delivery month. If the 25th is not a business day, trading ends on the fourth business day prior to the last business day before the 25th calendar day (NYMEX 2007b).

In a typical year, there are 252 trading days or 21 trading days each month. Thus, for example, a 3-month ahead futures contract corresponds to a contract with 63 days until delivery. A common problem in constructing monthly futures prices at a given horizon is that an h -month contract may not trade on a given day. One way of dealing with this problem is to treat futures prices from a window in the middle of the month as a proxy for the futures price for a given month (see Chernenko et al. 2004). Another approach is to substitute the price of a j -month contract for a given day for the missing price of the h -month contract on that day where $j \neq h$, (see Bailey and Chan 1993). Our approach is different. At the end of a given month, we identify the contract closest to the last trading day of the month that corresponds to the h -month ahead futures contract and use the price associated with this contract as the end-of-month value. For all horizons, we obtain a continuous monthly time series based on a backward-looking window of at most five days. For maturities up to three months, the backward-looking window was three days. Our approach is motivated by the objective of computing in a consistent manner end-of-month time series of futures prices for maturities ranging from one month to five years.

4.2. The Choice of Maturities in the Empirical Analysis

The perception that futures prices contain information about future spot prices implicitly relies on the assumption that futures contracts are actively traded at the relevant horizons. In this subsection we investigate how liquid futures markets at each maturity h . This question is important because one would not expect $F_t^{(h)}$ to have predictive content for future spot prices, unless the market is sufficiently liquid at the relevant horizon.

Policymakers and the public widely believe that the oil futures market provides effective insurance against risks associated with crude oil production shortfalls and conveys the market's assessment of the evolution of future supply and demand conditions in the crude oil market. As a recent article in the *New Yorker* stated:

...[T]raders don't just look at today's supply and demand. They also try to forecast the future. And if buyers think there's a chance that supply is going to be lower down the line – because, say, Iranian oil fields will be shut down – they will be willing to pay a higher price today in order to guarantee that they will have the oil they need. (Surowiecki 2007)

If this perception were correct, one would expect active trading at long horizons. The evidence below, however, suggests otherwise. Figure 2 shows the monthly trading volume corresponding to a futures contract with a fixed horizon that is closest to the last trading day of the month.

Volume refers to the number of contracts traded in a given month.⁵ As illustrated in Figure 2, over the past 25 years, trading volume in the futures market has grown significantly, particularly at the 1-month and 3-month horizon, and to a lesser extent at the 6-month horizon.

In 1989, the NYMEX introduced for the first time offered contracts exceeding twelve months and in 1999, a 7-year contract was first introduced (see Haigh et al. 2007). Although such contracts are available, the market remains illiquid at horizons beyond one year even in recent years. Trading volume falls sharply at longer maturities.

This observation is important for our forecast evaluation because one would not expect forecasts based on futures with long maturities to provide accurate predictions, when only a handful of contracts are trading. Given the evidence in Figure 2, we therefore will restrict ourselves to futures contracts of up to one year in the empirical analysis below.

In addition, the evidence in Figure 2 suggests that the public and policymakers have

⁵ In contrast to open interest, volume measures the total number of contracts, including those in a position that a trader closes or that reach delivery, and thus gives a good sense of the overall activity in the futures market. Our method of data construction is consistent with the conventions used in constructing the monthly futures prices.

overestimated the ability of oil futures markets to provide insurance against long-term risks such as political instability in the Middle East or the development of oil resources in the Caspian Sea. For example, Greenspan (2004a) suggests that oil futures are pricing long-term risks associated with the future supply of oil:

“More worrisome [than upgrading refinery capacity] are the longer-term uncertainties that in recent years have been boosting prices in distant futures markets for oil [...] Prices for delivery in 2010 of light, low-sulphur crude rose to more than \$35 per barrel when spot prices touched near \$49 per barrel in late August. Rising geopolitical concerns about insecure reserves and the lack of investment to exploit them appear to be the key sources of upward pressure on distant future prices...”

Given that Greenspan’s speech was made in October of 2004 and referred to futures prices for delivery in 2010, we may infer that his assessment of risks must have been based on the 6-year futures contract. Greenspan (2004b) is even more explicit. In referring to the 6-year oil futures contract, Greenspan states that:

“... futures prices at that horizon can be viewed as effective long-term supply prices.”

As our volume data in Figure 2 show, there is very little information contained in futures prices beyond one year, making it inadvisable to rely on such data. This conclusion is also consistent with prior studies of the crude oil futures market between 1983 and 1994 (see Neuberger 1999) and with perceptions of industry experts.⁶

4.3. Forecast Evaluation Criteria

Tables 1 through 5 assess the predictive accuracy of the current futures price against the benchmark of a random walk without drift for horizons of 1, 3, 6, 9, and 12 months. We also include the random walk with drift, various random walk models with local drift, the Hotelling model (to the extent that interest rates are available at the relevant maturities), the consensus survey forecast (at the horizons for which survey forecasts are available), and forecasts from the regressions on the futures spread discussed further below. The forecast evaluation period is 1991.1-2007.2. We exploit the full set of available oil futures price data. The assessment of

⁶ According to sources within the oil industry who wish to remain anonymous, oil companies are fully aware of how thin the market is at longer horizons and do not rely on futures price data for such maturities. The perception is that one trader signing a couple of contracts with a medium-term horizon may easily move the futures price by several dollars on a given day.

which forecasting model is best may depend on the loss function of the forecaster (see Elliott and Timmermann 2007). We present results for the mean-squared prediction error (MSPE) and the mean absolute prediction error (MAPE). We also report the bias of the forecasts, and we report the number of times that a forecast correctly predicts the sign of the change of the spot price based on the success ratio statistic of Pesaran and Timmermann (1992). In addition to ranking forecasting models by each loss function, we formally test the null that a given candidate forecasting model is as accurate as the random walk without drift. Suitably constructed p -values are shown in parentheses.

4.4. Oil Futures as Predictors of Oil Spot Prices

Tables 1 through 5 document that the no-change forecast has lower MSPE than the futures forecast at the 1-month, 6-month, 9-month and 12-month horizon. Only at the 3-month horizon the futures forecast is more accurate, but the improvement in accuracy is not statistically significant. Moreover, based on the MAPE metric, the random walk forecast is more accurate at *all* horizons. In all cases, the random walk forecast is less biased than the futures forecast. Nor do futures forecasts have important advantages when it comes to predicting the sign of the change in oil prices. Only at the 9-month and 12-month horizons, the success ratio is significant at the 10 percent level and 5 percent level, respectively, but the improvement is only 2.6 and 3.6 percentage points. These results are consistent with the presence of a large and time-varying risk premium in oil futures prices.

The observation that futures prices are worse predictors of the price of oil than simple no-change forecasts is important because it contradicts commonly held views that current futures prices are a good guide to the evolution of future spot prices. For example, Federal Reserve Board Chairman Ben Bernanke in 2004 inferred from the level of futures prices that the market expected the preceding increase in spot prices to be long-lived:

“...[P]robably more economically significant than near-term uncertainty about oil prices is the fact that traders appear to expect tight conditions in the oil market to continue for some years, with at best only a modest decline in prices. This belief on the part of traders can be seen in the prices of oil futures contracts. Throughout most of the 1990s, market prices of oil for delivery at dates up to six years in the future fluctuated around \$20 per barrel, suggesting that traders expected oil prices to remain at about that level well into the future. Today, futures markets place the expected price of a barrel of oil in the long run closer to \$39, a near doubling. Thus, although traders expect the price of oil to decline somewhat from recent highs, they also believe that a significant part of the recent

increase in prices will be long lived.” (see Bernanke 2004)⁷

That conclusion is questionable for two reasons. First, as we showed in section 3, using long-horizon futures prices is problematic because few contracts are traded at these horizons and hence these prices convey little information. Second, Bernanke’s statements rely on the assumption that the risk premium is either approximately constant or small, which seems unlikely given the empirical evidence in Tables 1-5. Given the apparent importance of the risk premium in determining futures prices, the nature of the risk premium will be further documented and studied in section 5.

4.5. Other Predictors of the Spot Price of Oil

An obvious question is whether the no-change forecast can be improved upon, for example, by using information on interest rates.⁸ Tables 2, 3, and 5 show that the random walk model has lower MSPE than the Hotelling model at horizons of 3 and 6 months, whereas at the 12-month horizon the ranking is reversed. This reversal is not statistically significant, however. Based on the MAPE, the no-change forecast is superior at all three horizons. The Hotelling forecasting model has systematically lower bias at all three horizons than the no-change forecast. It also is systematically better at predicting the sign of the change in oil prices than futures forecasts, although we cannot assess the statistical significance of the improvement, given that there is no variability at all in the sign forecast.

Allowing for a drift in no case significantly improves the MSPE or MAPE of the random walk model, whether the drift is estimated based on rolling or recursive regressions. In some cases, allowing for a drift does improve significantly the ability to predict the sign of the change of the oil price. In general, the results for the random walk with drift are quite sensitive to the model specification and forecast horizon, and they do not account for the “specification mining” implicit in considering a large number of alternative models (see Inoue and Kilian (2004) and the references therein). There is no evidence that such models dominate the no-change forecast.

Finally, the consensus survey forecast has much higher MSPE than the no-change forecast at the both the 3-month and 12-month horizon. It also has a larger bias and higher

⁷ In the endnotes to this speech, Bernanke qualifies his comments by adding that his conclusions are conditional on futures prices being unbiased predictors of future spot prices and not containing a significant risk premium.

⁸ The interest rate data are the 3-month, 6-month, and 12-month constant-maturity Treasury bill rates from the Federal Reserve Board’s website <http://federalreserve.gov/releases/H15/data.htm>

MAPE and there is no statistically significant evidence that it is better at predicting signs than a coin flip. The survey forecast also is inferior to the futures forecasts, suggesting that survey respondents do not rely on futures data alone in forming their expectations.

4.6. Oil Future Spreads as Predictors of Future Spot Prices

The preceding subsection suggested that oil price forecasts based on futures prices, while clearly superior to the consensus survey forecast, are in turn outperformed by simple no-change forecasts. This subsection considers forecasting models based on the futures spread. Tables 1-5 show that the no-change forecast has lower MSPE than spread-based forecasts at horizons of 6, 9 and 12 months. At horizons 1 and 3 in some cases one of the spread models has lower MSPE, but the improvement is never statistically significant and no one spread model performs well systematically. Based on the MAPE, the no-change forecast is superior at all horizons. These results are consistent with the earlier evidence for the futures forecasts, and the presence of a large and time-varying risk premium. Tables 1 through 5 also show robust evidence that the spread models help predict the direction of change at horizons of 9 and 12 months. The gains in accuracy are statistically significant, but quite moderate. There is no such evidence at shorter horizons.

