What Do We Learn from the Price of Crude Oil Futures?

Ron Alquist*
University of Michigan

Lutz Kilian
University of Michigan and CEPR

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

Abstract: Based on a two-country, multi-period general equilibrium model of the spot and futures markets for crude oil, we show that there is no theoretical support for the common view that oil futures prices are accurate predictors of the spot price in the mean-squared prediction error (MSPE) sense; yet under certain conditions there is support for the view that oil futures prices are unbiased predictors. Our empirical analysis documents that futures-based forecasts typically are less accurate than the no-change forecast and biased, although the bias is small. Much of the MSPE is driven by the variability of the futures price about the expected spot price, as captured by the basis. Empirically, the fluctuations in the oil futures basis are larger and more persistent than fluctuations in the basis of foreign exchange futures. Within the context of our theoretical model, this anomaly can be explained by the marginal convenience yield of oil inventories. We show that increased uncertainty about future oil supply shortfalls under plausible assumptions causes the basis to decline and precautionary demand for crude oil to increase, resulting in an immediate increase in the real spot price that is not necessarily associated with an accumulation of oil inventories. Our main result is that the negative of the basis may be viewed as an index of fluctuations in the price of crude oil driven by precautionary demand for oil. An empirical analysis of this index provides independent evidence of how shifts in market expectations about future oil supply shortfalls affect the spot price of crude oil. Such expectation shifts have been difficult to quantify, yet have been shown to play an important role in explaining oil price fluctuations. Our empirical results are consistent with related evidence in the literature obtained by alternative methodologies.

Key words: Crude oil; futures market; spot market; spread; 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, Lucas Davis, Ana-María Herrera, Steve Salant, and Ken West as well as seminar participants at the University of Michigan and the Federal Reserve Bank of Dallas 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

The surge in the price of crude oil since 2002 has renewed interest in the question of what determines the spot and futures price of crude oil and has highlighted the importance of being able to predict as accurately as possible the evolution of the spot price of oil (see, e.g., Greenspan 2004a,b, 2005; Bernanke 2004, 2006; Gramlich 2004; Davies 2007; Kohn 2007). In this paper, we use insights provided by a theoretical model of the spot and futures market for crude oil in conjunction with empirical analysis to shed light on the relationship between the spot price of crude oil, expectations of future oil prices, the price of crude oil futures, and the oil futures basis (defined as the percent deviation of the futures price from the spot price of oil).

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. Futures-based forecasts of the price of crude oil play an important role in monetary policy decisions and affect financial markets' perceptions of the risks to price stability and sustainable growth. It is widely believed that oil futures prices can be viewed as effective long-term supply prices (see, e.g., Greenspan 2004a) or as the expected price of oil (see, e.g., Bernanke 2004). We examine this common practice both from a theoretical and an empirical point of view.

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. Oil importers may insure against uncertainty about stochastic oil endowments by holding above-ground oil inventories or buying oil futures. Oil producers may sell oil futures to protect against endowment uncertainty. The model abstracts from oil below the ground. The spot and futures prices of oil are determined endogenously and simultaneously. Within this theoretical framework, we study the predictive power of futures prices for the spot price of oil. We show that there is no presumption that the price of oil futures will be an accurate predictor of the spot price of oil according to the conventional mean-squared prediction error (MSPE) metric, nor will it necessarily be unbiased. In the model, the current futures price equals the expected value of the spot price only under risk neutrality.

In section 4, we assess empirically whether forecasts based on the price of oil futures are more accurate than forecasts from benchmark models excluding futures prices. 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 and includes observations up to February of 2007, we show that forecasts based on futures prices and forecasts based on the futures spread in practice tend to be less accurate than forecasts from alternative easy-to-use models such as the no-change forecast under standard loss functions including the mean-squared prediction error (MSPE).

In section 5, we assess the unbiasedness of futures prices as predictors of spot prices. Unbiasedness may be assessed only with reference to an explicit measure of the expected spot price. A reasonable proxy for the expected spot price is the best possible predictor available to real-time forecasters. Hence, we conduct a systematic evaluation of the out-of-sample predictive accuracy of other forecasting methods based on the forecast evaluation period 1991.1-2007.2. The forecast accuracy comparison suggests that the no-change forecast (based on the random walk model without drift) is the best proxy for the expected spot price. A robust finding across all horizons from 1 month to 12 months is that the no-change forecast tends to be more accurate than forecasts based on other econometric models as well as professional forecasts of the price of crude oil such as the 3-month and 12-month survey forecasts provided by *Consensus Economics Inc*. Using the no-change forecast as the proxy for the expected spot price, we formally test and reject the unbiasedness hypothesis, although the bias is small. We obtain qualitatively similar results under the assumption of perfect foresight.

The result that futures prices are neither unbiased predictors nor the best possible predictors in the MSPE sense is new and surprising because it contradicts widely held views among policymakers and financial analysts. It also differs from some earlier empirical results in the academic literature based on shorter samples. Finally, it differs from related results in the foreign exchange literature. Although the no-change forecast has been shown to work well in other contexts such as exchange rate forecasting, there are important differences between the foreign exchange market and the crude oil market. Forecast efficiency regressions for oil markets generate the expected signs and magnitudes for all coefficients, whereas similar regressions for foreign exchange markets generate coefficients of the wrong sign and magnitude (see, e.g., Froot and Thaler 1990). Thus, the superiority of the random walk predictor of oil prices compared with futures prices is by no means expected.

Our analysis establishes that the cause of the large MSPE of futures-based forecasts is not the bias, but rather the variability of the futures price about the expected spot price, as captured by the basis. We document that there are large and persistent fluctuations in the oil futures basis that are unlike the fluctuations observed in the basis of foreign exchange futures (see, e.g., Taylor 1989). In section 6, we show that these differences can be linked to the existence of a marginal convenience yield for crude oil that is absent in foreign exchange markets. Oil inventories, unlike inventories of many financial assets, may serve to avoid interruptions of the production process or to meet unexpected shifts in demand. This option value is reflected in a convenience yield (see, e.g., Brennan 1991; Pindyck 1991, 2001). We study the implications of the marginal convenience yield for the basis in the context of the theoretical model of section 3. Using comparative statics, we formally establish that under plausible conditions increased uncertainty about future oil supply shortfalls cause the oil futures basis to fall. This fact helps explain the poor forecasting performance of futures prices and spreads. Such uncertainty shifts also raise the current spot price of oil, as precautionary demand for oil inventories increases in response to increased uncertainty. Hence, the model implies a negative correlation between the basis and the component of the spot price of oil driven by precautionary demand for crude oil. These results suggest that the negative of the oil futures basis may be viewed as an *index* of fluctuations in the spot price of crude oil driven by shifts in precautionary demand for oil.

Our empirical analysis suggests that the sharp spike in oil prices during the Persian Gulf War was associated with an expectations shift reflected in higher precautionary demand for crude oil. We also find evidence of shifts in the basis associated with the Asian Financial Crisis, with 9/11 and with the 2003 Iraq War, for example. Our findings corroborate earlier results in the literature based on regression dummies as well as historical decompositions derived from structural vector autoregressive models.

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 alternative estimate is based on a structural VAR model of the global crude oil market and does not rely on data from the oil futures market. We show that the VAR-based measure and the futures-based measure have a correlation of between 58 percent and 79 percent during 1989.1-2003.12, depending on the maturity of the futures data used. The correlation monotonically increases in the horizon, consistent with the view that precautionary demand is mainly concerned with risk beyond the short run. The close correlation of these two measures up to

2003.12 further corroborates the credibility of our results.

While the oil futures basis can be useful in identifying expectations shifts in general, our results suggest caution in interpreting the basis in the absence of further information about the market structure. Specifically, the ability of the basis to capture fluctuations in the spot price of oil driven by shifts in the precautionary demand for oil weakens after 2003.12. That weakening coincides with an unprecedented increase in speculative activities in the oil futures market. To the extent that increased speculative trading tends to raise futures prices more than spot prices (and hence raises the basis), this weakening is not unexpected, as our theoretical model does not allow for speculation. Establishing such a link is left for future research.

In section 7 we address the reasons for the lack of a tight relationship between crude oil inventories and the precautionary demand component of the spot price of oil. Using a three-period extension of the baseline model, we establish that increases in the precautionary demand for crude oil need not go hand in hand with the accumulation of oil inventories. The concluding remarks are in section 8.

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

It is commonplace in policy institutions, including many central banks and the International Monetary Fund (IMF), to use the price of NYMEX oil futures as a proxy for the market's expectation of the spot price of crude oil.[1] A widespread view is that prices of NYMEX futures contracts are not only good proxies for the expected spot price of oil, but also better predictors of oil prices than econometric forecasts. Forecasts of the spot price of oil are used, for example, as 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 2005, p. 67; 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).

[1] 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. Futures contracts can be retraded between inception and maturity on a futures exchange such as the New York Mercantile Exchange (NYMEX). 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).

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 sections 4 and 5), 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 spot 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 section 6.

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) who established that competitive futures and spot markets under weak conditions will be Pareto efficient as long as agents can trade in the spot market after the uncertainty about the state of the world is resolved. Our objective is to obtain a stylized model of trading in the spot market and futures market for crude oil. We adapt the framework used by Townsend as follows. First, we add production of a final good. In our model a representative agent for the oil-importing economy, the United States, uses oil as an input in producing a final consumption good. Second, we allow the United States to carry above-ground inventories of oil from the first period to the second period. Third, all agents have access to a bond market that allows them to save at a risk-free rate.[2] Adding production of a final consumption good is central to deriving some of the results in section 5. The introduction of inventories is necessary for characterizing the conditions for arbitrage between the spot price of oil and the futures price, also known as the cost-of-carry relationship. Finally, the bond market serves to pin down endogenously the interest rate used in

[2] 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.

deriving the cost-of-carry relationship.

3.1. Model Description

There are two countries, the United States and Saudi Arabia. Saudi Arabia trades its oil endowment 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 is treated as an endowment economy in recognition of the fact that capacity constraints have been binding in global crude oil production in recent years (see Kilian 2007c). The existence of capacity constraints implies that extracting less oil today does not permit more oil to be extracted in the future. Saudi Arabia's oil endowment in the first period is known to be ω. Saudi Arabia's second period oil endowment is uncertain. There are two states of nature. With probability θ, Saudi Arabia receives the random endowment $\omega + \varepsilon$ (state 1). With probability $1-\theta$, it receives the endowment $\omega - \hat{\varepsilon}$ (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 denoted σ_{ε}^2.

In the first period, the United States chooses: (1) above-ground inventory holdings of oil to carry forward to period 2; (2) the number of oil futures contracts that deliver one barrel of oil in period 2; (3) the number of risk-free bonds that yield $1+R$ in period 2, and (4) the quantity of oil to use in the production of the consumption good. Saudi Arabia chooses the number of oil futures contracts and the number of risk-free bonds it wishes to hold.

The optimization problem for the United States can be solved by backward induction. In period 2, the United States' problem is

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

where Z_{2s} is the quantity of oil the United States chooses for production in period 2 and state $s = 1,2$; $U(.)$ is the United States' utility function defined over the consumption good; $C_{US,2s}$ is the quantity of the consumption good available to the United States in period 2 and state s; $F(.)$ is the strictly concave production function that generates the consumption good from oil; $F' > 0, F'' < 0, F''' > 0$; I_{US} is the quantity of oil the United States holds as inventory; N_{US} is the number of oil futures contracts the United States buys in period 1 for delivery in period 2; 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 period 2 regardless of the state; R is the risk-free interest rate; and B_{US} is the number of bonds the United States holds. In period 2, I_{US}, N_{US}, and B_{US} are state variables inherited from period 1.

In period 1, the United States chooses the quantity of oil for immediate use, the amount of oil to store as inventories, the number of oil futures contracts, and the number of bonds to maximize its utility:

$$\max_{\{Z_1, I_{US}, N_{US}, B_{US}\}} U(C_{US,1}) + g(I_{US}, \sigma_\varepsilon^2) + \beta\left[\theta J_{21}(I_{US}, N_{US,} B_{US}) + (1-\theta) J_{22}(I_{US}, N_{US,} B_{US})\right] \text{ s.t.}$$
$$C_{US,1} = F(Z_1) - \frac{S_1}{P_1}(Z_1 + I_{US}) - \frac{B_{US}}{P_1},$$

where Z_1 is the quantity of oil the United States uses to produce in period 1; $C_{US,1}$ is the quantity of the consumption good available to the United States in period 1; $g(I_{US}, \sigma_\varepsilon^2)$ measures the convenience yield accruing to the United States from holding inventory between period one and period two;[3] it is an increasing function of inventories and a decreasing function of the variance of the oil endowment such that $g_1 > 0,\ g_{11} < 0,$ and $g_{12} > 0,$ where g_i denotes the derivative of g with respect to its *ith* argument and g_{ij} the cross-partial derivative of g with respect to the arguments i and j. Throughout the paper we postulate that the Inada condition

$$\lim_{I_{US}\to 0} g_1\left(I_{US}, \sigma_\varepsilon^2\right) = \infty$$

holds. This condition ensures that the U.S. holds strictly positive inventories. $\beta \in (0,1)$ is the subjective discount factor; S_1 is the spot price of oil in period 1; P_1 is the price of the consumption good in period 1; and J_{ts} is the U.S. value function in period t and state s.