4.7. Relationship with Forecast Efficiency Regressions

It is interesting and useful to compare our results for the spread model in Tables 1 through 5 to the closely related literature on forecast efficiency regressions (see, e.g., Chernenko et al. 2004; Chinn, LeBlanc, and Coibion 2005). Consider the full-sample regression model:

$$\Delta s_{t+h} = \alpha + \beta (f_t^{(h)} - s_t) + u_{t+h}, \quad h = 1, 3, 6, 9, 12,$$

where lower-case letters denote variables in logs and u_{t+h} denotes the error term. Forecast efficiency in the context of the oil futures spread means that the hypothesis $H_0 : \alpha = 0, \beta = 1$ holds. Chernenko et al. report that the hypothesis of forecast efficiency cannot be rejected at conventional significance levels. Despite important differences in the timing conventions used in constructing the monthly futures price data, we are able to replicate their results qualitatively using our data (see Table 6), although one p -value is close to the rejection region. It is important to note that such evidence does not mean that oil prices are forecastable based on the spread in practice. First, non-rejection of a null hypothesis does not imply that the null model is true. In fact, we showed that the forecasting model (12) that imposes this null does not dominate the no-

change forecasts in out-of-sample forecasts. Second, as our results show, relaxing one or more of the restrictions implied by forecast efficiency may either improve or worsen the forecast accuracy of the spread model, depending on the bias-variance trade-off. In particular, such models require the estimation of additional parameters compared with the no-change forecast, and the resulting loss in forecast precision may outweigh the benefits from reduced forecast bias. Thus, there is no contradiction between our results and the forecast efficiency results in the literature.

5. Understanding the Risk Premium

The forecast accuracy comparison of section 4 established that the no-change forecast is the best available predictor of the spot price of crude oil in practice, at least if the objective is to minimize the MSPE or the MAPE. Based on this benchmark for $E_t[S_{t+h}]$, it is straightforward to compute a direct measure of the risk premium in expression (2):

$$\rho_t^{(h)} \equiv \frac{E_t[S_{t+h}] - F_t^{(h)}}{S_t}.$$

in terms of the observables S_t and $F_t^{(h)}$. This puts us in a position to assess the magnitude of the risk premium and the extent to which the risk premium has been fluctuating historically.

Before assessing this evidence, it is useful to address the question of why the risk premium should be fluctuating over time. The two-country, two-period general equilibrium model of section 3 sheds light on this question. In the model, the fundamental driving the risk premium embodied in oil futures prices is shifts in the uncertainty about future oil supply shortfalls. Whereas concerns about future supply shortfalls may in principle arise in any commodity market, there is reason to believe that such concerns historically have been particularly relevant in the crude oil market and may explain the existence both of large risk premia and of sharp fluctuations in risk premia over time.

It is important to note that unanticipated shifts in uncertainty about future oil supply shortfalls differ from actual supply shortfalls. In particular, it is not necessary for any of the underlying expectation shifts actually to be realized. In fact, a common situation in practice is that concerns about possible shortfalls of oil production driven, for example, by fears that OPEC would reinstitute an oil embargo in the 1970s, that Khomeini would attack moderate Gulf states in the 1980s, or that Saddam Hussein would invade Saudi Arabia in 1990 have evaporated over

time, only to be replaced by new and different concerns such the nuclear threat from Iran or the instability of the Saudi throne. This situation is not unlike that of a peso problem in foreign exchange markets.

As we will demonstrate below, a theoretical link can be established between shifts in expectations about future oil supply and the *basis* of oil futures defined by $basis_t = (F_t^{(h)} - S_t)/S_t$. Under the random walk without drift model, $E_t[S_{t+h}] = S_t$, so the risk premium reduces to $\rho_t^{(h)} \equiv (S_t - F_t^{(h)})/S_t = -basis_t$. Thus, fluctuations in the risk premium are directly linked to shifts in uncertainty about future oil supply shortfalls. The model also implies that the spot price of oil increases in response to adverse expectations shifts, which implies a positive correlation between the random walk risk premium and the component of the spot price of oil driven by concerns about future oil supply shortfalls. In other words, the model implies a positive relationship between the risk premium and increases in the spot price of oil driven by precautionary demand for oil in response to shifting concerns about future oil supply shortfalls. Thus, we can interpret the random walk risk premium as an index of the expectations-driven component of oil prices and compare that estimate to independently constructed estimates of this component in the literature. In section 5.1, we derive the comparative statics results that form the basis of the subsequent analysis.

5.1. Comparative Statics

In this subsection, we derive two comparative statics results. The first result is that an increase in uncertainty about the availability of oil supplies in the second period raises the first period's spot price; the second result is that this increase in uncertainty lowers the basis in the first period. In both cases, we model the increase in uncertainty as a mean-preserving increase in the spread of the second period oil endowment shock.

To derive the comparative statics results, we make a simplifying assumption. We assume that both the United States and Saudi Arabia are risk neutral and hence that their marginal utilities are constants normalized to one. This assumption makes the results transparent, because it avoids introducing assumptions about the curvature of the utility functions associated with different degrees of risk aversion. If we did not assume risk-neutrality, the results would hinge

on traders' risk preferences.⁹

5.2.1. The Effect of an Increase in Uncertainty on the First-Period Spot Price

From the United States' first-order condition, the marginal efficiency conditions, and the market clearing conditions, we obtain

$$F'(\omega - I_{US} - I_{SA}) - g_{I_{US}}(I_{US}, \sigma) = \beta[\theta F'(\omega + \varepsilon + I_{US} + I_{SA}) + (1 - \theta)F'(\omega - \hat{\varepsilon} + I_{US} + I_{SA})],$$

where the notation $g_{I_{US}}(\cdot)$ denotes the partial derivative of the function $g(\cdot)$ with respect to U.S. inventory holdings. From Saudi Arabia's first order conditions, we obtain the analogous condition:

$$F'(\omega - I_{US} - I_{SA}) - h_{I_{SA}}(I_{SA}, \sigma) = \beta[\theta F'(\omega + \varepsilon + I_{US} + I_{SA}) + (1 - \theta)F'(\omega - \hat{\varepsilon} + I_{US} + I_{SA})],$$

where the notation $h_{I_{SA}}(\cdot)$ denotes the partial derivative of the function $h(\cdot)$ with respect to Saudi Arabia's inventory holdings. Totally differentiating both of these expression with respect to I_{US} , I_{SA} , and ε yields two equations in two unknowns, which can be compactly represented in matrix notation

$$\begin{bmatrix} A - g_{I_{US}I_{US}} & A \\ A & A - h_{I_{SA}I_{SA}} \end{bmatrix} \begin{bmatrix} dI_{US} \\ dI_{SA} \end{bmatrix} = \begin{bmatrix} B \\ B \end{bmatrix} d\varepsilon,$$

where

$$A = -[F''(\omega - I_{US} - I_{SA}) + \beta\theta F''(\omega + \varepsilon + I_{US} + I_{SA}) + \beta(1 - \theta)F''(\omega - \hat{\varepsilon} + I_{US} + I_{SA})] > 0$$

$$B = \beta\theta[F''(\omega + \varepsilon + I_{US} + I_{SA}) - F''(\omega - \hat{\varepsilon} + I_{US} + I_{SA})] > 0.$$

Solving for the change in U.S. and Saudi oil inventories yields

$$\begin{bmatrix} \frac{dI_{US}}{d\varepsilon} \\ \frac{dI_{SA}}{d\varepsilon} \end{bmatrix} = \begin{bmatrix} A - g_{I_{US}I_{US}} & A \\ A & A - h_{I_{SA}I_{SA}} \end{bmatrix}^{-1} \begin{bmatrix} B \\ B \end{bmatrix}$$

$$= \frac{1}{g_{I_{US}I_{US}}h_{I_{SA}I_{SA}} - A(g_{I_{US}I_{US}} + h_{I_{SA}I_{SA}})} \begin{bmatrix} -Bh_{I_{SA}I_{SA}} \\ -Bg_{I_{US}I_{US}} \end{bmatrix} > 0,$$

which shows that both the United States and Saudi Arabia accumulate inventories in response to

⁹ As Varian (1985) and Barsky (1989) show, even in asset pricing models with no production, deriving unambiguous conclusions about how asset prices change in general equilibrium requires strong assumptions on traders' risk preferences and the distribution of shocks.

a mean-preserving increase in the spread. Since $\frac{S_1}{P_1} = F'(\omega - I_{US} - I_{SA})$ is a decreasing convex function of U.S. and Saudi inventories, the increase in inventories raises the first-period spot price in real terms. Thus the first implication of the model is that an increase in uncertainty raises the spot price of oil in terms of the consumption good.

5.2.2. The Effect of an Increase in Uncertainty on the Oil Futures Basis

We can express the oil futures basis as a function of the marginal convenience yield, the real spot price of oil, and the risk-free interest rate:

$$\frac{F_1 - S_1}{S_1} = R - (1 + R) \frac{g_{I_{US}}(I_{US}, \sigma^2)}{\frac{S_1}{P_1}}.$$

Totally differentiating the right-hand side of this equation, we obtain the expression

$$\frac{dR}{d\varepsilon} - \frac{dR}{d\varepsilon} \left[\frac{g_{I_{US}}}{\frac{S_1}{P_1}} \right] - (1 + R) \frac{[g_{I_{US}} F''(\omega) + g_{I_{US}I_{US}} F'(\omega)] \frac{dI_{US}}{d\varepsilon} + g_{I_{US}} F''(\omega) \frac{dI_{SA}}{d\varepsilon} + g_{I_{US}\sigma^2} F'(\omega) \frac{d\sigma^2}{d\varepsilon}}{[F''(\omega)]^2},$$

where we use the shorthand $g_{I_{US}} = g_{I_{US}}(I_{US}, \sigma^2)$ and $F(\omega) = F(\omega - I_{US} - I_{SA})$. Assuming that $dR/d\varepsilon \approx 0$, the sign of the expression depends on the relative magnitudes of: (1) the decrease in the marginal convenience yield associated with the increase in inventories triggered by the shock to the endowment distribution; and (2) the increase in the marginal convenience yield associated with the increase in σ triggered by the same shock. Algebraically, the basis will decline if and only if

$$g_{I_{US}\sigma} F'(\omega) \frac{d\sigma^2}{d\varepsilon} > -[g_{I_{US}I_{US}} F'(\omega) + g_{I_{US}} F''(\omega)] \frac{dI_{US}}{d\varepsilon} - g_{I_{US}} F''(\omega) \frac{dI_{SA}}{d\varepsilon},$$

It follows that a sufficiently large increase in uncertainty will cause the basis to decrease. Thus, one would expect a negative correlation between the basis and the component of the spot price driven by precautionary demand for oil at least during major shifts in uncertainty. A decrease in the basis (or equivalently an increase in the random walk risk premium) seems the more likely, the smaller the decrease in the marginal convenience yield associated with an increase in inventory holdings.

In traditional models that incorporate the marginal convenience yield such as Brennan

(1958), Telser (1958), and Fama and French (1988) an increase in inventory holdings has an unambiguously positive effect on the spot price because of the decreasing returns to holding inventories. As inventory holdings increase, the marginal convenience yield declines, raising the basis. Our conclusion is different. The reason is that we follow Pindyck (2001) in incorporating a *direct* effect from increased uncertainty on the marginal convenience yield, which raises the value of holding inventories when uncertainty is high and pushes up the marginal convenience yield for each unit of inventory. The distinction is that the indirect effect through inventories reflects the accumulation of additional inventories in response to the increase in uncertainty; whereas the effect operating through σ^2 captures the notion that each unit of inventories now has greater value, as uncertainty has increased. This increase in the value of inventories reflects increased demand for oil in the United States in response to greater uncertainty. We refer to this demand as precautionary demand for crude oil.

5.3. Using the Model to Identify Expectations Shifts

The model of section 3 suggests that fluctuations in the basis (or equivalently in the risk premium implied by the random walk model) are indicative of fluctuations in the spot price of oil driven by precautionary demand for crude oil. Strictly speaking, this link holds if and only if a change in demand for oil inventories is confronted with an inelastic supply of oil. In the model, this inelasticity is represented in the form of an endowment structure. While this assumption may be unrealistic for the early 1980s, throughout much of the sample that we consider below this is a reasonable assumption. Kilian (2007c) documents that capacity constraints in world crude oil production have been binding since the early 1990s. Since the risk premium at the 12-month horizon becomes available only in 1989.1, we will focus on the period 1989.1-2006.12.

Table 7 summarizes the time series properties of the risk premium implied by our analysis in section 4. For expository purposes, we focus on the 3-month and 12-month horizon. Table 7 confirms, first of all, that the risk premium tends to be large in magnitude and that the futures price is a biased predictor of the spot price. Using heteroskedasticity and autocorrelation consistent (HAC) standard errors, both the 3-month and 12-month risk premium are statistically different from zero at the 1% level. To get a sense of the average magnitude of the risk premium, Table 7 also presents the mean absolute deviation. This measure shows that, on average, the magnitude of the risk premium is substantial. Equation (3) allows us to get a sense of the economic significance of the estimated risk premium. Taking the spot price of crude oil to be

\$65, about its level in late March 2007, the minimum (maximum) value of the 3-month ahead risk premium implies that the futures price can trade at a premium (discount) of as much as \$19 (\$27). The mean implies that it trades at about a \$3 premium. Somewhat smaller values are obtained for the 12-month risk premium. Overall, these calculations suggest that the magnitude of the risk premium can be economically significant.

Even if there is a risk premium, future prices may still provide unbiased forecasts as long as it is constant and one includes an intercept in the forecasting model. As Table 7 shows, however, there is strong evidence of persistent time variation in the risk premium that will cause futures prices to be biased predictors of spot prices. We capture this time variation by modeling the risk premium as an autoregression with the lag order selected by the Akaike Information Criterion. Based on the fitted autoregressive models, we compute the sum of the autoregressive coefficients as measure of persistence as suggested by Andrews and Chen (1994). The sum of the autoregressive coefficients for the 3-month risk premium is 0.68, whereas that for 12-month risk premium is 0.84. This evidence confirms that there is considerable persistence in the risk premium. The long-horizon risk premium is more persistent than the short-horizon risk premium.

It is even more interesting to identify the episodes associated with shifts in the risk premium and hence expectations. Figure 3 plots the random walk risk premium for 1989.1-2007.2 by horizon. The plot confirms that the sharp spike in oil prices during the Persian Gulf War was mainly driven by an expectations shift reflected in higher precautionary demand, corroborating earlier results in Kilian (2007c) based on regression dummies representing expectations shifts as well as results in Kilian (2007a,b) based on historical decompositions from structural VAR models. The plot also confirms that the temporary decline in oil prices following the Asian crisis and its reversal after 1999 reflected fluctuations in precautionary demand. In addition, the plot suggests a persistent decline in precautionary demand in recent years, a point that we will discuss in more detail below.

5.4. The Term Structure of Risk Premia

Further information may be obtained from the term structure of risk premia. The term structure is informative about whether the market expects higher risks to oil prices at long horizons than at short horizons. Figure 4 plots the (demeaned) difference between the 12-month and 3-month risk premium. A positive value of this term structure indicates that the market expects risks to

increase over time, whereas a negative value indicates a decline in the compensation required for risks associated with future uncertainty.

For example, the outbreak of the Persian Gulf War of August 1990 was preceded by expectations of increasing risks. As the spot price jumped upon the invasion of Kuwait, the term structure turned sharply negative, indicating a greater short-term risk than long-term risk. By April 1991, after the war had ended, the term structure turned positive again. Similarly, the Iraq War of early 2003 was associated with a negative spike in the term structure, whereas the most recent period (as well as the period leading up to the 2003 Iraq War) was characterized mostly by a positive term structure, indicating the expectation of increasing risks. Wars are not the only periods characterized by a negative term structure. There also was an episode from late 1993 until late 1996, during which the market expected risks to decline at longer horizons. The intermittent period of 1997-2001 appears volatile, but the term structure is less persistent than in earlier periods.

5.5. Alternative Measures of Precautionary Demand Shifts

The index of expectations-driven oil price increases proposed in this paper is not the only possible approach. Recently, an alternative measure of the component of the spot price of crude oil that is driven by shocks to precautionary demand has been proposed by Kilian (2007a,b). Unlike the measure developed in this paper, that estimate was based on a structural VAR decomposition of the price of crude oil. The structural representation of the underlying trivariate autoregressive model is

$$A_0 z_t = \alpha + \sum_{i=1}^{24} A_i z_{t-i} + \varepsilon_t ,$$

where ε_t denotes the vector of serially and mutually uncorrelated structural innovations and z_t a vector variable including the percent change in global crude oil production, a (suitably detrended) measure of real economic activity, as it relates to industrial commodity markets, and the real price of oil (in that order) at monthly frequency. Let e_t denote the reduced form VAR innovations such that $e_t = A_0^{-1} \varepsilon_t$. The structural innovations are derived from the reduced form innovations by imposing exclusion restrictions on A_0^{-1} . The identifying assumptions are that (1) crude oil supply will not respond to oil demand shocks within the month, given the costs of adjusting oil production and the uncertainty about the state of the crude oil market; that (2)

increases in the real price of oil driven by shocks that are specific to the oil market will not lower global real economic activity within the month; and that (3) innovations to the real price of oil that cannot be explained by oil supply shocks or shocks to the aggregate demand for industrial commodities must be demand shocks that are specific to the oil market. As documented by Kilian (2007a), the latter shock by construction captures fluctuations in precautionary demand for oil driven by fears about the availability of future oil supplies. From the structural moving average decomposition, one can construct a time series of the component of the real price of oil that can be attributed to shifts in the precautionary demand for crude oil in response to shifting concerns about the future oil supply shortfalls.

While it is not possible to compare this VAR-based measure to the futures-based measure for the entire sample period of 1973-2006 considered in Kilian's work, given the limited availability of oil futures data, we may compare these two independently obtained measures for the period 1989.1-2006.12, which includes several major oil price spikes. Table 8 shows that the two measures in general are highly correlated notwithstanding the differences in their method of construction. From 1989.1 through 2001.12, the correlation ranges from 62% at the 3-month horizon to 80% at the 12 month horizon. The fit improves monotonically with the horizon, consistent with the view that shifts in precautionary demand are primarily concerned with expectations beyond the short run.

Table 8 also shows, however, that starting in 2002 the positive correlation between the basis and the VAR-based estimate of the precautionary demand component of the spot price vanishes. Whereas the VAR-based measure suggests an increase in concerns about future oil supply shortfalls, especially in 2005 and 2006 (consistent with an increase in geopolitical uncertainty, strong expected demand for crude oil and expectations of tight oil supplies), the futures-based estimate suggests a rather implausible steady decline in precautionary demand after 2002. This pattern is reflected in a near-reversal of the correlations in Table 8 relative to the pre-2002 sample. This reversal is also clearly seen in Figure 5. Starting in 2002, the basis and the VAR-based estimate of the precautionary demand component of the spot price begin to move in opposite directions.

5.6. Explaining the Breakdown

The fact that the theoretical model of section 3 generates economically implausible predictions after 2001 and predictions that depart from independent estimates of precautionary demand,

suggests that a structural change may have occurred around 2001/2002 that is beyond the scope of the theoretical model in section 3. Indeed, it has been suggested in the financial press that the nature of the oil futures markets has changed in recent years, as hedge funds and other investors with no ties to the oil industry attempted to capitalize on rising oil prices. Verifying this claim is difficult. By its nature, the NYMEX market for crude oil futures is anonymous, and it is commensurately difficult to pin down exactly the extent to which the hedging or the speculative motive of trading predominates. The Commodity Futures Trading Commission (CFTC) collects data that speak to this issue, however. In order to monitor the amount of speculation in commodity futures markets, the CFTC requires that traders who hold positions in the futures market report the nature of their business. Its weekly Commitments of Traders (COT) report records the positions of traders in every market for which 20 or more traders hold positions at or above a threshold required by the Commission. These data comprise 70-90% of the oil futures contracts traded on the NYMEX (CFTC 2007).

The CFTC uses these reports to classify the traders into “commercial” and “non-commercial”. Commercial traders are entities that use futures contracts in the commodity as defined by the CFTC’s regulations (CFTC 2007). The precise wording from CFTC regulation 1.3(z) reads:

Bona fide hedging transactions and positions shall mean transactions or positions in a contract for future delivery on any contract market, or in a commodity option, where such transactions or positions normally represent a substitute for transactions to be made or positions to be taken at a later time in a physical marketing channel, and where they are economically appropriate to the reduction of risks in the conduct and management of a commercial enterprise... (CFTC 2007)

More succinctly, a trader is considered commercial by the CFTC as long as it is engaged in business activities hedged by the use of futures. Such traders are identified as hedgers.¹⁰ Among the types of firms classified as commercial by the CFTC are merchants, manufacturers, producers, and commodity swaps and derivative dealers (see Haigh, Harris, Overdahl, and Robe 2007). Non-commercial traders are all other traders, such as hedge funds, floor brokers and traders, and non-reporting traders, and are identified as speculators. To guarantee that traders are classified accurately, the CFTC can re-classify a trader at its discretion if it has additional information about how the trader uses the future market.