Saudi Arabia's decision problem can be solved in the first period, because it faces no decision in period 2. Saudi Arabia's problem is

$$\max_{\{N_{SA}, B_{SA}\}} V(C_{SA,1}) + \beta\left[\theta V(C_{21}) + (1-\theta) V(C_{22})\right] \text{ s.t.}$$

[3] The term *convenience yield* in the literature refers to the benefits arising from access to crude oil in the form of inventories such as the ability to avoid disruptions of the production process or the ability to meet unexpected demand for the final good. As increases in the variance make production shortfalls more likely, the marginal convenience yield from holding inventories is increasing in the variance.

$$C_{SA,1} = \frac{S_1}{P_1}\omega - \frac{B_{SA}}{P_1}$$
$$C_{SA,21} = \frac{S_{21}}{P_{21}}(\omega+\varepsilon) + N_{SA}\left(\frac{S_{21}}{P_{21}} - \frac{F_1}{P_{21}}\right) + \frac{(1+R)}{P_{21}}B_{SA}$$
$$C_{SA,22} = \frac{S_{22}}{P_{22}}(\omega-\hat{\varepsilon}) + N_{SA}\left(\frac{S_{22}}{P_{22}} - \frac{F_1}{P_{22}}\right) + \frac{(1+R)}{P_{22}}B_{SA}$$

where $V(.)$ is Saudi Arabia's utility function defined over the consumption good; $C_{SA,1}$ is the quantity of the consumption good available to Saudi Arabia in period 1; $C_{SA,2s}$ is the quantity of the consumption good available to Saudi Arabia in period 2 and state $s = 1,2$; B_{SA} is the number of bonds Saudi Arabia holds; and N_{SA} is the number of futures contracts sold.

These two optimization problems jointly yield 14 equations. There are six first-order conditions and three budget constraints for the United States; and there are two first-order conditions and three budget constraints for Saudi Arabia. These 14 equations define the optimal choices of the two agents for any price vector. In addition, there are eight market-clearing conditions:

$$C_{SA,1} + C_{US,1} = F(Z_1)$$
$$C_{SA,21} + C_{US,21} = F(Z_{21})$$
$$C_{SA,22} + C_{US,22} = F(Z_{22})$$
$$Z_1 = \omega - I_{US}$$
$$Z_{21} = \omega + \varepsilon + I_{US}$$
$$Z_{22} = \omega + \hat{\varepsilon} + I_{US}$$
$$N_{US} + N_{SA} = 0$$
$$B_{US} + B_{SA} = 0$$

These 22 equations allow us to solve for the 22 endogenous variables (one of which is implied by the remaining 21 by Walras' Law), the six consumption variables $(C_{US,1}, C_{US,21}, C_{US,22}, C_{SA,1}, C_{SA,21}, C_{SA,22})$; inventory holdings (I_{US}); 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 $\left(S_1/P_1,\ F_1/P_1, S_{21}/P_{21}, S_{22}/P_{22}, F_1/P_{21}, F_1/P_{21}, R\right)$.

The model has direct implications for the relationship between the futures price and the expected spot price of crude oil. Suppose, for expository purposes, that the Saudi and U.S. utility functions are identical and consider the United States' decision problem. Taking the first-order condition of $U\left(C_{US,2s}\right)$ with respect to U.S. holdings of futures contracts (N_{US}), we obtain

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

Rearranging this expression as

$$\left[\theta\frac{S_{21}}{P_{21}}U'(C_{US,21})+(1-\theta)\frac{S_{22}}{P_{22}}U'(C_{US,22})\right]=\left[\theta\frac{F_1}{P_{21}}U'(C_{US,21})+(1-\theta)\frac{F_1}{P_{22}}U'(C_{US,22})\right]$$

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

(1) $$F_1=\frac{\operatorname{cov}\left(U'\left(C_{US,2}\right),\frac{S_2}{P_2}\right)}{E\left[\frac{U'\left(C_{US,2}\right)}{P_2}\right]}+\frac{E\left[U'\left(C_{US,2}\right)\right]E\left[\frac{S_2}{P_2}\right]}{E\left[\frac{U'\left(C_{US,2}\right)}{P_2}\right]}.$$

An analogous expression could be derived for Saudi Arabia. In equilibrium, the futures price adjusts to ensure that both of these expressions are equal. From the perspective of modern asset pricing theory, the first term in expression (1) may be interpreted as a *risk premium*. Clearly, under risk aversion there is no simple linear relationship between the futures price and the spot price nor is there any reason to believe that $F_1=E\left[S_2\right].$ Under risk neutrality, the covariance term is zero, and based on a Taylor series approximation about the mean we obtain:

(1′) $$F_1\approx E\left[S_2\right] \quad ,$$

Thus, the futures price will be an approximately unbiased predictor of the spot price. Even under risk neutrality, however, the model provides no assurances that the futures price will be an accurate predictor in the MSPE sense. In the next two sections, we will explore both the ability of oil futures prices to forecast spot prices and the extent to which these forecasts are unbiased. In section 6, 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 basis over time.

4. Do Oil Futures Prices Help Predict the Spot Price of Oil?

4.1. Forecasting Models

4.1.1. The Benchmark Model

Let $F_t^{(h)}$ denote the current nominal 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*. A natural benchmark for forecasts based on the price of oil futures is provided by the random walk model without drift. This model implies that changes in the spot price are unpredictable, so the best forecast of the future spot price of crude oil is simply the current spot price:

$$\hat{S}_{t+h|t} = S_t \qquad h = 1, 3, 6, 9, 12 \tag{2}$$

Below we consider two types of forecasting models based on the price of oil futures. The first model simply treats the current level of futures prices as the predictor; the second model is based on the futures spread.

4.1.1. Futures Prices as Future Spot Prices

The Greenspan (2004a) quote of the introduction implies the forecasting model:

$$\hat{S}_{t+h|t} = F_t^{(h)} \qquad h = 1, 3, 6, 9, 12\,. \tag{3}$$

4.1.2. 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 (see, e.g., Gramlich 2004). The rationale for this approach is clear from dividing expression (1′) by S_t. If the futures price equals the expected spot price, the spread should be an indicator of the expected change in spot prices, although not necessarily an accurate predictor of the change in spot prices in the MSPE sense. We will explore the forecasting accuracy of the spread based on a variety of alternative forecasting models. The baseline model is:

$$\hat{S}_{t+h|t} = S_t\left(1 + \ln(F_t^{(h)}/S_t)\right), \qquad h = 1, 3, 6, 9, 12 \tag{4}$$

To allow for the possibility that the spread may be a biased predictor, it is common to relax the assumption of a zero intercept:

$$(5) \qquad \hat{S}_{t+h|t} = S_t\left(1+\hat{\alpha}+\ln(F_t^{(h)}/S_t)\right), \qquad h = 1, 3, 6, 9, 12$$

Alternatively, one can relax the proportionality restriction:

$$(6) \qquad \hat{S}_{t+h|t} = S_t\left(1+\hat{\beta}\ln(F_t^{(h)}/S_t)\right), \qquad h = 1, 3, 6, 9, 12$$

Finally, we can relax both the unbiasedness and proportionality restrictions:

$$(7) \qquad \hat{S}_{t+h|t} = S_t\left(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(h)}/S_t)\right), \qquad h = 1, 3, 6, 9, 12\,.$$

4.2. Data Description and Timing Conventions

4.2.1. Data Construction

In section 4.3, we will compare the real-time forecast accuracy of models (2)-(7). Our empirical 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. Contracts are for delivery at Cushing, OK. Trading ends 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).

A common problem in constructing monthly futures prices of a given maturity 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 in a given month (see Chernenko, Schwarz, and Wright 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. We identify the *h*-month futures contract trading closest to the last trading day of the month and use the price associated with that 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 is at most three days. Our approach is motivated by the objective of computing in a consistent manner end-of-month time series of futures prices for different maturities. This allows us to match up end-of-month spot prices and futures prices as closely as possible. The daily spot price data are obtained from *Datastream* and refer to the price of West Texas Intermediate crude oil available for delivery at Cushing, OK. Figure 1 plots the monthly prices of oil futures contracts for maturities of 1 through 12 months and the spot price of crude oil starting in 1983.1. Note that contracts of longer maturities only gradually became available over the course of the sample period.

4.2.2. 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 are 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. If the market were effectively pricing the possibility of, say, a shutdown of the Iranian oil fields or the demise of the Saudi monarchy within the next five years, one would expect active trading at such 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.[4] 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 contracts exceeding twelve months and in 1999, a 7-year contract was first introduced. Although such contracts are available, the market remains illiquid at horizons beyond one year even in recent years. Trading volumes fall 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.

---

[4] 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.

In addition, the evidence in Figure 2 suggests that the public and policymakers have 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. Policymakers routinely rely on futures prices for long maturities in predicting future oil prices. For example, Greenspan (2004a), in referring to the 6-year oil futures contract, states that "… [oil] futures prices at that horizon can be viewed as effective long-term supply prices." For similar statements also see Greenspan (2004b), Gramlich (2004) and Bernanke (2004). 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.[5]

4.3. Real-Time Forecast Accuracy of Futures-Based Forecasting Models

Tables 1 through 5 assess the predictive accuracy of various forecasting models against the benchmark of a random walk without drift for horizons of 1, 3, 6, 9, and 12 months. The forecast evaluation period is 1991.1-2007.2. The assessment of which forecasting model is best may depend on the loss function of the forecaster (see Elliott and Timmermann 2007). We present results for the 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.

4.3.1. Oil Futures as Predictors of Oil Spot Prices

The first two rows of 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 is the futures forecast 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

[5] 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.

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 is the success ratio significant at the 10 percent level and 5 percent level, respectively, but the improvement is only 2.6 and 3.6 percentage points. 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, as exemplified by the Greenspan (2004a) quote.

4.3.2. Oil Future Spreads as Predictors of Future Spot Prices

Rows 3-6 in Tables 1-5 document 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 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 rankings, the no-change forecast is superior at all horizons. These results are broadly consistent with the earlier evidence for the futures forecasts. Finally, rows 3-6 reveal some evidence that spread models may 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.[6]

4.3.3. Relationship with Forecast Efficiency Regressions

It is 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:

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

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. A rejection of these restrictions is interpreted as evidence of the existence of a time-varying risk premium (see, e.g., Fama and French 1987, 1988; Chernenko et al. 2004).[7] Chernenko et al. report that the hypothesis of forecast efficiency cannot be rejected at

[6] Motivated by term-structure models, we also experimented with models including a weighted average of spreads at different horizons. These models consistently performed so poorly that no results will be reported.
[7] Such tests implicitly postulate that the trader's loss function coincides with the econometrician's quadratic loss function. If that is not the case, forecast efficiency tests tend to be biased in favor of the alternative hypothesis (see Elliott, Komunjer, and Timmermann 2005).

conventional significance levels. It is important to bear in mind that such evidence does not necessarily 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 (4) that imposes this null does not dominate the no-change forecasts in out-of-sample forecasts. Second, as our forecasting 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.

In addition, it can be shown that the results in Chernenko et al. are not robust to updating the sample. Despite differences in the timing conventions used in constructing the monthly futures price data, we are able to replicate their results qualitatively using our data, but their sample period. For the full sample, however, we do reject the hypothesis of forecast efficiency at horizons 6 and 12 (see Table 6). This pattern is consistent with the earlier forecasting results. This rejection of forecast efficiency occurs despite the fact that $\hat{\alpha}$ is close to zero and $\hat{\beta}$ fairly close to 1, as suggested by theory, and very much unlike in the foreign exchange literature (see, e.g., Froot and Thaler 1990).

5. Are Futures Prices Unbiased Predictors of the Spot Price?

As equation ($1'$) shows, unbiasedness may be assessed only with reference to an explicit measure of the expected spot price. A reasonable proxy for the expected spot price is the best possible predictor available to real-time forecasters. The first part of section 5 is devoted to identifying this predictor. One possible approach to measuring market expectations of the spot price is the use of survey data. 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). An alternative approach is the use of econometric forecasting model of the spot price of crude oil. Some of these models are atheoretical, whereas others can be motivated based on economic theory.

5.1. Candidate Forecasting Models

5.1.1 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 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 standard 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; 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 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:

$$(8) \qquad \hat{S}_{t+h|t} = S_{t,h}^{CF} \qquad h = 3, 12$$

5.1.2. Expected Future Spot Prices from Econometric Models

An alternative to modeling expectations of spot prices for crude oil is based on econometric models. One example of such econometric models is the random walk model without drift introduced earlier. In this section, we introduce the random walk with drift and the Hotelling model as additional competitors.

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

$$(9) \qquad \hat{S}_{t+h|t} = S_t \left(1+\alpha\right) \qquad h = 1, 3, 6, 9, 12$$

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

(10) $$\hat{S}_{t+h|t} = S_t(1+\Delta\bar{s}_t^{(l)}) \qquad h = 1, 3, 6, 9, 12,\ l = 1, 3, 6, 9, 12$$

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. The local drift model is appealing in that it may capture "short-term forecastability" that arises from local trends in the oil price data.

An alternative forecasting model is motivated by Hotelling's (1931) model, which predicts that the price of an exhaustible resource such as oil appreciates at the risk free rate of interest:

(11) $$\hat{S}_{t+h|t} = S_t(1+i_{t,h}) \qquad h = 3, 6, 12$$

where $i_{t,h}$ refers to the interest rate at the relevant maturity h.[8] 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). We use the Treasury bill rate as a proxy for the risk free rate.[9]

5.2. Real-Time Forecast Accuracy of Candidate Forecasting Models

In this subsection, we compare the real time forecast accuracy of models (8)-(11) to that of the no-change forecast in (2).

5.2.1. Hotelling Model

Section 4.3 established that the no-change forecast tends to be more accurate than models based on the price of oil futures. An obvious question is whether the no-change forecast can be improved upon, for example, by using information on interest rates. Row 7 in Tables 2, 3, and 5 shows that the random walk model has lower MSPE than the Hotelling model at horizons of 3

---

[8] 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 risk free rate. In competitive equilibrium, oil companies extract crude oil at the socially optimal rate.

[9] Specifically, we use 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

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.

5.2.2. Random Walk Models with Drift

The next six rows in Tables 1-5 document that allowing for a drift in no case significantly lowers the MSPE of the random walk model, when the drift is estimated based on rolling regressions, and only in one case, when the drift is estimated recursively. Allowing for a drift lowers the MAPE at some horizons and for some models, but the gains are not systematic and different models work well for different horizons. Again, the Clark and West (2005) test rejects the null of no predictability in several cases (mainly at the nine-month horizon). As discussed earlier, that rejection does not necessarily translate into accuracy gains in real time forecasting, as evidenced by the MAPE rankings. In some cases, allowing for a drift also improves significantly the ability to predict the sign of the change of the oil price at longer horizons, but only when the drift is estimated recursively. 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 with drift (see Inoue and Kilian (2004) and the references therein). There is no evidence that such models dominate the no-change forecast.

5.2.3. Professional Survey Forecasts

The last row in Tables 2 and 5 shows that the consensus survey forecast has much higher MSPE than the no-change forecast at both the 3-month and 12-month horizons. It also has a larger bias and higher MAPE and there is no statistically significant evidence that it is better at predicting signs than a coin flip. The survey forecast is also inferior to the futures forecasts, suggesting that survey respondents do not rely on futures price data alone in forming their expectations.

5.3. Why the No-Change Forecast is a Plausible Measure of Expectations

The central result of section 5.2 is that no-change forecasts of the price of oil tend to be more accurate than forecasts based on econometric models and more accurate than survey forecasts.[10] This result is consistent with views among oil experts. For example, Peter Davies, chief economist of British Petroleum, has noted that "we cannot forecast oil prices with any degree of accuracy over any period whether short or long" (see Davies 2007). The favorable forecasting performance of the no-change forecast also is consistent with the observed high persistence of the nominal spot price of oil (see, e.g., Diebold and Kilian 2000). The high autocorrelation of commodity prices in general has been widely recognized in the literature (see, e.g., Deaton and Laroque 1992, 1996). Finally, it is important to stress that the superior forecast accuracy of the random walk model without drift does not contradict theoretical results in the literature that oil prices are endogenous with respect to global macroeconomic conditions (see, e.g., Barsky and Kilian 2002). The first point to keep in mind is that macroeconomic determinants such as U.S. interest rates, U.S. inflation, or global economic growth are but one of many determinants of the price of oil. For example, many of the major oil price increases in recent decades have been associated with unforeseen political disturbances in the Middle East and rising concerns about future oil supply shortfalls. Hence, one would not expect forecasting models based on macroeconomic fundamentals alone to be successful in practice. The second point to bear in mind is that predictability that exists in population may be difficult to exploit in real time in finite samples. The link from macroeconomic fundamentals to the price of oil is complicated and may even be nonlinear. Even if the spot price of crude oil does not truly follow a random walk, random walk forecasts tend to be attractive in terms of their mean-squared prediction error (MSPE) since the reduction in variance from excluding other predictors in small samples will typically more than offset the omitted variable bias. Thus, the superior forecast accuracy of the no-change forecast does not invalidate economic models of the crude oil market.

5.4. Testing the Unbiasedness of Futures Prices

We now turn to the question of whether futures prices are unbiased predictors of the spot price. For this purpose it is convenient to express the deviation of the futures price from the expected

[10] This result differs from at least some earlier studies. For example, Chernenko et al. (2004) report evidence that futures-based forecasts have marginally lower MSPE than the no-change forecast at horizons of 3, 6 and 12 months. In related work, Wu and McCallum (2005) find that futures prices are generally inferior to the no-change forecast, but report that spread regressions have lower MSPE than the no-change forecast at short horizons. These findings do not contradict our results. The differences in MSPE rankings can be traced mainly to differences in the sample period. The sample period considered in our paper is longer than in any previous study. Further sensitivity analysis suggests that evidence of accuracy gains, sometimes obtained in samples shorter than ours, tends to vanish when the full sample is examined.

spot price in percentage terms as:

(12) $$\frac{F_t^{(h)} - E_t\left[S_{t+h}\right]}{S_t}$$

Evidence that this expression differs from zero on average would cast doubt on the hypothesis of unbiasedness. Table 7 evaluates this expression under alternative assumptions about $E_t\left[S_{t+h}\right]$. For expository purposes, we focus on the 3-month and 12-month horizon. Our sample period is 1989.1-2007.2, as a contiguous time series for the 12-month basis becomes available only starting in 1989.1.

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. Substituting the no-change forecast S_t for $E_t\left[S_{t+h}\right]$, expression (12) can be expressed in terms of observables as:

(12′) $$\frac{F_t^{(h)} - S_t}{S_t} \equiv basis_t$$

The resulting expression coincides with the definition of the *basis* of oil futures (see Hull 2005). The spread term used in the spread regressions (4)-(7) is the log approximation to the basis. Table 7 shows, first of all, that the basis tends to be large in magnitude based on the mean absolute deviation. Formal tests of the hypothesis of a zero mean confirm that the futures price is a biased predictor of the spot price. Using heteroskedasticity and autocorrelation consistent (HAC) standard errors, on average both the 3-month and 12-month basis are statistically different from zero at the 1% level. Although the effective sample size is small, the rejection is decisive.

This result is robust to alternative assumptions about $E_t\left[S_{t+h}\right]$. The next column of Table 7 shows results under the alternative assumption of perfect foresight. Substituting the (in practice infeasible) perfect-foresight forecast S_{t+h} for $E_t\left[S_{t+h}\right]$, expression (12) can be expressed in terms of observables as:

(12″) $$\frac{F_t^{(h)} - S_{t+h}}{S_t}.$$

The quantitative results for the perfect foresight case are similar to the no-change forecast benchmark. There is evidence that futures prices systematically underpredict the spot price on average.[11]

Table 7 shows that the bias is negative, comparatively small in magnitude (under 5% at the 12 month horizon), and that it increases in the forecast horizon. Since risk aversion in the model of section 3 would imply a bias in the opposite direction, it cannot explain the results in Table 7. A more plausible explanation is the existence of a peso problem in the data (for example, traders assigning a small positive probability to a collapse of the Chinese economy). Moreover, not only would it have been difficult to exploit this bias for financial gain, given the large intermittent losses associated with such a trading strategy, but the sign of the bias was presumably unpredictable in real time. Thus, the data are not obviously inconsistent with the premise of risk neutrality. In section 6, we will show that the model of section 3 with risk neutrality imposed is indeed helpful in understanding the variability over time of the basis and its persistence.

Even if there is bias, futures prices may still provide accurate forecasts as long as the bias is approximately constant and one includes an intercept in the forecasting model. As Table 7 shows, however, the basis is time-varying and highly persistent, consistent with the inferior forecasting performance of spread models that allow for a nonzero intercept. We capture this persistence by modeling the basis 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 a measure of persistence as suggested by Andrews and Chen (1994). The sum of the autoregressive coefficients for the 3-month basis in the first column is 0.74, whereas that for 12-month basis is 0.81. This evidence demonstrates that there is considerable persistence in the basis.

These results are important in that they cast doubt on widely held views among policymakers. 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:

---

[11] We can also link the evidence in Table 7 to the forecast efficiency regressions underlying Table 6. It has been customary to focus on tests of $H_0 : \beta = 1$ and joint tests of $H_0 : \alpha = 0, \beta = 1$ in the literature. Evidence that these hypotheses cannot be rejected has been interpreted as evidence that that the oil futures spread contains information about changes in spot prices (see, e.g., Chernenko et al. 2004). Typically, no direct tests of the unbiasedness of futures prices have been reported in the literature. Unbiasedness implies that $\alpha = 0$ in forecast efficiency regressions. As Table 6 shows, direct tests of $H_0 : \alpha = 0$ tend to reject the null hypothesis of unbiasedness at all horizons. This result is consistent with the evidence in Table 7. Unlike in Chernenko et al., even the joint hypothesis is rejected at horizons 6 and 12. This difference is driven by the sample period.

> "…[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)[12]

That conclusion is questionable for two reasons. First, as we showed in section 4.2.2, 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 premise that futures prices are approximately unbiased predictors of spot prices, which seems inconsistent with the empirical evidence presented in this section. Third, whereas the bias in futures prices may seem small overall, the variability about the expected value is not. At any point in time, the discrepancy between the futures price and the expected spot price may be very large and go in either direction. Taking the spot price of crude oil to be $65, about its level in late March 2007, for example, the minimum and maximum value of the 12-month basis implies that the futures price may differ from the expected spot price by as much as $20 in one direction or by as much as $18 in the other (see Table 7). It is this variability of the deviation of expected spot prices and futures prices about the mean rather than the bias that drives the large MSPE of futures-based forecasts and that makes the use of such oil price forecasts inadvisable. The remainder of the paper is devoted to understanding the cause of these fluctuations.

6. Understanding Large Fluctuations in the Basis

The evidence in Table 7 illustrates that there is a fundamental difference between markets for financial assets such as the foreign exchange market and the market for crude oil. Whereas the basis of oil futures experiences large and persistent fluctuations over time that are poorly understood, the corresponding basis for foreign exchange futures is well explained by the interest rate differential which captures the opportunity cost of holding assets in one currency as opposed to another. This covered interest rate parity result has been documented, for example, by Taylor (1989). Considering the typical size of interest rate differentials, the basis in the major foreign exchange markets tends to be small compared with the basis in oil markets. This point is

[12] 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.

illustrated in Figure 3. In this section we propose an explanation of this discrepancy based on observation that the difference between the futures price and the expected spot price of oil is not accounted for by the interest rate alone, but that it also reflects the value that firms derive from having ready access to oil, a fact commonly referred to as the convenience yield. The presence of this convenience yield makes the analysis of oil futures markets fundamentally different from the analysis of the market for foreign exchange futures.

We demonstrate, based on the two-country, two-period general equilibrium model of section 3, that the persistent and large fluctuations in the basis documented in Table 7 and in Figure 3 can be understood in terms of fluctuations in the marginal convenience yield. The model implies that fluctuations in the marginal convenience yield can be directly linked to shifting fundamentals in the form of expectation shifts 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 both large and sharp fluctuations in the basis over time.

We model shifts in uncertainty as a mean-preserving spread. In other words, we abstract from changes in the conditional mean of oil supplies and focus on changes in the conditional variance. The motivation for this modeling choice is best seen by focusing on the example of the Persian Gulf War. Events such as the invasion of Kuwait in August of 1990 have two distinct effects. First, they cause a reduction in expected oil supply. This oil supply shock represents a change in the conditional mean of oil supplies. It has been documented in the literature that such a shock indeed occurred in 1990, but that this supply shock fails to explain the bulk of the movements in the real price of oil in 1990/91. Second, there is an increase in uncertainty about future oil supply shortfalls. Indirect evidence that the price spike of 1990/91 was driven by increased uncertainty about future oil supply shortfalls has been presented in Kilian (2007c). We model this increased uncertainty as an increase in the conditional variance of oil supplies, implicitly abstracting from the global business cycle.

An increase in the variance allows for the possibility of realizations higher or lower than the expected value. Such an increased spread of outcomes makes sense, following an event such as the invasion of Kuwait, since we need to allow for the possibility that Iraq may choose to withdraw from Kuwait to avoid a military conflict or that other countries will choose to increase their oil supplies to offset the reduction of Kuwaiti output. We also need to allow for the possibility that Saddam Hussein may choose to invade Saudi Arabia, leading to a further cutback of oil supplies. Thus, the uncertainty increases in both directions, following a reduction in the conditional mean.

6.1. Comparative Statics

In this subsection, we state two comparative statics results derived under risk neutrality. The first result is that an increase in uncertainty about the availability of oil supplies in the second period raises the first period's real spot price; the second result is that under plausible assumptions 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. The thought experiment is an increase in ε.

6.1.1. 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

$$F'(\omega - I_{US}) - g_1(I_{US}, \sigma_\varepsilon^2) = \beta[\theta F'(\omega + \varepsilon + I_{US}) + (1-\theta)F'(\omega - \hat{\varepsilon} + I_{US})].$$

Totally differentiating this expression with respect to ε yields

$$\left(A - g_{11}\right) dI_{US} - g_{12} d\sigma_\varepsilon^2 = B d\varepsilon$$

where

$$A = -[F''(Z_1) + \beta\theta F''(Z_{21}) + \beta(1-\theta)F''(Z_{22})] > 0$$

$$B = \beta\theta[F''(Z_{21}) - F''(Z_{22})] > 0$$

$B > 0$ by $F''' > 0.$ Solving for the change in U.S. inventories yields

$$\text{(13)} \qquad \frac{dI_{US}}{d\varepsilon} = \frac{B}{A - g_{11}} + \frac{g_{12}}{A - g_{11}} \frac{d\sigma_\varepsilon^2}{d\varepsilon} > 0,$$

which shows that the United States accumulates inventories in response to a mean-preserving increase in the spread. This result follows from the fact that $A > 0$ and $B > 0,$ combined with the observation that each unit of inventory becomes more valuable in response to an increase in σ^2, implying that $g_{12} > 0,$ and the fact that $g_{11} < 0,$ which follows from the strict concavity of $g(\cdot)$ in inventories (see, e.g., Pindyck 2001). Since

$$\frac{S_1}{P_1} = F'(\omega - I_{US})$$

is an increasing function of U.S. 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 about future oil supply shortfalls raises the spot price of oil in terms of the consumption good.