¹⁰ Appendix 1 contains all of section 1.3(z).

The CFTC records the open interest of each type of trader for short and long positions. “Open interest” is the total number of long positions outstanding in a futures contract, which is equal to the total number of short positions in a futures contract, and measures the number of outstanding contracts that exist at a point in time (see Hull 2006). The open interest held by an individual trader is the trader’s position. When a trader opens a new position, the seller writes a new contract, increasing open interest by one. If the trader closes his position at or prior to delivery, the contract is terminated and open interest decreases by one. For non-commercial traders, the CFTC also records the number of contracts associated with a spread position. A “spread” position entails the simultaneous purchase of one contract and the sale of another with the intention of exploiting the relative price differential between the two contracts. For example, a trader may have a long position in the near crude oil contract and a short position in the long-term crude oil contract in anticipation of rising spot prices in the near term and declining spot prices in the long term. These two positions offset one another and hence count as a single spread position.

These data are likely to be reliable. Firms that use the futures market to hedge price risk can deduct the positions as a business expense from their taxes, so misreporting the nature of one’s business has significant legal implications and is considered tax fraud. Moreover, although hedge funds in particular have an incentive to conceal their positions among a variety of different brokers, the CFTC monitors and enforces a regulation that considers multiple positions subject to common ownership as a single position (see CFTC 2007). These institutional mechanisms reduce the scope for misreporting among firms active in the crude oil futures market. Moreover, there is no incentive for firms to misreport themselves as non-commercial firms, so any increase in such positions is likely to be genuine.

A natural measure of the relative importance of speculative activities is the number of noncommercial spread positions expressed as a percentage of the reportable open interest positions. Figure 6 shows a market increase in the percent share of the noncommercial spread positions after 2001.¹², suggesting that speculation has recently intensified. In fact, Figure 6 suggests that there have been spikes in speculative activity, whenever the spot price of oil moved persistently (or predictably) up or down, notably in 1990/91 during the Persian Gulf War, when the price of oil collapsed after the Asian Crisis, and when it recovered in 1999. However, the recent increase in the non-commercial spread position is clearly unprecedented in the sample.

The percentage in question has more than tripled since 2001.

This evidence matters because, as increased demand for oil futures from speculators drives up futures prices, the basis increases, and the random walk risk premium falls, creating the mistaken impression that precautionary demand has fallen. This breakdown is expected, as our theoretical model in this section does not allow for speculation. Whereas the VAR-based approach appears immune from the presence of speculators in the oil futures market, the futures-price based approach breaks down starting in 2002. We conclude that oil futures prices (or related variables such as the basis or the spread), although they are not useful for predicting future spot prices and in fact as predictors are dominated by simple no-change forecasts, may contain useful information about the determinants of the current spot price. They will not be reliable, however, unless major structural shifts in the composition of traders are accounted for.

6. Conclusion

In this paper, we introduced a two-country, two-period general equilibrium model of both the spot market and the futures market for crude oil to provide fresh insights about the interpretation of oil futures prices and related statistics such as the futures spread or basis. The key insights can be summarized as follows:

- Our theoretical model illustrates that prices of oil futures would be expected to forecast spot prices only under certain restrictive conditions. How well these forecasts work in practice therefore is an empirical question and depends on the properties of the risk premium.
- Using a newly constructed data set of monthly oil futures prices and oil spot prices that takes careful account of the exact dating of the underlying daily time series, we documented that prices of crude oil futures are not useful predictors of the spot price of crude oil in practice. We showed that the presence of a large and time-varying risk premium invalidates such forecasts. For example, given a spot price of \$60, the futures price may deviate as much as \$43 from the expected spot price or as little as \$0.
- Many users of futures-based forecasts are aware of these caveats and understand that futures-based forecasts may be poor, but still believe that futures-based forecasts provide the *best* available forecast of spot prices of crude oil. We showed this not to be the case. Futures-based forecasts are inferior to simple and easy-to-use forecasting methods such as the no-change forecast. Such forecasts are also more accurate than commercial survey-based forecasts.
- Given our evidence that simple no-change forecasts are most accurate in forecasting the spot

price of crude oil, we can express the empirical counterpart of the risk premium in the theoretical model in terms of observables.

- In our theoretical model, fluctuations in this risk premium are driven by shifts in the uncertainty about future oil supply shortfalls. Our theoretical analysis implies a positive correlation between the risk premium and the component of the spot price of oil driven by precautionary demand for crude oil. Thus, the risk premium may be viewed as an index of fluctuations in the price of crude oil driven by precautionary demand for oil.
- The estimated random walk risk premia suggest major shifts in precautionary demand for oil during the Persian Gulf War and following the Asian crisis, for example. These results provide independent evidence on how shifts in market expectations about future oil supply shortfalls affect the spot price of crude oil. They are also consistent with related evidence in the literature obtained by alternative methodologies (see Kilian 2007a,b,c).
- The ability of the risk premium to capture shifts in precautionary demand sharply deteriorates after 2002, however, consistent with evidence of an unprecedented increase in speculative activities in the oil futures market. This breakdown is expected, as our theoretical model does not allow for speculation and increased speculative trading will tend to depress the observed risk premium. Thus, while futures price data can be useful in identifying expectations shifts in general, our results suggest caution in interpreting oil futures data in the absence of further information about the market structure.
- We also proposed a term structure of random walk risk premia as a measure of the market's expectation of whether risks in the crude oil market are increasing or decreasing in the forecast horizon. This measure suggests, for example, that futures traders have been expecting risks to increase at longer horizons for the most part since mid-2004.
- Finally, we showed that oil futures markets are too illiquid at horizons beyond one year to provide effective insurance against risks of future oil supply shortfalls at medium and long-term horizons.

References

- Andrews, D.W.K. and H.-Y. Chen (1994), "Approximately Median-Unbiased Estimation of Autoregressive Models," *Journal of Business and Economic Statistics*, 12(2), 187-204.
- Bailey, W. and K.C. Chan (1993), "Macroeconomic Influences and the Variability of the Commodity Futures Basis," *Journal of Finance*, 48(2), 555-573.
- Barsky, R.B. (1989), "Why Don't the Prices of Stocks and Bonds Move Together?," *American Economics Review*, 79(5), 1132-1145.
- Bernanke, B.S (2004), "Oil and the Economy," Speech presented at Darton College, Albany, GA, <http://www.federalreserve.gov/boarddocs/speeches/2004/20041021/default.htm>
- Brennan, M.J. (1958), "The Supply of Storage," *American Economic Review*, 48(1), 50-72.
- Britto, R. (1984), "The Simultaneous Determination of Spot and Futures Prices in a Simple Model with Production Risk," *Quarterly Journal of Economics*, 99(2), 351-365
- Chernenko, S.V., K.B. Schwarz, and J.H. Wright (2004), "The Information Content of Forward and Futures Prices," *International Finance Discussion Paper No. 808*.
- Chinn, M.D., M. LeBlanc, and O. Coibion (2005), "The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline," NBER Working Paper No. 11033.
- Chinn, M.D. and J.A. Frankel (1995), "Patterns on Exchange Rates for 25 Currencies," *Journal of Money, Credit, and Banking*, 26, 759-770.
- Clark, T.E. and K.D. West (2006), "Using Out-of-Sample Mean Squared Prediction Errors to Test the Martingale Difference Hypothesis," *Journal of Econometrics*, 135(1-2), 155-186.
- Commodity Futures Trading Commission (2007), "Backgrounder: The Commitments of Traders Report," <http://www.cftc.gov/opa/backgrounder/opacot596.htm>.
- Davies, P. (2007), "What's the Value of an Energy Economist?" Speech presented at the International Association for Energy Economics, Wellington, New Zealand.
- Diebold, F.X. and R.S. Mariano (1995), "Comparing Predictive Accuracy," *Journal of Business and Economic Statistics*, 13(3), 253-263.
- Duffie, D. (1989), *Futures Markets*. Prentice Hall, New Jersey.
- Elliott, G., and A. Timmermann (2007), "Economic Forecasting," forthcoming: *Journal of Economic Literature*.

- Elliott, G., I. Komunjer, and A. Timmermann, 2005, "Estimation and Testing of Forecast Rationality under Flexible Loss," *Review of Economic Studies*, 72, 1107-1125.
- Fama, E.F. and K.R. French (1987), "Commodity Futures Prices: Some Evidence on Forecast Power, Premia, and the Theory of Storage," *Journal of Business*, 60(1), 55-73.
- Fama, E.F. and K.R. French (1988), "Business Cycles and the Behavior of Metals Prices," *Journal of Finance*, 43(5), 1075-1093.
- Gorton, G. and K.G. Rouwenhorst (2006), "Facts and Fantasies about Commodity Futures," *Financial Analysts Journal*, 62(2), 47-68.
- Gramlich, E.M. (2004), Oil Shocks and Monetary Policy, Annual Economic Luncheon, Federal Reserve Bank of Kansas City, Kansas City, Missouri.
- Greenspan, A. (2004a), "Oil," Speech presented at the National Italian American Foundation, Washington, DC.
<http://www.federalreserve.gov/boarddocs/speeches/2004/200410152/default.htm>
- Greenspan, A. (2004b), "Energy," Remarks by Chairman Alan Greenspan Before the Center for Strategic & International Studies, Washington, D.C.
<http://www.federalreserve.gov/boarddocs/speeches/2004/20040427/default.htm>
- Haigh, M.S., Harris, J.H., Overdahl, J.A. and Robe, M.A. (2007), "Market Growth, Trader Participation and Derivative Pricing," mimeo, Kogod School of Business, American University, see <http://ssrn.com/abstract=966692>.
- Hotelling, H. (1931), "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39(2), 137-175.
- Hull, J.C. (2006), *Options, Futures, and Other Derivatives*, Pearson Prentice Hall.
- Inoue, A., and L. Kilian (2004), "In-Sample or Out-of-Sample Tests of Predictability: Which One Should We Use?" *Econometric Reviews*, 23(4), 371-402.
- International Monetary Fund (2007), *World Economic Outlook*, Washington, DC.
- Kilian, L. (2007a), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," mimeo, Department of Economics, University of Michigan.
- Kilian, L. (2007b), "The Economic Effects of Energy Price Shocks," invited for publication in the *Journal of Economic Literature*.
- Kilian, L. (2007c), "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" forthcoming: *Review of Economics and Statistics*.

- Neuberger, A. (1999), "Hedging Long-Term Exposures with Multiple Short-Term Futures Contracts," *Review of Financial Studies*, 12(3), 429-459.
- New York Mercantile Exchange (2007a), "Light Sweet Crude Oil: Market Data," http://www.nymex.com/lSCO_fut_descri.aspx.
- New York Mercantile Exchange (2007b), "Light Sweet Crude Oil: Contract Detail," http://www.nymex.com/WS_spec.aspx.
- Pagano, P., and M. Pisani (2006), "Risk-Adjusted Forecasts of Oil Prices," Discussion Paper No. 585, Banca D'Italia.
- Pesaran, M.H. and A. Timmermann (1992), "A Simple Nonparametric Test of Predictive Performance," *Journal of Business and Economic Statistics*, 10(4), 461-465.
- 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 and Futures Markets: A Primer," *Energy Journal*, 22(3), 1-29.
- Surowiecki, J. (2007), "Troubled Waters over Oil," *The New Yorker*, February 19-26, 64.
- Svensson, L.E.O. (2005), "Oil Prices and ECB Monetary Policy," mimeo, Department of Economics, Princeton University, see <http://www.princeton.edu/svensson/>
- Telser, L.G. (1958), "Futures Trading and the Storage of Cotton and Wheat," *Journal of Political Economy*, 66, 233-255.
- Townsend, R.M. (1978), "On the Optimality of Forward Markets," *American Economic Review*, 68(1), 54-66.
- Varian, H.R. (1985), "A Divergence of Opinion in Complete Markets: A Note," *Journal of Finance*, 40(1), 309-317.
- Wu, T., and A. McCallum (2005), "Do Oil Futures Prices Help Predict Future Oil Prices?" *Federal Reserve Bank of San Francisco Economic Letter*, 2005-38.

Table 1: 1-Month Ahead Recursive Forecast Error Diagnostics

$\hat{S}_{t+1 t}$	MSPE (<i>p</i> -value)	Bias	MAP E (<i>p</i> -value)	Success Ratio (<i>p</i> -value)
S_t	6.998	0.172	1.941	N.A.
$F_t^{(1)}$	7.106 (0.809)	0.210	1.949 (0.770)	0.443 (0.898)
$S_t(1 + (\hat{\alpha} + \hat{\beta} \ln(F_t^{(1)} / S_t)))$	6.994 (0.352)	0.157	1.954 (0.822)	0.479 (0.529)
$S_t(1 + \hat{\beta} \ln(F_t^{(1)} / S_t))$	6.975 (0.246)	0.156	1.960 (0.742)	0.423 (0.984)
$S_t(1 + \ln(F_t^{(12)} / S_t))$	7.106 (0.807)	0.212	1.949 (0.776)	0.443 (0.898)
$S_t(1 + \hat{\alpha})$	7.013 (0.527)	0.186	1.945 (0.525)	0.479 (0.497)
$S_t(1 + \Delta \bar{s}_t^{(1)})$	14.131 (0.450)	-0.164	2.619 (0.946)	0.490 (0.646)
$S_t(1 + \Delta \bar{s}_t^{(3)})$	10.061 (0.828)	-0.089	2.216 (0.750)	0.526 (0.265)
$S_t(1 + \Delta \bar{s}_t^{(6)})$	8.388 (0.797)	-0.106	2.042 (0.570)	0.505 (0.499)
$S_t(1 + \Delta \bar{s}_t^{(9)})$	8.238 (0.903)	-0.131	2.041 (0.696)	0.505 (0.509)
$S_t(1 + \Delta \bar{s}_t^{(12)})$	7.488 (0.281)	-0.141	1.929	0.546 (0.130)

Notes: The forecast evaluation period is 1991.1-2007.2. The initial estimation window is 1983.4-1990.12. For regressions based on 6-month futures prices the estimation window begins in 1983.10; for the 9-month futures price in 1986.12; for the 12-month futures price in 1989.1. $F_t^{(h)}$ is the futures price that matures in h periods; $i_{t,m}$ is the m month interest rate; and $\Delta \bar{s}_t^{(l)}$ denotes the trailing geometric average of the monthly percent change for l months. *p*-values are in parentheses. All *p*-values refer to pairwise tests of the null of a random walk without drift. Comparisons of nonnested models without estimated parameters are based on the *DM*-test of Diebold and Mariano (2005) with $N(0,1)$ critical values. Nested model comparisons with estimated parameters are based on Clark and West (2006). For the rolling regression estimates of the random walk with drift we use $N(0,1)$ critical values under quadratic loss; for recursive estimates under quadratic loss and for all estimates under absolute loss we use bootstrap critical values as described in Clark and West. The sign test in the last column is based on Pesaran and Timmermann (1992).

Table 2: 3-Month Ahead Recursive Forecast Error Diagnostics

$\hat{S}_{t+3 t}$	MSPE (<i>p</i> -value)	Bias	MAPE (<i>p</i> -value)	Success Ratio (<i>p</i> -value)
S_t	19.560	0.435	3.099	N.A.
$F_t^{(3)}$	19.038 (0.347)	0.631	3.172 (0.920)	0.479 (0.648)
$S_t(1+i_{t,3})$	19.811 (0.715)	0.167	3.111 (0.632)	0.541 N.A.
$S_t(1+(\hat{\alpha} + \hat{\beta} \ln(F_t^{(3)} / S_t)))$	24.217 (0.948)	0.253	3.610 (0.999)	0.407 (0.996)
$S_t(1 + \hat{\beta} \ln(F_t^{(3)} / S_t))$	22.826 (0.999)	0.804	3.541 (0.999)	0.407 (0.992)
$S_t(1 + \ln(F_t^{(12)} / S_t))$	19.036 (0.348)	0.649	3.176 (0.920)	0.479 (0.648)
$S_t(1 + \hat{\alpha})$	19.699 (0.493)	0.484	3.121 (0.500)	0.484 (0.716)
$S_t(1 + \Delta \bar{s}_t^{(1)})$	27.876 (0.684)	0.106	3.618 (0.816)	0.510 (0.418)
$S_t(1 + \Delta \bar{s}_t^{(3)})$	24.949 (0.947)	0.124	3.451 (0.829)	0.515 (0.376)
$S_t(1 + \Delta \bar{s}_t^{(6)})$	22.312 (0.860)	0.096	3.224 (0.669)	0.505 (0.509)
$S_t(1 + \Delta \bar{s}_t^{(9)})$	20.333 (0.510)	0.105	3.098 (0.545)	0.567 (0.044)
$S_t(1 + \Delta \bar{s}_t^{(12)})$	20.093 (0.439)	0.103	3.063 (0.459)	0.572 (0.032)
$S_{t,3}^{CF}$	30.726 (0.997)	-1.905	4.148 (0.999)	0.500 (0.338)

Notes: See Table 1.

Table 3: 6-Month Ahead Recursive Forecast Error Diagnostics

$\hat{S}_{t+6 t}$	MSPE (<i>p</i> -value)	Bias	MAP E (<i>p</i> -value)	Success Ratio (<i>p</i> -value)
S_t	34.058	0.937	4.466	N.A.
$F_t^{(6)}$	36.334 (0.716)	1.615	4.608 (0.906)	0.485 (0.483)
$S_t(1+i_{t,6})$	34.906 (0.713)	0.382	4.509 (0.708)	0.557 N.A.
$S_t(1+(\hat{\alpha} + \hat{\beta}\ln(F_t^{(6)} / S_t)))$	51.809 (0.926)	1.012	5.315 (0.960)	0.485 (0.983)
$S_t(1+\hat{\beta}\ln(F_t^{(6)} / S_t))$	47.143 (0.997)	1.959	5.200 (0.997)	0.464 (0.703)
$S_t(1+\ln(F_t^{(12)} / S_t))$	36.475 (0.721)	1.684	4.621 (0.910)	0.485 (0.483)
$S_t(1+\hat{\alpha})$	33.942 (0.486)	1.093	4.678 (0.487)	0.515 (0.022)
$S_t(1+\Delta\bar{s}_t^{(1)})$	46.176 (0.769)	0.418	4.912 (0.828)	0.505 (0.501)
$S_t(1+\Delta\bar{s}_t^{(3)})$	41.811 (0.851)	0.481	4.725 (0.858)	0.495 (0.641)
$S_t(1+\Delta\bar{s}_t^{(6)})$	36.154 (0.622)	0.550	4.520 (0.752)	0.521 (0.374)
$S_t(1+\Delta\bar{s}_t^{(9)})$	33.848 (0.293)	0.585	4.372 (0.988)	0.557 (0.091)
$S_t(1+\Delta\bar{s}_t^{(12)})$	34.372 (0.413)	0.588	4.450 (0.596)	0.515 (0.450)

Notes: See Table 1.

Table 4: 9-Month Ahead Recursive Forecast Error Diagnostics

\hat{S}_{t+9t}	MSPE (<i>p</i> -value)	Bias	MAPE (<i>p</i> -value)	Success Ratio (<i>p</i> -value)
S_t	46.574	1.791	5.161	N.A.
$F_t^{(9)}$	53.798 (0.887)	2.892	5.370 (0.926)	0.526 (0.080)
$S_t(1 + (\hat{\alpha} + \hat{\beta} \ln(F_t^{(9)} / S_t)))$	54.225 (0.802)	2.515	5.406 (0.684)	0.546 (0.035)
$S_t(1 + \hat{\beta} \ln(F_t^{(9)} / S_t))$	54.939 (0.901)	3.163	5.411 (0.850)	0.536 (0.026)
$S_t(1 + \ln(F_t^{(12)} / S_t))$	54.642 (0.898)	3.017	5.403 (0.948)	0.526 (0.080)
$S_t(1 + \hat{\alpha})$	46.107 (0.526)	2.090	5.150 (0.528)	0.600 (0.000)
$S_t(1 + \Delta \bar{s}_t^{(1)})$	60.122 (0.849)	1.284	5.629 (0.883)	0.495 (0.611)
$S_t(1 + \Delta \bar{s}_t^{(3)})$	51.476 (0.638)	1.368	5.248 (0.725)	0.536 (0.210)
$S_t(1 + \Delta \bar{s}_t^{(6)})$	46.268 (0.301)	1.437	5.101 (0.561)	0.515 (0.427)
$S_t(1 + \Delta \bar{s}_t^{(9)})$	45.204 (0.162)	1.463	5.066 (0.312)	0.515 (0.442)
$S_t(1 + \Delta \bar{s}_t^{(12)})$	45.949 (0.270)	1.460	5.123 (0.351)	0.490 (0.729)

Notes: See Table 1.