Expression (13) illustrates that the spot price increases for two related reasons. First, the U.S. accumulates more inventories in response to a mean-preserving spread. This direct effect is represented by the first term in expression (13). Second, the willingness to pay for a given unit of inventories is higher, as uncertainty increases. This indirect effect operates through σ_ε^2 and is represented by the second term. The combined effect is an increase in oil inventories, which by the concavity of the production function is associated with an increase in the marginal product and hence a higher real spot price. We refer to the shift in the demand for oil inventories in expression (13) as increased *precautionary demand* for crude oil. The increment in the real spot price of oil caused by this shift is referred to as the *precautionary demand component of the spot price*.

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

Under risk neutrality, 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:

$$\frac{F_1 - S_1}{S_1} = R - (1+R)\frac{g_1(I_{US}, \sigma_\varepsilon^2)}{S_1/P_1}.$$

Totally differentiating the right-hand side of this equation with respect to ε, we obtain

$$\frac{dR}{d\varepsilon} - \frac{d(1+R)}{d\varepsilon}\left[\frac{g_1}{\frac{S_1}{P_1}}\right] - (1+R)\left\{\frac{1}{F'}\left[g_{11}\frac{dI_{US}}{d\varepsilon} + g_{12}\frac{d\sigma_\varepsilon^2}{d\varepsilon}\right] + g_1\frac{F''}{F'^2}\frac{dI_{US}}{d\varepsilon}\right\},$$

Assuming that $dR/d\varepsilon \approx 0$, which seems plausible for our sample period, the sign of this expression depends on the relative magnitudes of (1) the decrease in the marginal convenience yield associated with the inventory accumulation associated with the mean-preserving spread; and (2) the increase in the marginal convenience yield associated with the increase in σ_ε^2 triggered by the same shock. Hence, in general a mean-preserving spread will move the basis. The direction of this effect depends on the curvature of F and g. The basis will decline if and only if

(14) $$\frac{d\sigma_\varepsilon^2}{d\varepsilon} > -\frac{1}{g_{12}}\left[g_{11} + g_1 \frac{F''}{F'}\right]\frac{dI_{US}}{d\varepsilon}.$$

We can express both $d\sigma_\varepsilon^2 / d\varepsilon$ and $dI_{US} / d\varepsilon$ in terms of the model's parameters and show that expression (14) is equivalent to:

(14′) $$g_{12} > \frac{\lambda(1-\theta)B}{2\theta\varepsilon(1-\lambda)},$$

where $\lambda \equiv -\left(g_{11}/g_1 + F''/F'\right) g_1/(A - g_{11})$, $0 < \lambda < 1$. Hence, for a given stock of inventories and given ε, the basis will decline, provided g_{12} is large enough. The term g_{12} measures the shift in the marginal convenience yield induced by the mean-preserving spread. This shift reflects the fact that following an increase in uncertainty each unit of inventory has greater value as an insurance against supply shortfalls. In other words, the basis will decline if agents' willingness to pay for an extra barrel of oil to be used as an insurance against oil supply shortfalls increases sufficiently in response to an unanticipated shift in uncertainty. It is well-documented that during past oil price shocks, traders were willing to pay exorbitant prices to procure extra stocks of oil (see, e.g., Penrose 1976; Terzian 1985). Thus, large values of g_{12} seem empirically plausible and one would expect a negative correlation between the basis and the component of the spot price driven by precautionary demand for oil, at least following major shifts in uncertainty. This implication is testable. In section 6.2 we will show that indeed exogenous events in the oil market are associated with sharp declines in the basis, as would be expected if g_{12} were large. These events provide an economic explanation for the large and persistent fluctuations in the basis that undermine the forecasting accuracy of oil futures prices.[13]

It is useful to contrast our conclusion to the earlier literature. In traditional models that incorporate the marginal convenience yield such as Working (1948), Brennan (1958), Telser (1958), and Fama and French (1988) an increase in inventory holdings has an unambiguously

[13] Earlier we documented that the basis is highly persistent, but mean reverting (see Table 7). We also documented that the no-change forecast (or random walk model) is the best predictor of the nominal spot price of oil. The conclusion that under plausible conditions the mean-reverting basis is correlated with the precautionary demand component of the spot price may seem to imply a contradiction of the random walk result. This is not the case. First, the result about the forecast accuracy refers to the nominal price of oil, whereas the comparative statics result is for the real price of oil. Second, the forecasting results are for total spot price of oil, whereas the results of this section are only for one of the components of the price of oil. Third, as Diebold and Kilian (2000) demonstrate, for autoregressive processes with degrees of persistence in the range documented in Table 7 an incorrectly specified random walk model will tend to have lower MSPE than the correct mean-reverting model in small samples.

positive effect on the basis because, as inventory holdings increase, the marginal convenience yield declines. Our conclusion is different. The reason is that we follow Pindyck (2001) in incorporating an *indirect* effect from increased uncertainty on the marginal convenience yield, which raises the value of holding inventories when uncertainty increases and pushes up the marginal convenience yield for each unit of inventory. In other words, the direct effect through inventories reflects the accumulation of additional inventories in response to the increase in ε; whereas the indirect effect operating through σ^2 captures the notion that the willingness to pay for a unit of inventories increases with uncertainty.

6.2. Using the Model to Identify Expectations Shifts

The comparative statics results above suggest that fluctuations in the basis will be indicative of fluctuations in the spot price of oil driven by precautionary demand for crude oil, provided g_{12} is large enough.[14] One way of judging the empirical content of the model is to verify that the basis moves in the expected direction at times of major unforeseen events such as the outbreak of the wars. In Figure 4, we focus on several clearly defined events in recent history that should have been associated with shifts in the market's uncertainty about future oil supply shortfalls such the Persian Gulf War and the 2003 Iraq War (which should have caused the basis to fall), and the Asian Financial Crisis and 9/11 (which should have caused the basis to increase as world demand for crude oil fell, making a shortfall less likely). Clearly, expectations shifts of the type embodied in our theoretical model are not the only possible reason for shifts in the basis, but arguably they are the most important reason.

Figure 4 plots the negative of the basis for 1989.1-2007.2 by horizon. This normalization allows us to interpret positive spikes as increases in the precautionary demand component of the real spot price. The plot confirms the conclusion in Kilian (2007a) that the sharp spike in oil prices during the Persian Gulf War was driven by expectations shifts reflected first in higher precautionary demand, as Iraq invaded Kuwait, and then in lower precautionary demand, as the U.S. troop presence in the region increased (also see Kilian 2007a). Likewise, the spike after mid-2002 in the period leading up to the 2003 Iraq War is as expected, given that the Iraq War

[14] 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.

was anticipated by the market starting in the summer of 2002 (see Barsky and Kilian 2004). The plot also indicates that the temporary decline in oil prices following the Asian crisis (and its reversal after 1999) reflected fluctuations in precautionary demand. There is a similar but smaller temporary decline following the adverse demand shock associated with 9/11. Anecdotal evidence suggests that the spike in 1996 was associated with concerns about tight oil supplies and that in 2000 with concerns arising from strong demand for crude oil. In addition, the plot suggests a persistent decline in precautionary demand in recent years. Such a decline seems highly implausible on a priori grounds, given that recent years have been characterized by widespread concerns about future oil supply shortfalls, a point to which we will return below.

6.3. Alternative Measures of Precautionary Demand Shifts

The index of expectations-driven oil price increases proposed in this paper is not the only possible measure. 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) based on different data and a different methodology. Unlike the measure developed in this paper, that estimate was based on a structural VAR decomposition of the real price of crude oil. The structural representation of the underlying trivariate autoregressive model is

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

Where p denotes the lag order, ε_t is 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) index of global real economic activity that captures fluctuations in the global demand for all industrial commodities, and the real price of oil (in that order), measured 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. In this model, innovations to the real price of oil that cannot be explained by oil supply shocks or demand shocks that are common to all industrial commodities by construction must be demand shocks that are specific to the oil market. The latter oil-specific demand shock by construction captures fluctuations in precautionary demand for oil driven by fears about the availability of future oil supplies. Kilian (2007a) makes the case that this shock effectively can be interpreted as a precautionary demand shock, given the absence of plausible alternative interpretations and given the time path of this shock during specific historical episodes, during which we would expect precautionary demand to shift.

The structural VAR model postulates a vertical short-run supply curve for crude oil and a downward sloping short-run demand curve that is being shifted by innovations to the business cycle in global industrial commodity markets as well as expectations shifts in the global crude oil market. Given these assumptions, one can use the structural moving average decomposition of the VAR model to 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 future oil supply shortfalls. While it is not possible to compare this VAR-based measure of the precautionary demand component of the spot price to the futures-based measure for the full sample period of 1973-2006 considered in Kilian (2007a), given the limited availability of oil futures price data, we may compare these two measures for the period 1989.1-2006.12, which includes several major oil price spikes. Since the futures-based measure is an index and the VAR-based measure is not, the appropriate metric of comparison is their contemporaneous correlation.

Table 8 shows that the two measures in general are highly correlated notwithstanding the differences in their method of construction. For the sample period of 1989.1 through 2006.12, the correlation ranges from 39% at the 3-month horizon to 61% 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. Thus, we focus on the 12-month basis. A correlation of 61% between two independently constructed measures of the fluctuations in the spot price of oil driven by precautionary demand is remarkably high. The correlation is even higher if we exclude the last three years of data, for which the basis seems implausibly high, as discussed above. In fact, for the post-2003 period, the basis data and the VAR-based measure of the precautionary demand component of the spot price of oil paint a somewhat different picture (see Figure 5). Whereas the VAR-based measure on average remains at a high level after 2003.12, consistent with the perception of sustained uncertainty about future oil supply shortfalls, the futures-based measure systematically declines. This evidence casts further doubt on the credibility of the negative of the basis as an index of fluctuations in the precautionary demand component of the spot price over the last three years of the sample.

Table 8 shows that, excluding the last three years, the correlation of the two measures rises to 79% at the 12 month horizon. A correlation of near 80% for most of the sample is evidence both of the predictive power of our theoretical model of the oil futures and spot markets and of the realism of the identifying assumptions underlying the VAR-based measure. The next subsection will discuss a possible reason for the weakening of this close relationship starting in 2004.

6.4. Toward an Explanation of the Weakening of the Empirical Relationship

The two facts that (1) the basis generates economically implausible predictions after 2003.12 and that (2) its path differs systematically from independent estimates of precautionary demand after 2003.12, both suggest that a structural change may have occurred around 2003.12 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 (also see United States Senate 2006). 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 oil futures contracts as defined by the CFTC’s regulations (CFTC 2007). A trader is considered commercial by the CFTC as long as the trader is engaged in business activities hedged by the use of oil futures. Such traders are identified as hedgers (as opposed to speculators).[15] 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 oil futures market. 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. 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.

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).[16] 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.[17] 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 marked and sustained increase in the percent share of noncommercial spread positions since 2003.12, suggesting that

[15] 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).

[16] 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.

[17] 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.

speculation has recently intensified. Although there have been other spikes in speculative activity in the past, the most recent increase in the non-commercial spread position is unprecedented in the sample. The percentage in question has more than doubled since early 2004. Given that the fluctuations in basis become implausible after 2003.12, it is natural to test formally whether there has been a structural break in the share of speculators shown in Figure 6 at that point in time. Since the potential breakpoint has not been determined based on inspecting the data in Figure 6 (but on the basis of the data in Figures 4 and 5), a statistical test of this proposition does not involve an endogenous breakpoint selection problem and inference is standard. We fit a deterministic trend polynomial of order 1 to the share of speculators, and test for the presence of a one-time break in the level and slope of this regression in 2004.1. The *p*-value of a Wald test of the no-break null hypothesis is 0.000, confirming that there has been a shift after 2003.12.

Are the two facts of an implausibly high basis after 2003.12 and the rapidly increasing share of speculators related? The correlation between the 12-month basis and the share of speculators in Figure 6 is 0.871 in the last three years of the sample. This evidence does not prove that increased speculation caused the unexpected decline in the negative of the basis and created the mistaken impression that precautionary demand has fallen, but it is suggestive. First, there are no other important exogenous events in the oil futures market that occurred near that date. Second, a weakening of the empirical relationship between the basis and the precautionary demand component of the spot price of oil seems likely in response to sharply increased speculative trading, because our theoretical model does not allow for speculators.

For an exogenous increase in speculation to cause a spurious increase in the basis (and hence a decline in the index), it would have to be the case that increased speculation causes the futures price to increase faster than the spot price. Establishing such a link will require a theoretical model that allows for speculation. Speculators enter the futures market when they expect the futures price to increase, regardless of the economic fundamentals of the market. They will buy futures contracts, as long as they believe that in the future someone else will be willing to pay even more for that contract. Thus, their expectations formation is inconsistent with the assumptions of our theoretical model. Incorporating speculative traders with model-inconsistent expectations into this theoretical framework is a nontrivial task that is left for future research. Indeed, there is reason to believe that a different class of models may be required to analyze such behavior.