Table 5: 12-Month Ahead Recursive Forecast Error Diagnostics

$\hat{S}_{t+12 t}$	MSPE (<i>p</i> -value)	Bias	MAPE (<i>p</i> -value)	Success Ratio (<i>p</i> -value)
S_t	65.978	2.540	5.885	N.A.
$F_t^{(12)}$	77.204 (0.898)	4.009	6.212 (0.767)	0.536 (0.021)
$S_t(1+i_{t,12})$	65.285 (0.408)	1.439	6.018 (0.804)	0.582 N.A.
$S_t(1+(\hat{\alpha} + \hat{\beta}\ln(F_t^{(12)} / S_t)))$	78.414 (0.849)	3.874	6.272 (0.768)	0.526 (0.032)
$S_t(1+\hat{\beta}\ln(F_t^{(12)} / S_t))$	84.275 (0.967)	4.352	6.411 (0.950)	0.541 (0.004)
$S_t(1+\ln(F_t^{(12)} / S_t))$	79.007 (0.916)	4.189	6.279 (0.789)	0.536 (0.021)
$S_t(1+\hat{\alpha})$	64.709 (0.502)	3.200	5.968 (0.562)	0.552 (0.001)
$S_t(1+\Delta\bar{s}_t^{(1)})$	72.181 (0.294)	2.095	6.179 (0.93)	0.505 (0.002)
$S_t(1+\Delta\bar{s}_t^{(3)})$	68.919 (0.469)	2.148	6.053 (0.941)	0.515 (0.421)
$S_t(1+\Delta\bar{s}_t^{(6)})$	65.409 (0.291)	2.204	5.955 (0.918)	0.459 (0.932)
$S_t(1+\Delta\bar{s}_t^{(9)})$	64.500 (0.197)	2.225	5.920 (0.862)	0.448 (0.967)
$S_t(1+\Delta\bar{s}_t^{(12)})$	64.488 (0.191)	2.232	5.899 (0.622)	0.490 (0.746)
$S_{t,12}^{CF}$	107.866 (0.979)	-4.808	6.957 (0.954)	0.515 (0.122)

Notes: See Table 1.

Table 6: Asymptotic p -Values for Forecast Efficiency Regressions

Horizon	$H_0 : \beta = 1$	$H_0 : \alpha = 0, \beta = 1$
3-month	0.339	0.456
6-month	0.543	0.186
12-month	0.491	0.106

Notes: For the 3- and 6-month regressions, the sample period is 1989.4-2003.12. For the 12-month regression, the sample is 1990.1-2003.12. All t - and $Wald$ -tests have been computed based on HAC standard errors.

Table 7: Time Series Features of the Random Walk Risk Premium

	3 Month	12 Month
Mean	4.68	4.36
<i>t</i> -stat. (p-value)	2.97 (0.00)	3.28 (0.00)
Mean Abs. Dev	11.00	7.78
Min	-71.2	-32.2
Max	41.0	31.3
Persistence	0.68	0.84

Notes: The sample for the Consensus-based risk premia is 1989.10-2007.1; the sample for the 3-month random walk and Hotelling-based risk premia is 1983.7-2007.1; and that for the 12-month random walk and Hotelling-based risk premia is 1989.1-2007.1, reflecting the data constraints. The HAC *t*-statistic is for $H_0 : \hat{\rho}_t^{(h)} = 0$. The measure of persistence is the sum of the autoregressive coefficients proposed by Andrews and Chen (1994).

Table 8: Contemporaneous Correlation of Random Walk Risk Premium and Precautionary Demand Component of Spot Price of Crude Oil (Percent)

Horizon	1989.1-2001.12	2002.1-2006.12
3	62.1	-42.6
6	73.4	-50.4
9	77.8	-52.6
12	79.9	-50.2

NOTES: Computed based on Figure 5.

Figure 1: Why $F_t^{(h)} > S_t$ Does Not Imply Rising Spot Prices

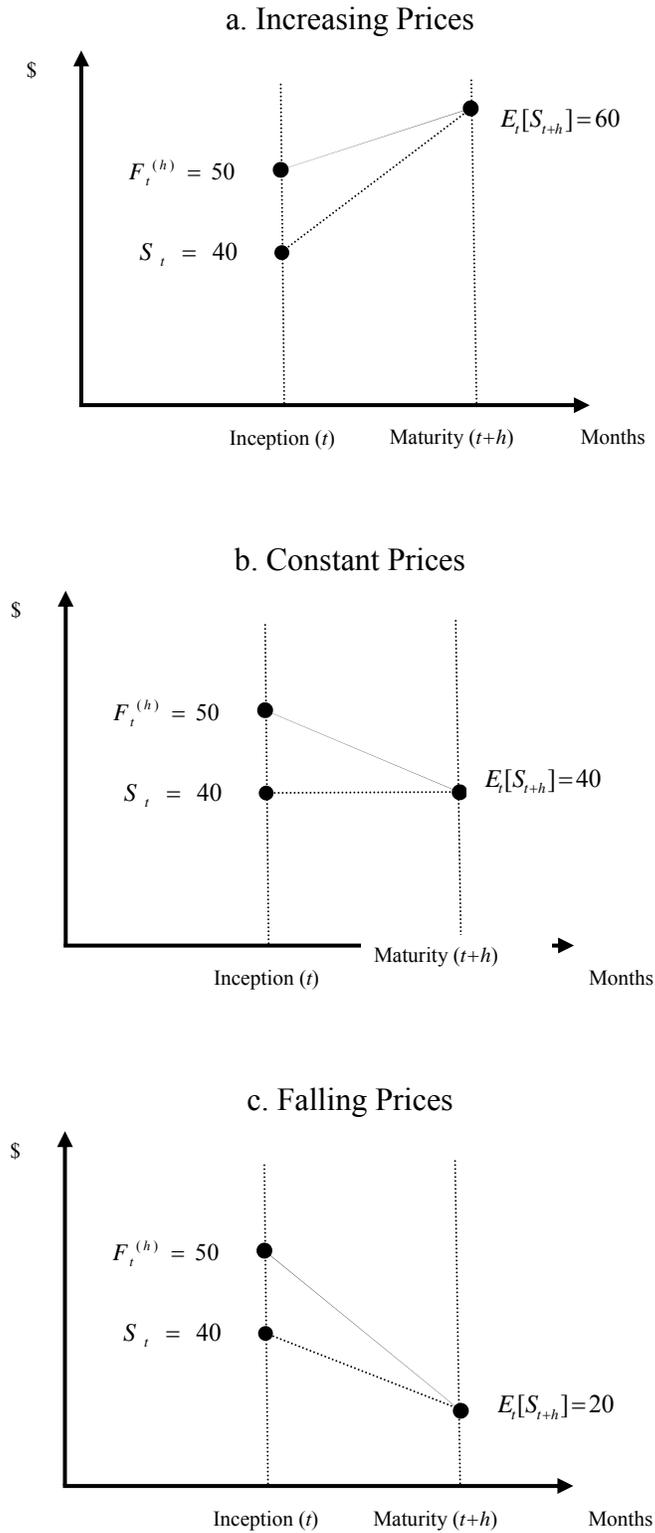
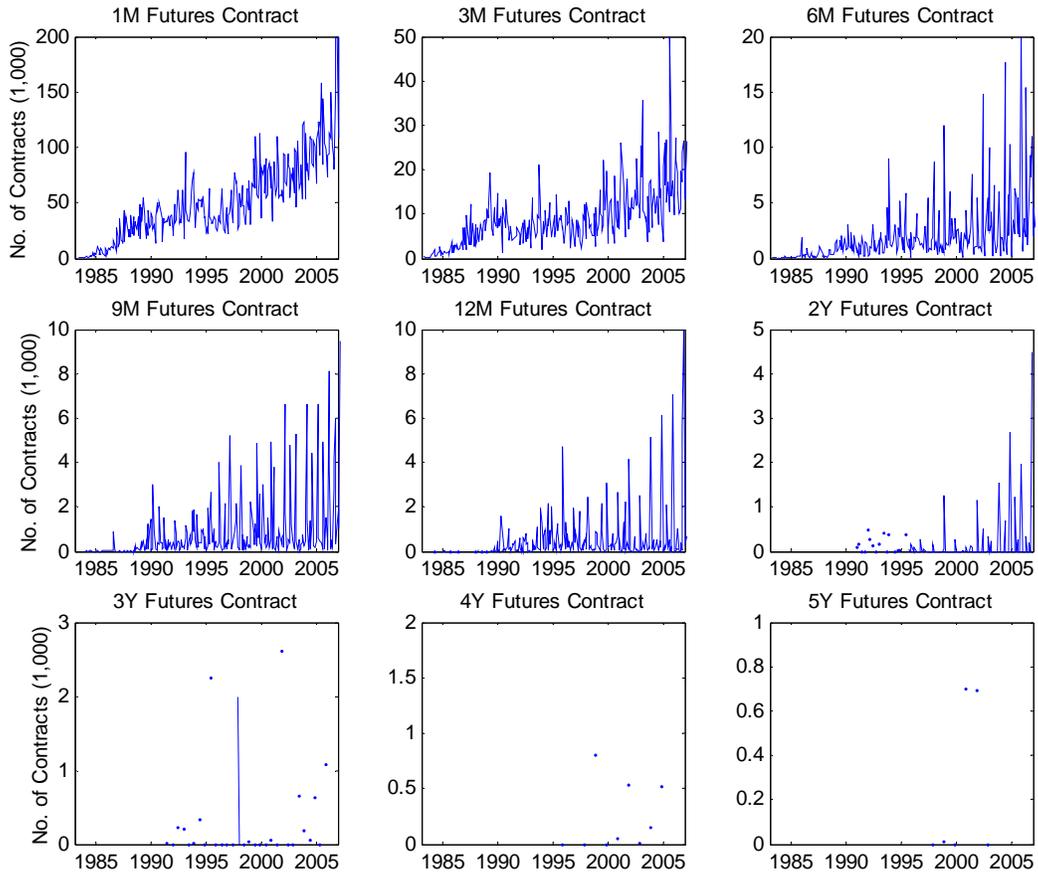
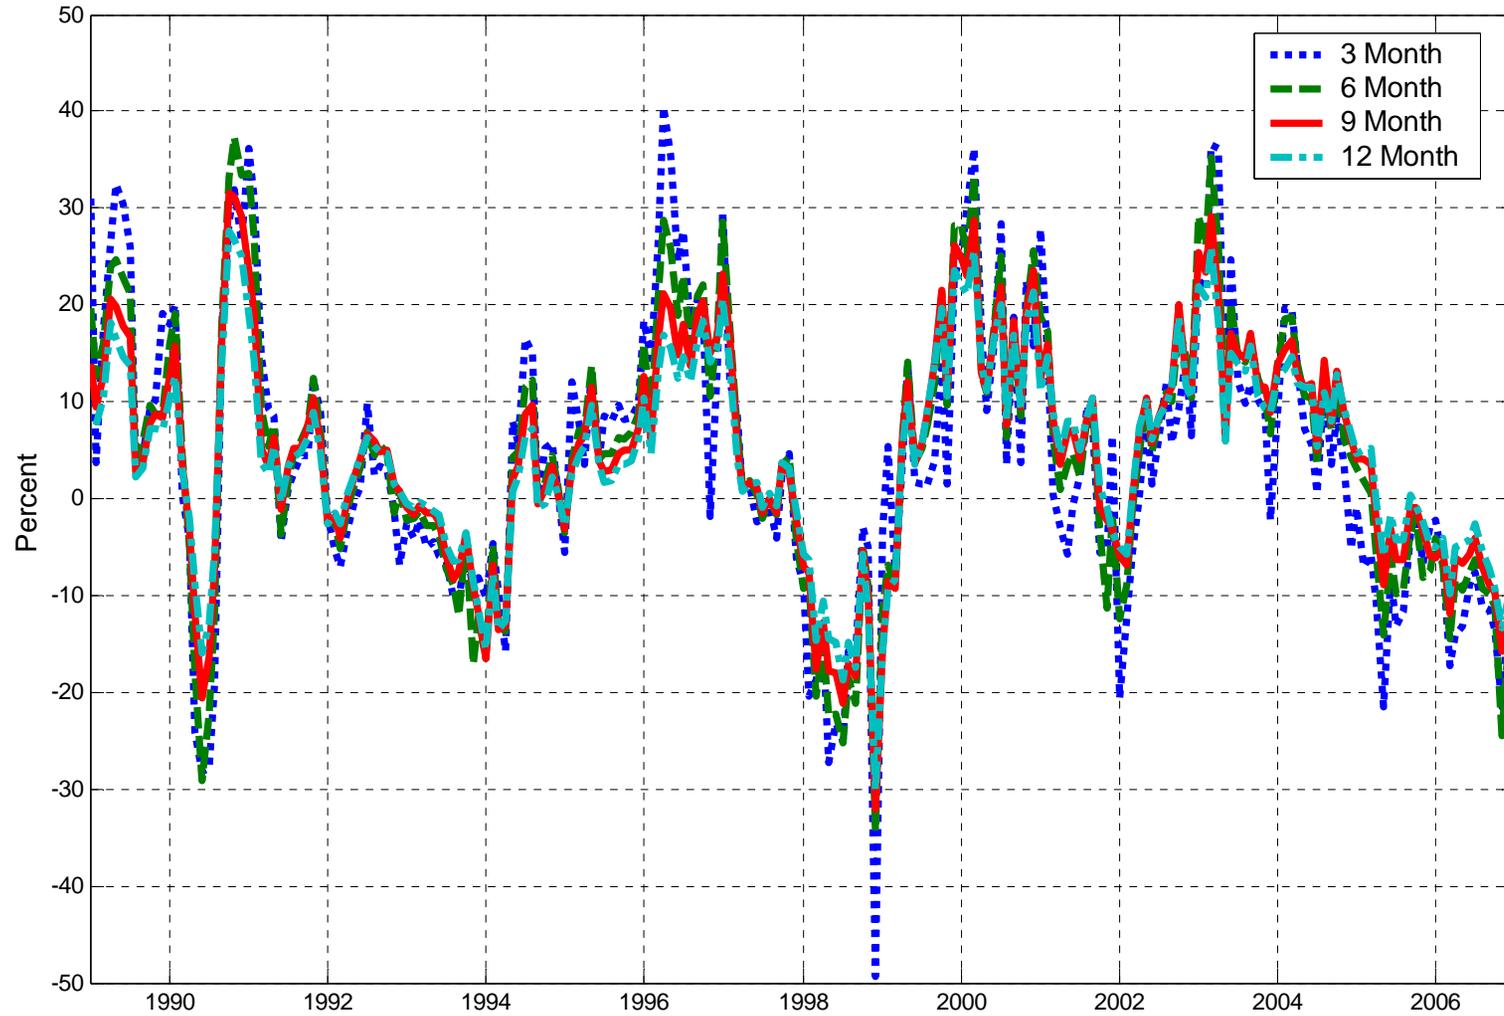