We conclude that oil futures prices (or related variables such as the futures 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 in the form of the basis. The basis of oil futures will not be a reliable indicator of fluctuations in the precautionary demand component of the spot price of oil, however, unless major structural shifts in the composition of traders are accounted for.

7. The Relationship between Oil Inventories and the Precautionary Demand Component of the Spot Price

Whereas we have focused on the empirical relationship between increased concerns about future oil supply shortfalls and the precautionary demand component of the spot price of oil, the model also has implications for the behavior of inventories in response to increased uncertainty. Testing these implications is not straightforward, given that inventories move for many reasons other than expectations shifts. First, for the spot price, the VAR decomposition allows us to focus specifically on the precautionary demand component of the spot price. No similar measure of the precautionary demand component of inventories exists. Second, inventory data are trending, and measures of the comovement between the precautionary demand component of the spot price and inventories tend to be sensitive to the method of detrending. Nevertheless, inspection of the inventory data around the time of the Persian Gulf War suggests that there was no inventory accumulation in the OECD immediately following the invasion of Kuwait in August of 1990. Rather inventory accumulation was delayed until 1991. At first glance, this result seems to contradict the two-period model of section 3. The purpose of the current section is to demonstrate that there is no contradiction.

7.1. Frictions in Inventory Markets

Our two-period model implies that inventories should increase along with the spot price of oil, as the mean-preserving spread increases precautionary demand for oil. This result can be thought of as the equilibrium outcome of a dynamic adjustment process that may extend over several periods. Thus, it does not convey information about the time path of the adjustment.

In practice, inventories are not likely to adjust instantaneously. Empirically, firms in possession of crude oil appear to be particularly reluctant to part with their crude oil stocks, when uncertainty about future oil supply shortfalls increases. This situation has been aptly described by Terzian (1985, p. 260) in the context of the 1979 oil price shock:

> "Spot deals became more and more infrequent. The independent refineries, with no access to direct supply from producers, began to look desperately for oil on the so-called 'free market'. But from the beginning of November, most of the big oil companies invoked *force majeure* and reduced their oil deliveries to third parties by 10% to 30%, when they did not cut them off altogether. Everybody was anxious to hang on to as much of their own oil as possible, until the situation had become clearer. The shortage was purely psychological, or 'precautionary' as one dealer put it."

Penrose (1976, p. 46) describes a similar hoarding phenomenon in the period leading up to the 1973 oil price shock, as oil companies became concerned with the possibility of being expropriated. In her words, "the major oil companies became increasingly cautious about outside sales as uncertainty increased."

In this view, the market's attempt to increase inventories in the aggregate will be frustrated, at least initially, and the accumulation of inventories will occur only gradually over time. To capture this delay one requires a multi-period model. In a multi-period setting, the spot price will overshoot, before declining to the new equilibrium level, as inventories accumulate over time.[18] In other words, in the multi-period model, there will be a pure valuation effect that drives up the spot price of oil instantaneously, without any movement in crude oil inventories. Thus, there is no reason for increases in the precautionary demand component of the spot price to be associated with inventory accumulation.

7.2. The Three-Period Model

This basic point can be made in a three-period model. Except for the timing of the decisions, the set up of this model is identical to the two-period model. We postulate that the United States and Saudi Arabia make their production, consumption, and inventory-holding decisions at the beginning of each period and do not have the opportunity to readjust these choices until the beginning of the subsequent period. At the beginning of the first period, both countries make their production and consumption decisions. At the end of the first period, the uncertainty about an oil supply shortfall in period 3 increases. Consistent with the anecdotal evidence discussed above, countries are permitted to adjust their production and consumption decision only at the beginning of the second period after prices have adjusted.

Thus, in contrast to the two-period model, inventories cannot adjust immediately in the

[18] This overshooting result is analogous to the overshooting of the exchange rate in the Dornbusch (1976) model. It is also clearly evident in the estimate of the response of the real price of oil to a precautionary demand shock in the VAR model of Kilian (2007a).

United States, but spot and futures prices can adjust. This adjustment occurs through a change in the marginal convenience yield. Figure 7a illustrates the immediate effect of an increase in uncertainty on the marginal convenience yield for a fixed level of U.S. inventory holdings in period 1. Each unit of inventory becomes more valuable as a means of avoiding disruptions to production of the consumption good in period 3. This increase in the value of the marginal unit of inventory is reflected in the increase in convenience yield. We refer to this price increase as the pure valuation effect, as the stock of inventories remains constant.

The effect of the increase in the marginal convenience yield on the spot price and futures price is evident from the equation for the basis in period 1, which can be derived by analogy to the corresponding equation in section 6.1.2:

$$\frac{S_1(1+R_1)-F_1}{P_1}=(1+R_1)g_{1,1}(I_{US,1},\sigma^2),$$

where $g_{1,1}(I_{US,1},\sigma^2)$ is the analogue of g_1 for period 1, and the other variables are defined as before. $I_{US,1}$ remains fixed after the increase in uncertainty because the stock of inventory holdings is predetermined in the three-period model. The difference between the capitalized value of the spot price and the futures price is increasing the marginal convenience yield. When the marginal convenience yield schedule shifts outward in response to the increase in uncertainty, this difference increases. Consequently, both the spot price and the futures price increase when there is an increase in uncertainty, but the spot price at the end of period 1 increases relative to the futures price. Indeed, this is what we observe in the second half of 1990 (see Figure 4).

Figure 7b illustrates the dynamic path of the economy starting at the beginning of period 2, when the United States accumulates inventory holdings in response to the increase in uncertainty about an oil supply shortfall in period 3. The argument that establishes this result is identical to the one in the two-period model and relies upon the concavity of the U.S. production function. As a consequence of the accumulation of inventories, both the marginal convenience yield in period 2 and the spot price in period 2 decline relative to their values at the end of period 1, although the spot price will remain higher than before the initial increase in uncertainty, consistent with the predictions of the two-period model. Hence, we expect negative comovement between inventories and the spot price, in the second period following the mean-preserving spread. Again, this model prediction is consistent with the inventory data for 1991.

In sum, the effects of an increase in uncertainty in the three-period model are staggered as follows: (1) There is an instantaneous valuation effect that shifts up the marginal convenience yield schedule for all levels of inventories and raises the spot price of oil; (2) subsequently, there will be a gradual adjustment of inventories that reduces the convenience yield and the spot price relative to their period 1 values. This result provides a formal justification for the delayed adjustment of OECD oil inventories in response to large shifts in expectations about future oil supply shortfalls. It suggests that inventories need not respond immediately to expectation shifts about future oil supplies when one allows for frictions that prevent the immediate reallocation of oil from current consumption to inventories.

8. Conclusion

We introduced a two-country, multi-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 the basis. The key insights can be summarized as follows:

Our theoretical model illustrates that there is no support for the common perception that prices of oil futures are good predictors of forecast spot prices in the MSPE sense. Using observations up to February of 2007, we confirmed that the price of crude oil futures is not an accurate predictor of the spot price of crude oil. Many users of futures-based forecasts are aware of this caveat 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. No-change forecasts are also more accurate than commercial survey-based forecasts. Our evidence that no-change forecasts tend to be most accurate in forecasting the spot price of crude oil and hence are reasonable proxies for the expected spot price, allowed us to test the proposition that oil futures prices are unbiased predictors of the spot price. We rejected the unbiasedness hypothesis, although the bias is small overall.

We showed that the large MSPE of futures-based forecasts is driven not by the bias, but by the variability of the futures price about the expected spot price. We documented large and time-varying deviations of futures prices from expected spot prices, as measured by the basis of oil futures. For example, given a spot price of $65, the 12-month futures price may deviate as much as $20 from the expected spot price or as little as $0. Our analysis demonstrates that fluctuations in the oil futures basis are larger and more persistent than fluctuations in the basis of foreign exchange futures, which helps explain the poor forecasting performance of futures prices. We showed that this anomaly is linked to the presence of a marginal convenience yield in the oil futures market that is absent in the foreign exchange futures market. Our theoretical model incorporates this marginal convenience yield. The model implies that the oil futures basis is directly linked to shifts in oil market fundamentals. Specifically, we established that shifts in the uncertainty about future oil supply shortfalls cause fluctuations in the oil futures basis not found in models of the foreign exchange futures market. Under plausible conditions, there is a negative correlation between the oil futures basis and the component of the spot price of oil driven by precautionary demand for crude oil. We also showed that increases in the precautionary demand for crude oil need not go hand in hand with an accumulation of oil inventories.

Our main empirical result is that the negative of the basis may be viewed as an index of fluctuations in the price of crude oil driven by precautionary demand for oil. The time path of the basis since 1989 suggests major shifts in precautionary demand for oil during the Persian Gulf War and following the Asian crisis, for example. These results provide independent evidence of how shifts in market expectations about future oil supply shortfalls affect the spot price of crude oil. Such expectation shifts have been difficult to quantify, yet have been shown to play an important role in explaining oil price fluctuations. Our empirical results are consistent with related evidence in the literature obtained by alternative methodologies.

While the oil futures basis can be useful in identifying expectations shifts in general, our results suggest caution in interpreting the basis data in the absence of further information about the market structure. Specifically, the ability of the basis to capture fluctuations in the spot price of oil driven by shifts in the precautionary demand for oil weakens after 2003.12. That weakening coincides with an unprecedented increase in speculative activities in the oil futures market. To the extent that increased speculative trading tends to raise futures prices more than spot prices (and hence raises the basis), this weakening is not unexpected, as our theoretical model does not allow for speculation. Establishing such a link is left for future research.

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., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.

Barsky, R.B., and L. Kilian (2004), "Oil and the Macroeconomy Since the 1970s," *Journal of Economic Perspectives*, 18(4), 115-134.

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

Bernanke, B.S. (2006). "Energy and the Economy," remarks before the Economic Club of Chicago, Chicago, IL, http://www.federalreserve.gov/boarddocs/speeches/2006/200606152/default.htm

Brennan, M.J. (1958), "The Supply of Storage," *American Economic Review*, 48(1), 50-72.

Brennan, M.J. (1991), "The Cost of Convenience and the Pricing of Commodity Contingent Claims," in D. Lund and B. Øksendal (eds.), *Stochastic Models and Option Values,* North Holland: Amsterdam.

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.

Deaton, A., and G. Laroque (1992), “On the Behavior of Commodity Prices,” *Review of Economic Studies*, 59, 1-23.

Deaton, A., and G. Laroque (1996), “Competitive Storage and Commodity Price Dynamics,” *Journal of Political Economy*, 104, 896-923.

Diebold, F.X., and L. Kilian (2000), “Unit Root Tests are Useful for Selecting Forecasting Models,” *Journal of Business and Economic Statistics*, 18(3), 265-273.

Diebold, F.X. and R.S. Mariano (1995), “Comparing Predictive Accuracy,” *Journal of Business and Economic Statistics*, 13(3), 253-263.

Dornbusch, R. (1976), “Expectations and Exchange Rate Dynamics,” *Journal of Political Economy*, 84, 1161-1176.

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.

Froot, K.A., and R.H. Thaler (1990), “Anomalies: Foreign Exchange,” *Journal of Economic Perspectives*, 1(3), 179-192.

Gibson, R., and E.S. Schwartz (1990), “Stochastic Convenience Yield and the Pricing of Oil Contingent Claims,” *Journal of Finance*, 45(3), 959-976.

Gramlich, E.M. (2004), Oil Shocks and Monetary Policy, Annual Economic Luncheon, Federal Reserve Bank of Kansas City, Kansas City, Missouri. http://www.federalreserve.gov/boarddocs/speeches/2004/20040916/default.htm

Greenspan, A. (2004a), “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

Greenspan, A. (2004b), “Oil,” Speech presented at the National Italian American Foundation, Washington, DC. htttp://www.federalreserve.gov/boarddocs/speeches/2004/200410152/default.htm

Greenspan, A. (2005), “Energy,” remarks before the Japan Business Federation, Tokyo, Japan, http://www.federalreserve.gov/boarddocs/speeches/2005/20051017/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*, Prentice Hall: Upper Saddle River, NJ.

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 (2005), *World Economic Outlook*, Washington, DC.

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,” mimeo, Department of Economics, University of Michigan.

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*.

Kohn, D.L. (2007), “The Economic Outlook,” remarks at the Atlanta Rotary Club, Atlanta, GA, http://www.federalreserve.gov/boarddocs/speeches/2007/20070108/default.htm

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.

Penrose, E. (1976), “The Development of Crisis,” in R. Vernon (ed.), *The Oil Crisis*, New York: Norton.

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.

Svensson, L.E.O. (2005), "Oil Prices and ECB Monetary Policy," mimeo, Department of Economics, Princeton University, see http://www.princeton.edu/svensson/

Taylor, M.P. (1989), "Covered Interest Arbitrage and Market Turbulence," *Economic Journal*, 99, 376-391.

Telser, L.G. (1958), "Futures Trading and the Storage of Cotton and Wheat," *Journal of Political Economy*, 66, 233-255.

Terzian, P. (1985), *OPEC. The Inside Story*. London: Zed Books.

Townsend, R.M. (1978), "On the Optimality of Forward Markets," *American Economic Review*, 68(1), 54-66.

United States Senate (2006), *The Role of Market Speculation in Rising Oil and Gas Prices: A Need to Put the Cop back on the Beat.* Staff Report, Permanent Subcommittee of Investigations, United States Senate, Washington, DC, June 27.

Working, H. (1948), "Theory of the Inverse Carrying Charge in Futures Markets," *Journal of Farm Economics*, 30(1), 1-28.