Figure 2: Volume of NYMEX Futures Contracts

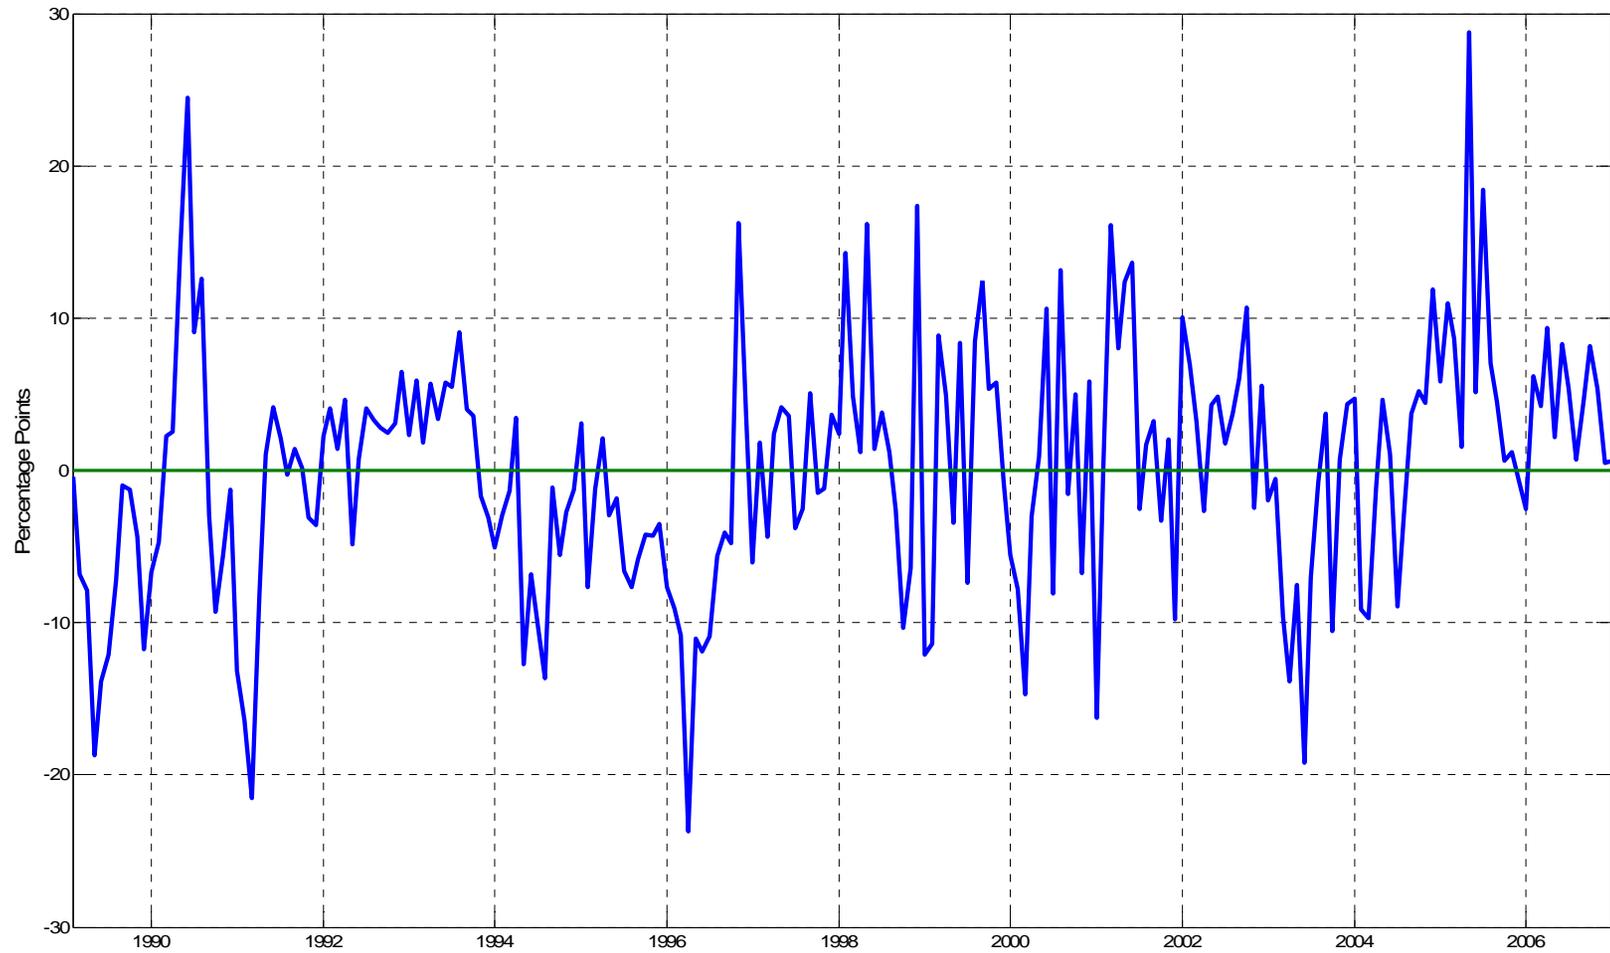


Source: *Price-data.com*

**Figure 3: Random Walk Risk Premium by Horizon
1989.1-2007.2**



**Figure 4: Term Structure of Random Walk Risk Premia
12-Month Risk Premium – 3-Month Risk Premium
1989.1-2007.2**



**Figure 5: Random Walk Risk Premium by Horizon and Precautionary Demand Component of Spot Price
1989.1-2006.12**

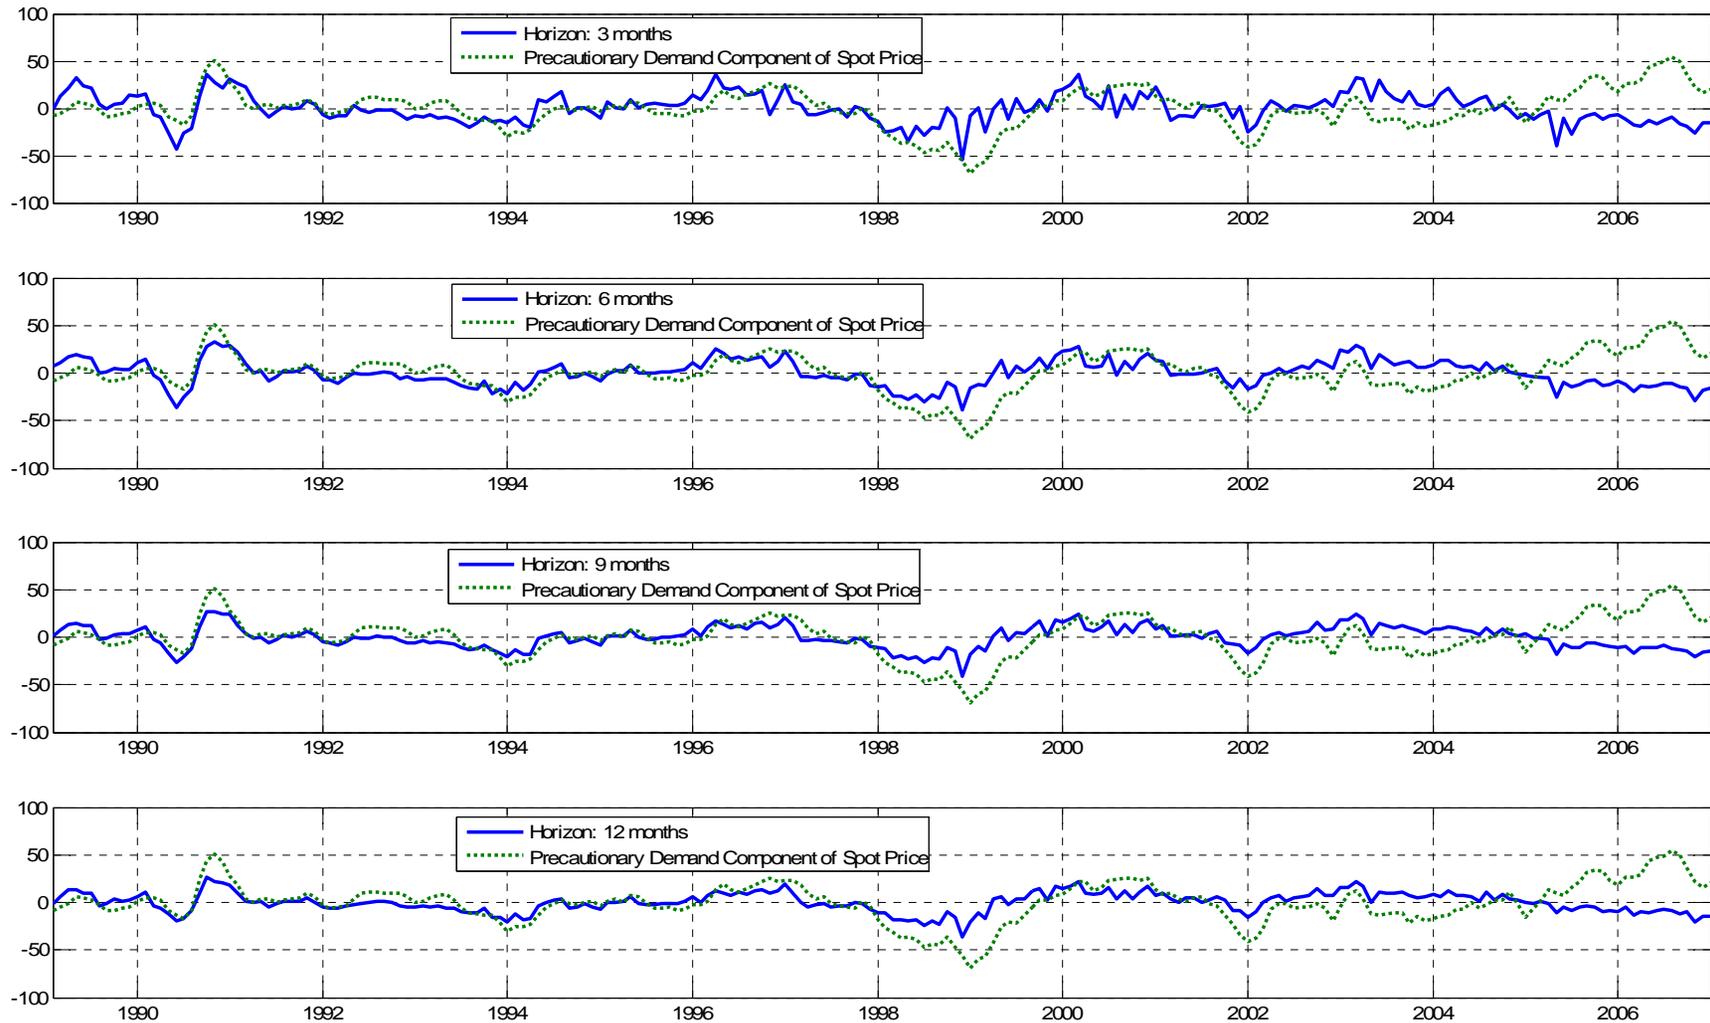
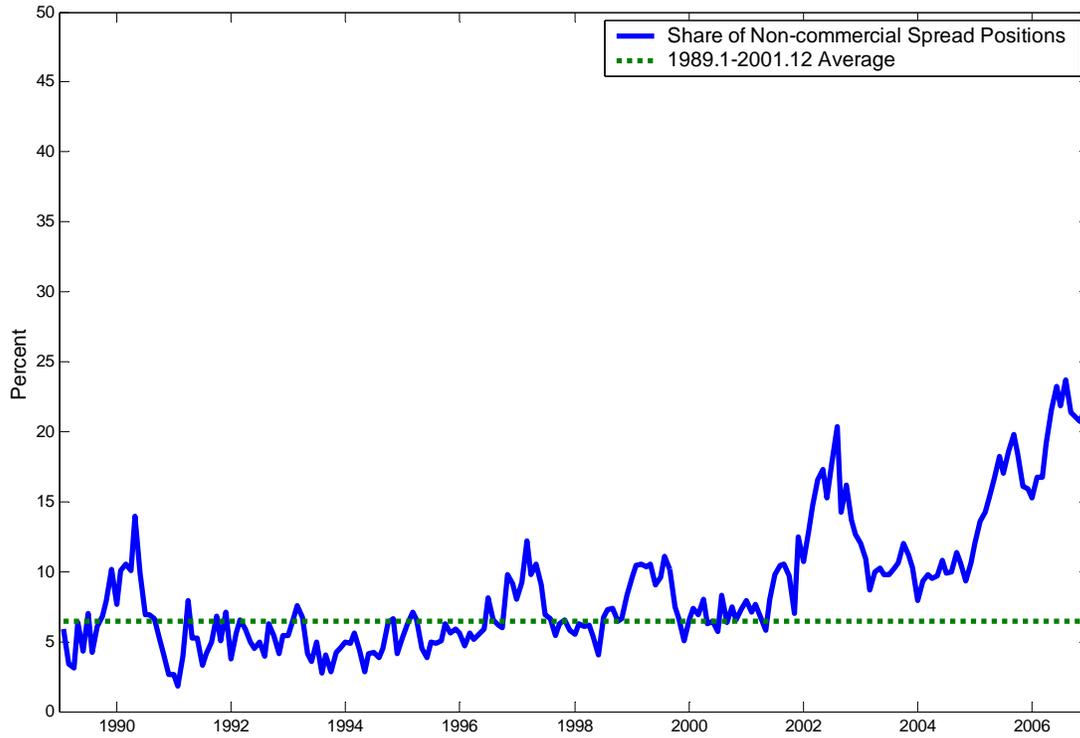


Figure 6: Share of Non-Commercial Spread Positions in Reportable Open Interest: Oil Futures Market 1989.1-2007.2



Notes: Data are end-of-month observations. Commercial firms are engaged in business activities that can be hedged in the futures market. Non-commercial firms are all other firms. Percentages refer to the share of reportable positions, which comprise 70-90% of all positions.

Appendix 1
CFTC Rule 1.3(z) Bona Fide Hedging Transactions and Positions¹¹

(z) *Bona fide hedging transactions and positions*—(1) *General definition*. Bona fide hedging transactions and positions shall mean transactions or positions in a contract for future delivery on any contract market, or in a commodity option, where such transactions or positions normally represent a substitute for transactions to be made or positions to be taken at a later time in a physical marketing channel, and where they are economically appropriate to the reduction of risks in the conduct and management of a commercial enterprise, and where they arise from:

- (i) The potential change in the value of assets which a person owns, produces, manufactures, processes, or merchandises or anticipates owning, producing, manufacturing, processing, or merchandising,
- (ii) The potential change in the value of liabilities which a person owns or anticipates incurring, or
- (iii) The potential change in the value of services which a person provides, purchases, or anticipates providing or purchasing.

Notwithstanding the foregoing, no transactions or positions shall be classified as bona fide hedging unless their purpose is to offset price risks incidental to commercial cash or spot operations and such positions are established and liquidated in an orderly manner in accordance with sound commercial practices and, for transactions or positions on contract markets subject to trading and position limits in effect pursuant to section 4a of the Act, unless the provisions of paragraphs (z) (2) and (3) of this section and §§1.47 and 1.48 of the regulations have been satisfied.

(2) *Enumerated hedging transactions*. The definitions of bona fide hedging transactions and positions in paragraph (z)(1) of this section includes, but is not limited to, the following specific transactions and positions:

- (i) *Sales* of any commodity for future delivery on a contract market which do not exceed in quantity:
 - (A) Ownership or fixed-price purchase of the same cash commodity by the same person; and
 - (B) Twelve months' unsold anticipated production of the same commodity by the same person provided that no such position is maintained in any future during the five last trading days of that future.
- (ii) *Purchases* of any commodity for future delivery on a contract market which do not exceed in quantity.

¹¹ See Electronic Code of Federal Regulations at <http://ecfr.gpoaccess.gov>.

- (A) The fixed-price sale of the same cash commodity by the same person.
- (B) The quantity equivalent of fixed-price sales of the cash products and by-products of such commodity by the same person; and
- (C) Twelve months' unfilled anticipated requirements of the same cash commodity for processing, manufacturing, or feeding by the same person, provided that such transactions and positions in the five last trading days of any one future do not exceed the person's unfilled anticipated requirements of the same cash commodity for that month and for the next succeeding month.
- (iii) Offsetting sales and purchases for future delivery on a contract market which do not exceed in quantity that amount of the same cash commodity which has been bought and sold by the same person at unfixed prices basis different delivery months of the contract market, provided that no such position is maintained in any future during the five last trading days of that future.
- (iv) Sales and purchases for future delivery described in paragraphs (z)(2)(i), (ii), and (iii) of this section may also be offset other than by the same quantity of the same cash commodity, provided that the fluctuations in value of the position for future delivery are substantially related to the fluctuations in value of the actual or anticipated cash position, and provided that the positions in any one future shall not be maintained during the five last trading days of that future.
- (3) *Non-enumerated cases.* Upon specific request made in accordance with §1.47 of the regulations, the Commission may recognize transactions and positions other than those enumerated in paragraph (z)(2) of this section as bona fide hedging in such amount and under such terms and conditions as it may specify in accordance with the provisions of §1.47. Such transactions and positions may include, but are not limited to, purchases or sales for future delivery on any contract market by an agent who does not own or who has not contracted to sell or purchase the offsetting cash commodity at a fixed price, *provided* That the person is responsible for the merchandising of the cash position which is being offset.