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 (*p*-value)	Bias	MAPE (*p*-value)	Success Ratio (*p*-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)
$S_t\left(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(1)}/S_t)\right)$	6.994 (0.175)	0.200	1.954 (0.580)	0.479 (0.529)
$S_t\left(1+\hat{\beta}\ln(F_t^{(1)}/S_t)\right)$	6.975 (0.104)	0.156	1.948 (0.462)	0.423 (0.984)
$S_t\left(1+\hat{\alpha}+\ln(F_t^{(1)}/S_t)\right)$	7.138 (0.799)	0.162	1.948 (0.439)	0.526 (0.257)
$S_t\left(1+\ln(F_t^{(1)}/S_t)\right)$	7.106 (0.807)	0.212	1.949 (0.676)	0.443 (0.898)
$S_t(1+\hat{\alpha})$	7.013 (0.384)	0.186	1.945 (0.522)	0.479 (0.497)
$S_t(1+\Delta\overline{s}_t^{(1)})$	13.946 (0.457)	-0.061	2.604 (0.003)	0.490 (0.646)
$S_t(1+\Delta\overline{s}_t^{(3)})$	10.044 (0.717)	0.015	2.235 (0.151)	0.521 (0.294)
$S_t(1+\Delta\overline{s}_t^{(6)})$	8.293 (0.835)	0.005	2.050 (0.087)	0.495 (0.567)
$S_t(1+\Delta\overline{s}_t^{(9)})$	8.155 (0.932)	-0.016	2.057 (0.806)	0.495 (0.567)
$S_t(1+\Delta\overline{s}_t^{(12)})$	7.405 (0.305)	-0.023	1.943 (0.521)	0.505 (0.443)

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 $\Delta\overline{s}_t^{(l)}$ 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 (*p*-value)	Bias	MAPE (*p*-value)	Success Ratio (*p*-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)
$S_t\left(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(3)}/S_t)\right)$	24.217 (0.870)	0.253	3.610 (0.990)	0.407 (0.996)
$S_t\left(1+\hat{\beta}\ln(F_t^{(3)}/S_t)\right)$	22.826 (0.983)	0.804	3.541 (0.998)	0.407 (0.992)
$S_t\left(1+\hat{\alpha}+\ln(F_t^{(3)}/S_t)\right)$	22.090 (0.747)	0.315	3.365 (0.965)	0.397 (0.998)
$S_t\left(1+\ln(F_t^{(3)}/S_t)\right)$	19.036 (0.348)	0.649	3.176 (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.
$S_t(1+\hat{\alpha})$	19.699 (0.351)	0.484	3.114 (0.345)	0.485 (0.413)
$S_t(1+\Delta\overline{s}_t^{(1)})$	27.857 (0.710)	0.210	3.620 (0.119)	0.510 (0.418)
$S_t(1+\Delta\overline{s}_t^{(3)})$	24.702 (0.961)	0.238	3.461 (0.707)	0.500 (0.524)
$S_t(1+\Delta\overline{s}_t^{(6)})$	22.098 (0.893)	0.213	3.231 (0.315)	0.485 (0.685)
$S_t(1+\Delta\overline{s}_t^{(9)})$	20.242 (0.531)	0.224	3.105 (0.023)	0.557 (0.061)
$S_t(1+\Delta\overline{s}_t^{(12)})$	20.013 (0.454)	0.223	3.071 (0.005)	0.546 (0.101)
$S_{t,3}^{CF}$	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 (*p*-value)	Bias	MAPE (*p*-value)	Success Ratio (*p*-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)
$S_t\left(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(6)}/S_t)\right)$	51.809 (0.738)	1.012	5.315 (0.794)	0.485 (0.613)
$S_t\left(1+\hat{\beta}\ln(F_t^{(6)}/S_t)\right)$	47.143 (0.917)	1.959	5.200 (0.904)	0.464 (0.703)
$S_t\left(1+\hat{\alpha}+\ln(F_t^{(6)}/S_t)\right)$	40.640 (0.710)	1.074	4.692 (0.528)	0.485 (0.576)
$S_t\left(1+\ln(F_t^{(6)}/S_t)\right)$	36.475 (0.721)	1.684	4.621 (0.910)	0.485 (0.483)
$S_t(1+i_{t,6})$	34.906 (0.713)	0.382	4.509 (0.708)	0.557 N.A.
$S_t(1+\hat{\alpha})$	33.942 (0.132)	1.093	4.678 (0.155)	0.515 (0.021)
$S_t(1+\Delta\overline{s}_t^{(1)})$	44.981 (0.780)	0.543	4.898 (0.275)	0.505 (0.501)
$S_t(1+\Delta\overline{s}_t^{(3)})$	41.100 (0.874)	0.605	4.738 (0.571)	0.479 (0.762)
$S_t(1+\Delta\overline{s}_t^{(6)})$	35.936 (0.691)	0.671	4.531 (0.170)	0.510 (0.424)
$S_t(1+\Delta\overline{s}_t^{(9)})$	33.812 (0.293)	0.585	4.372 (0.988)	0.557 (0.091)
$S_t(1+\Delta\overline{s}_t^{(12)})$	34.379 (0.437)	0.708	4.465 (0.085)	0.510 (0.411)

Notes: See Table 1.

Table 4: 9-Month Ahead Recursive Forecast Error Diagnostics

$\hat{S}_{t+9\|t}$	MSPE (*p*-value)	Bias	MAPE (*p*-value)	Success Ratio (*p*-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)
$S_t\left(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(9)}/S_t)\right)$	54.225 (0.471)	2.515	5.406 (0.296)	0.546 (0.035)
$S_t\left(1+\hat{\beta}\ln(F_t^{(9)}/S_t)\right)$	54.939 (0.632)	3.163	5.411 (0.452)	0.536 (0.026)
$S_t\left(1+\hat{\alpha}+\ln(F_t^{(9)}/S_t)\right)$	55.042 (0.725)	2.502	5.313 (0.361)	0.546 (0.025)
$S_t\left(1+\ln(F_t^{(9)}/S_t)\right)$	54.642 (0.898)	3.017	5.403 (0.948)	0.526 (0.080)
$S_t(1+\hat{\alpha})$	46.107 (0.111)	2.090	5.150 (0.130)	0.557 (0.000)
$S_t(1+\Delta\overline{s}_t^{(1)})$	59.202 (0.876)	1.408	5.623 (0.342)	0.495 (0.611)
$S_t(1+\Delta\overline{s}_t^{(3)})$	51.025 (0.658)	1.492	5.258 (0.245)	0.510 (0.431)
$S_t(1+\Delta\overline{s}_t^{(6)})$	46.300 (0.303)	1.556	5.116 (0.092)	0.595 (0.581)
$S_t(1+\Delta\overline{s}_t^{(9)})$	45.428 (0.168)	1.581	5.082 (0.048)	0.510 (0.401)
$S_t(1+\Delta\overline{s}_t^{(12)})$	46.229 (0.315)	1.578	5.139 (0.109)	0.500 (0.516)

Notes: See Table 1.

Table 5: 12-Month Ahead Recursive Forecast Error Diagnostics

$\hat{S}_{t+12\|t}$	MSPE (*p*-value)	Bias	MAPE (*p*-value)	Success Ratio (*p*-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)
$S_t\left(1+\hat{\alpha}+\hat{\beta}\ln(F_t^{(12)}/S_t)\right)$	78.414 (0.523)	3.874	6.272 (0.362)	0.526 (0.032)
$S_t\left(1+\hat{\beta}\ln(F_t^{(12)}/S_t)\right)$	84.275 (0.768)	4.352	6.411 (0.623)	0.541 (0.004)
$S_t\left(1+\hat{\alpha}+\ln(F_t^{(12)}/S_t)\right)$	76.682 (0.710)	3.839	6.138 (0.427)	0.515 (0.028)
$S_t\left(1+\ln(F_t^{(12)}/S_t)\right)$	79.007 (0.916)	4.189	6.279 (0.789)	0.536 (0.021)
$S_t(1+i_{t,12})$	65.285 (0.480)	1.439	6.018 (0.804)	0.582 N.A.
$S_t(1+\hat{\alpha})$	64.709 (0.108)	3.200	5.968 (0.269)	0.552 (0.001)
$S_t(1+\Delta\overline{s}_t^{(1)})$	71.550 (0.282)	2.218	6.181 (0.303)	0.505 (0.499)
$S_t(1+\Delta\overline{s}_t^{(3)})$	68.673 (0.484)	2.268	6.056 (0.478)	0.490 (0.668)
$S_t(1+\Delta\overline{s}_t^{(6)})$	65.632 (0.314)	2.321	5.964 (0.355)	0.438 (0.966)
$S_t(1+\Delta\overline{s}_t^{(9)})$	64.931 (0.234)	2.340	5.929 (0.274)	0.469 (0.816)
$S_t(1+\Delta\overline{s}_t^{(12)})$	64.986 (0.238)	2.346	5.906 (0.199)	0.479 (0.728)
$S_{t,12}^{CF}$	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	$\hat{\alpha}$	$\hat{\beta}$	$H_0 : \alpha = 0$	$H_0 : \beta = 1$	$H_0 : \alpha = 0, \beta = 1$
3-month	0.029	1.160	0.063	0.398	0.247
6-month	0.057	0.766	0.037	0.685	0.037
12-month	0.111	0.731	0.008	0.777	0.004

Notes: For the 3- and 6-month regressions, the sample period is 1989.4-2007.2. For the 12-month regression, the sample is 1990.1-2007.2. All *t*-and *Wald*-tests have been computed based on HAC standard errors.

Table 7: Time Series Features of $\left(F_t^{(h)} - E_t\left[S_{t+h}\right]\right)/S_t$ (Percent)

	No-Change Forecast $E_t\left[S_{t+h}\right] = S_t$		Perfect Foresight $E_t\left[S_{t+h}\right] = S_{t+h}$	
	3 Month	12 Month	3 Month	12 Month
Mean	-1.12	-4.88	-3.29	-15.02
(*p*-value)	(0.00)	(0.00)	(0.02)	(0.00)
Mean Abs. Deviation	2.72	8.89	12.35	24.19
Max	12.3	30.1	55.1	42.5
Min	-10.1	-27.7	-124.9	-139.4
Persistence	0.74	0.81	0.60	0.81

Notes: The sample for the 3-month forecasts is 1983.4-2007.2; and that for the 12-month forecast is 1990.1-2007.2, reflecting the data constraints. The *p*-values of the test for a zero mean are based on HAC standard errors. The measure of persistence is the sum of the autoregressive coefficients proposed by Andrews and Chen (1994). The autoregressive lag order is determined using the AIC with an upper bound of 24 lags.

Table 8: Contemporaneous Correlation of $Basis \cdot (-1)$ and Precautionary Demand Component of Spot Price of Crude Oil (Percent)

Horizon	1989.1-2006.12	1989.1-2003.12	2004.1-2006.12
3	39.1	57.9	24.9
6	49.7	69.9	33.5
9	56.4	75.8	32.5
12	61.4	79.4	39.1

NOTES: Computed based on Figure 5 and the VAR estimates of the precautionary demand component of the spot price of crude oil in Kilian (2007).

Figure 1: Prices of Oil Futures Contracts and Spot Price
1983.3-2007.2

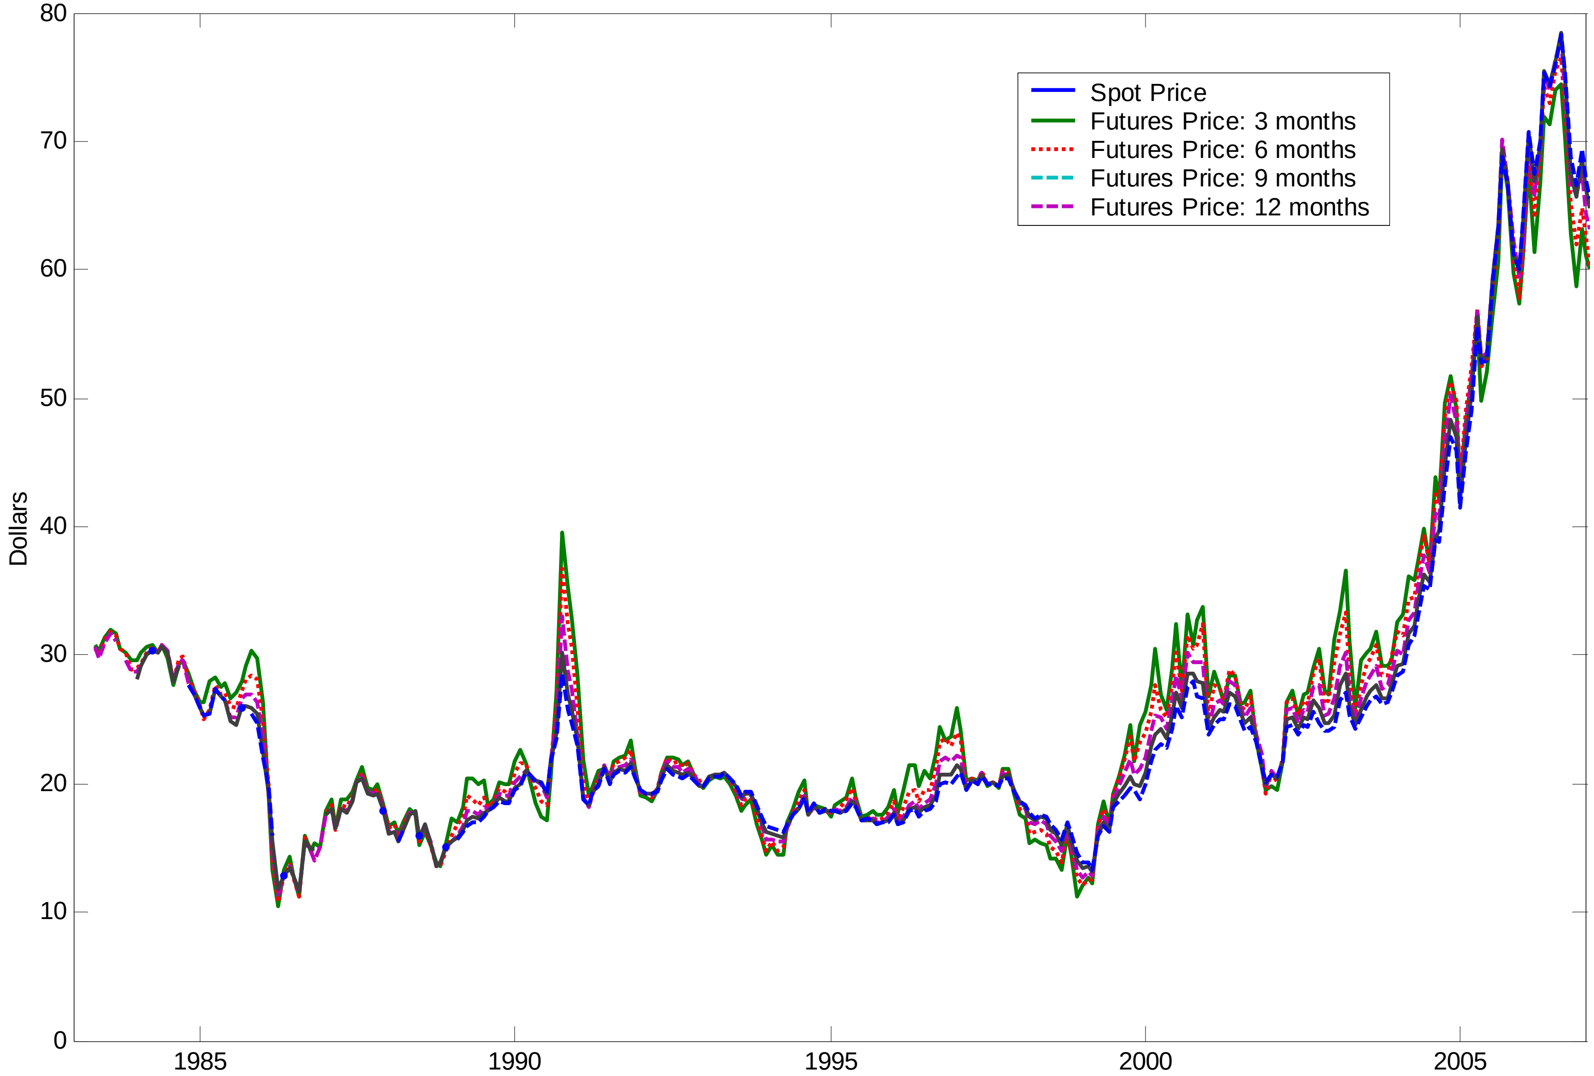

Source: Computed as described in the text based on daily NYMEX oil futures prices and the daily WTI spot price.

Figure 2: Volume of NYMEX Oil Futures Contracts

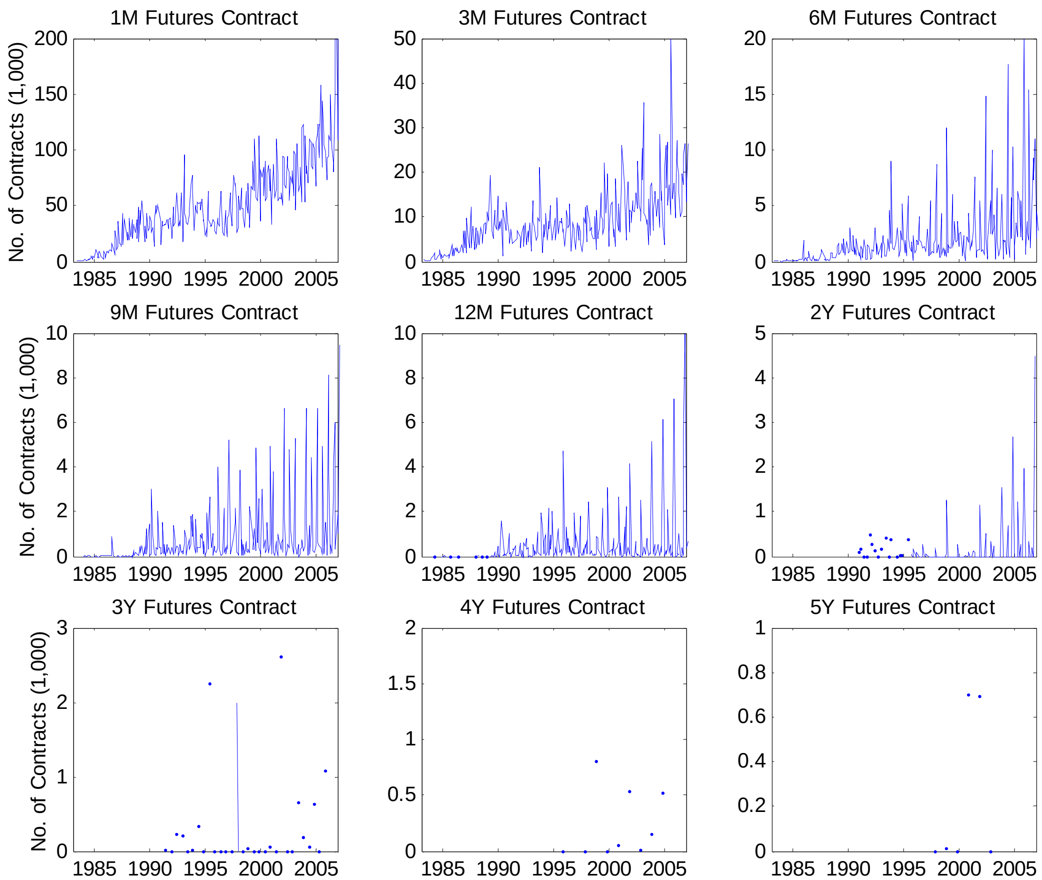

Source: *Price-data.com*

Figure 3: Basis of Oil Futures and Basis of Foreign Exchange Futures
3-Month Horizon

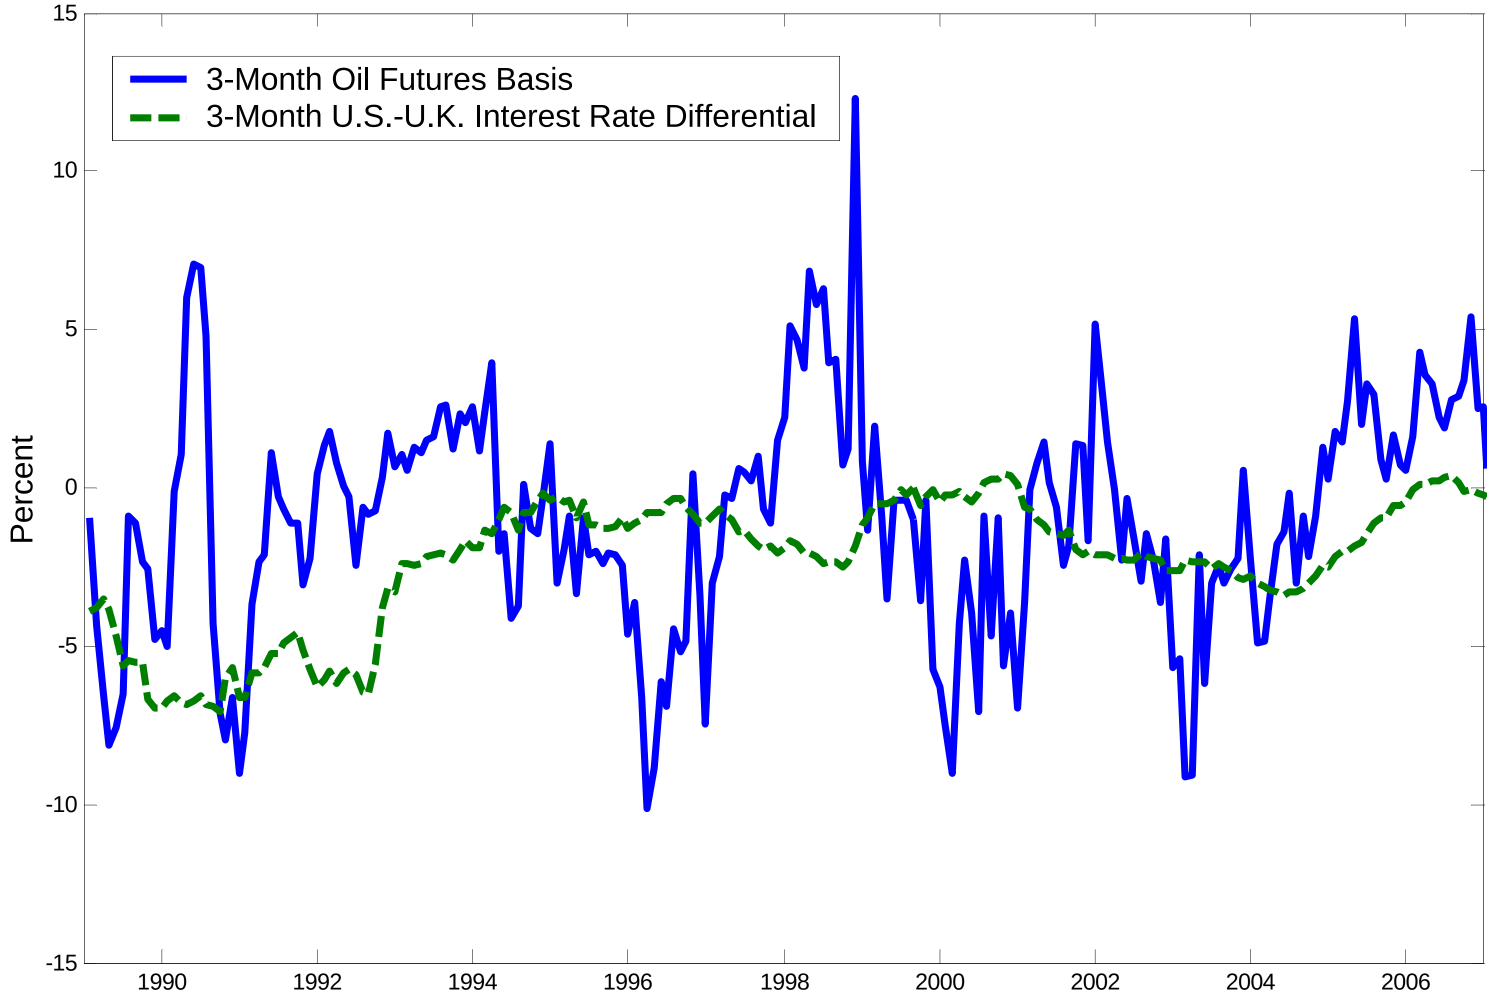

NOTES: The interest rates are end-of-month Treasury bill rates from the Bank of England and the Federal Reserve Board.

Figure 4: $Basis \cdot (-1)$ **by Horizon**
1989.1-2007.2

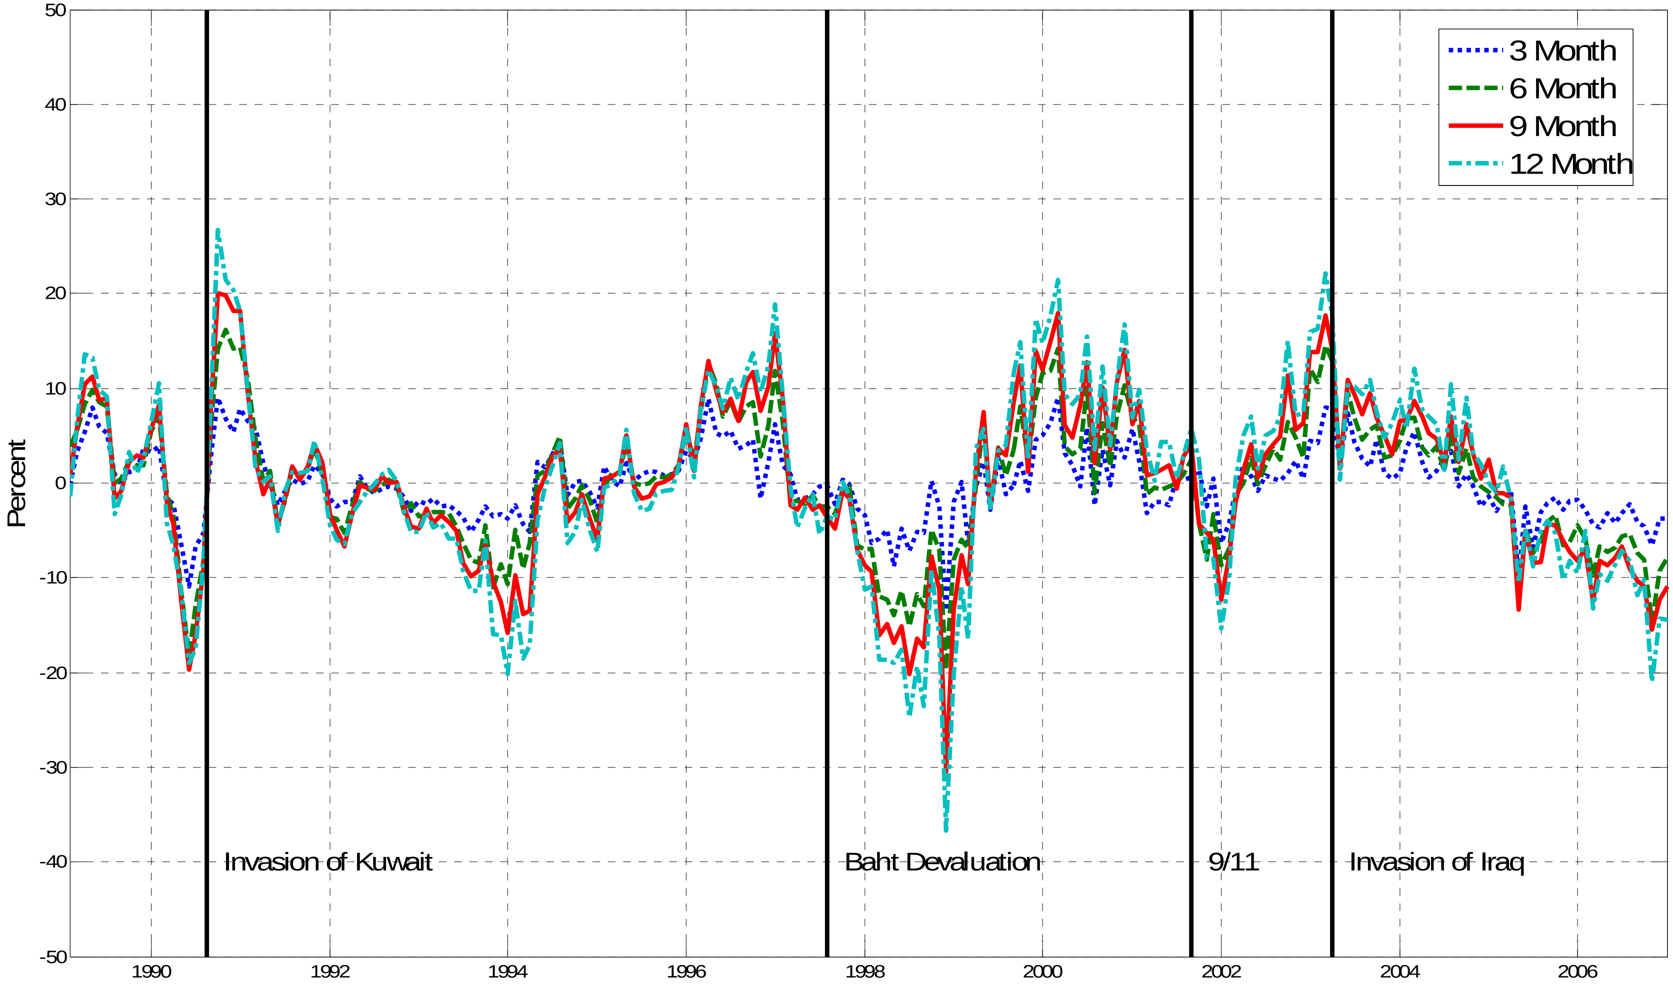

Figure 5: *Negative Basis* **and VAR-Based Estimate of Precautionary Demand Component of Spot Price 1989.1-2006.12**

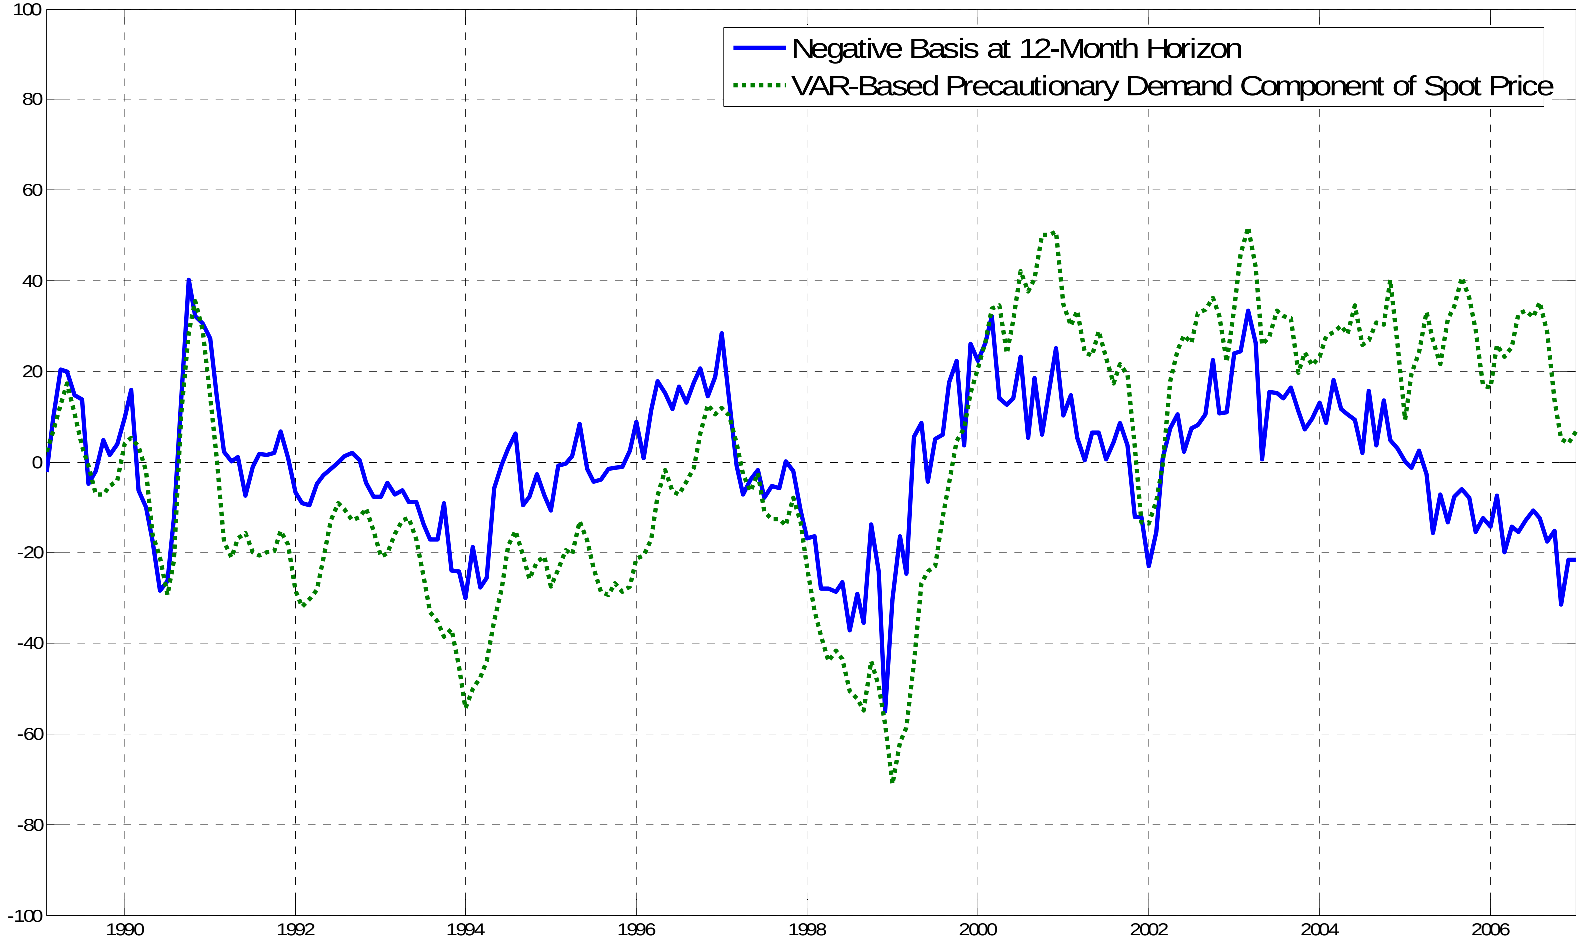

NOTES: The basis has been scaled by -1.5 to improve the readability of the graph. Since the basis is an index that transformation does not involve any loss of generality.

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

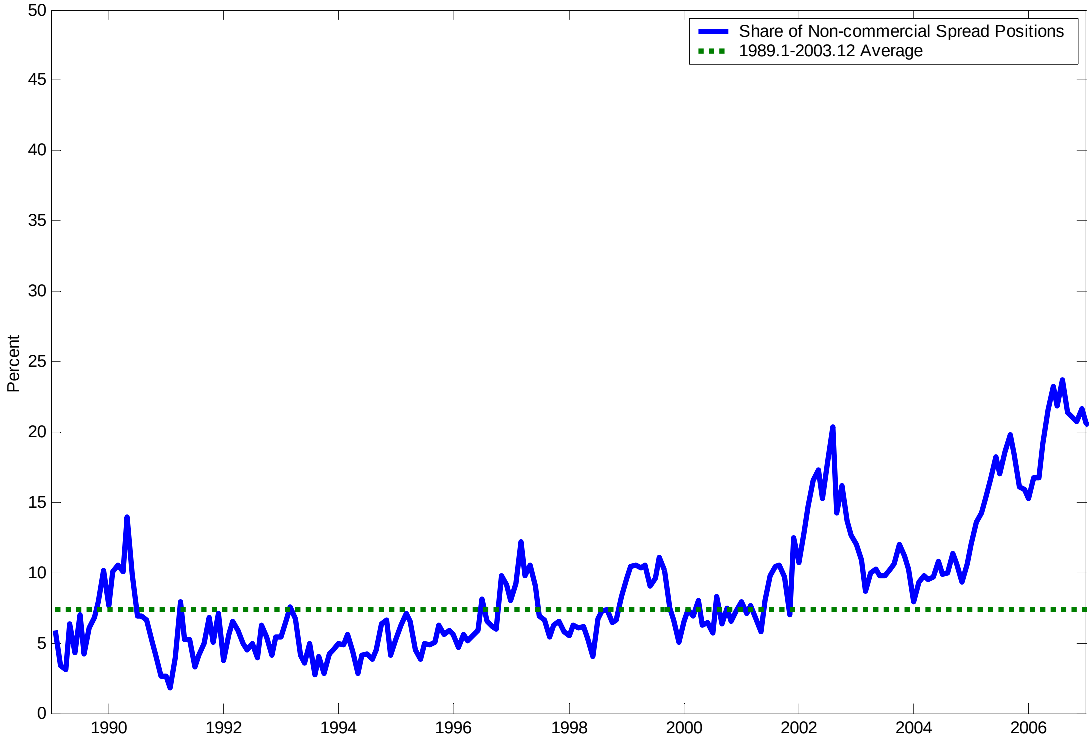

Notes: Data are end-of-month observations. Commercial firms are engaged in business activities that can be hedged in the futures market. Non-commercial firms are all other firms. Percentages refer to the share of reportable positions, which comprise 70-90% of all positions.

Figure 7a: The Instantaneous Effect of an Increase in Uncertainty on the Marginal Convenience Yield in Periods 1 and 2

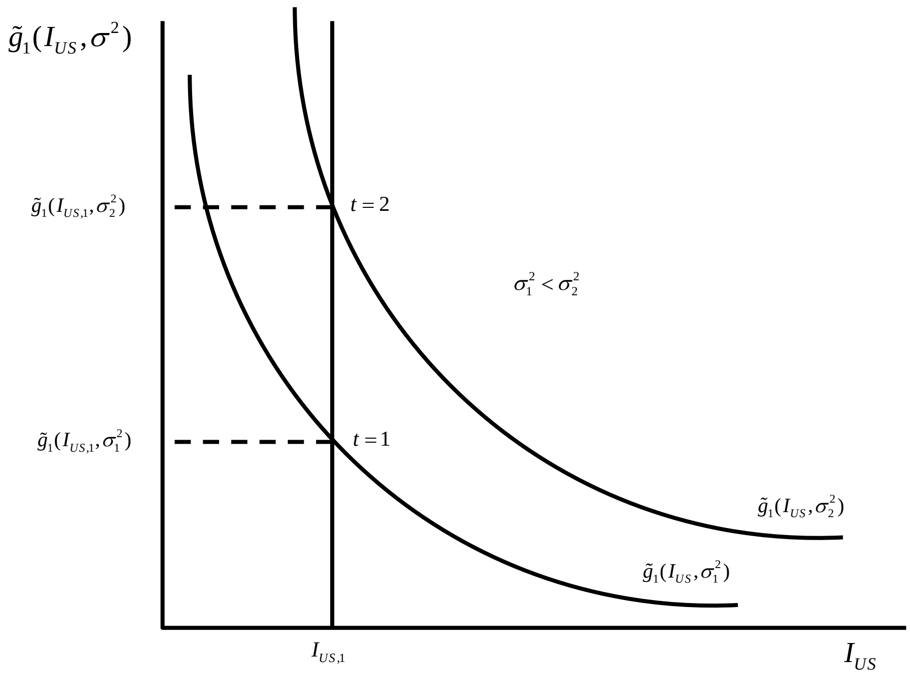

Figure 7b: The Evolution of the Marginal Convenience Yield in Period 3