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: February 1, 2008

Abstract: Despite their widespread use as predictors of the spot price of oil, oil futures prices tend to be less accurate in the mean-squared prediction error (MSPE) sense than no-change forecasts. This result is driven by the variability of the futures price about the spot price, as captured by the oil futures spread. This variability can be explained by the marginal convenience yield of oil inventories. Using a two-country, multi-period general equilibrium model of the spot and futures markets for crude oil we show that increased uncertainty about future oil supply shortfalls under plausible assumptions causes the spread to decline. Increased uncertainty also causes precautionary demand for oil to increase, resulting in an immediate increase in the real spot price. Thus the negative of the oil futures spread may be viewed as an indicator of fluctuations in the price of crude oil driven by precautionary demand. An empirical analysis of this indicator provides independent evidence of how shifts in the uncertainty about future oil supply shortfalls affect the spot price of crude oil and how they undermine the forecast accuracy of oil futures prices. Our model is consistent with a number of empirical regularities and results obtained by alternative methodologies.

Key words: Crude oil; futures market; spot market; spread; 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. In addition, we benefited from comments received at the St. Louis Fed Conference on Macroeconometrics in August 2007 and in the ECB Workshop on Forecasting Techniques in November 2007. 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: lkilian@umich.edu. Phone: (734) 647-5612. Fax: (734) 764-2769.

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 spread (defined as the percent deviation of the oil futures price from the spot price of oil).

The paper is organized as follows. In section 2, we document the use of oil futures prices as predictors of spot prices at central banks and international organizations. Futures-based forecasts of the price of crude oil inform 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 put this common practice to the test. Using a newly constructed data set of oil futures prices and oil spot prices that includes data up to February 2007, we assess empirically whether forecasts based on the price of oil futures are more accurate than forecasts from alternative models excluding futures prices. We show that forecasts based on oil futures prices and forecasts based on the oil futures spread 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). They also are more biased than the no-change forecast.

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. Moreover, it contrasts with 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.

In section 3, we conduct a systematic evaluation of the out-of-sample predictive accuracy of a broader set of oil price forecasting approaches based on the forecast evaluation period 1991.1-2007.2. 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 and much more accurate than professional survey forecasts of the price of crude oil. This makes the no-change forecast a natural benchmark.

In section 4 we show that the cause of the large mean-squared prediction error (MSPE) of futures-based forecasts relative to the no change forecast is not so much that these forecasts are so different on average, but rather the variability of the futures price about the spot price, as captured by the spread of oil futures. We document that there are large and persistent fluctuations in the oil futures spread that are unlike the fluctuations observed in the spread of foreign exchange futures (see, e.g., Taylor 1989).

In section 5, 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 1994, 2001; Routledge, Seppi and Spatt 2000; Schwartz 1997). We study the implications of the marginal convenience yield for the oil futures spread in the context of a multi-period, two-country general equilibrium model of the spot and futures markets for crude oil. We show that shifts in the uncertainty about futures oil supply shortfalls may explain the excess variability of oil futures prices relative to the spot price that is responsible for their poor predictive accuracy.

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. Using comparative statics, we establish that under plausible conditions increased uncertainty about future oil supply shortfalls causes the oil futures spread to fall. Such uncertainty shifts also raise the current real spot price of oil, as precautionary demand for oil inventories increases in response to increased uncertainty. Increased uncertainty about future oil supply shortfalls in the model will cause the real price of oil to overshoot initially with no response of oil inventories on impact, followed by a gradual decline of the real price of oil, as inventories are gradually accumulated gradually over time. The model implies that the oil futures spread declines, as the component of the real spot price of oil driven by precautionary demand for crude oil increases. Hence, the negative of the oil futures spread may be viewed as an *indicator* of fluctuations in the spot price of crude oil driven by shifts in precautionary demand for oil.

In section 6, we evaluate these predictions of our model empirically. First, we show that the proposed indicator moves as expected during events such as the Persian Gulf War that a priori should be associated with large shifts in precautionary demand for crude oil. We also find evidence of shifts in the spread 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 (VAR) models. Second, our indicator is highly correlated with an independent estimate of the precautionary demand component of the spot price of crude oil proposed in Kilian (2007a,b). 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 as high as 79 percent during 1989.1-2003.12. Third, we show that the overshooting pattern of the response of the real price of oil to a precautionary demand shock in the Kilian (2007a) VAR model is consistent with the predictions of our theoretical model. The concluding remarks are in section 7.

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

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

[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

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 as inputs in the macroeconomic forecasting exercises that these institutions produce. For example, the European Central Bank (ECB) employs oil futures prices 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. e.g., 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). 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.

There are subtle differences in how oil futures prices are interpreted by policymakers. In its strongest form, the price of oil futures is viewed as the best predictor of the spot price of oil. This interpretation is exemplified by Greenspan's (2004a) remark that "… oil futures prices … can be viewed as effective long-term supply prices." A weaker interpretation is that oil futures prices represent the expected spot price of oil. That view is illustrated by Bernanke's (2004) statement that "… futures prices of $20 a barrel suggest that traders expect oil prices to remain at that level". Before studying the theoretical support for these statements, in his section we examine their empirical support. We formally evaluate the predictive power of oil futures prices for the spot price of oil since the creation of oil futures markets in the 1980s.

2.1. Forecasting Models

2.1.1. 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:

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).

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

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.

2.1.2. Futures Prices as Future Spot Prices

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

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

2.1.3. 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 the price of oil is likely to go up or down (see, e.g., Gramlich 2004). If the futures price equals the expected spot price, as stated by Bernanke (2004), 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. The rationale for this approach is clear from dividing $F_t^{(h)} = E_t[S_{t+h}]$ by S_t, which results in $E_t[S_{t+h}]/S_t = F_t^{(h)}/S_t$. We explore the forecasting accuracy of the spread based on several alternative forecasting models. The baseline model is:

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

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

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

Alternatively, one can relax the proportionality restriction:

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

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

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

2.2. Data Description and Timing Conventions

2.2.1. Data Construction

In section 2.3, we will compare the real-time forecast accuracy of models (1)-(6). 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. 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.[2] 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.

2.2.2. 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

[2] Our approach differs from that in Chernenko, Schwarz, and Wright 2004). Their approach 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. Yet 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).

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.[3] 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. 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) explicitly referred to the 6-year oil futures contract in assessing 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.[4]

[3] 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.

[4] 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.

2.3. 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.

2.3.1. 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 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) and Bernanke (2004) quotations.

2.3.2. 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.[5]

2.3.3. 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).[6] Chernenko et al. report that the hypothesis of forecast efficiency cannot be rejected at 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 (3) 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

[5] 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.

[6] 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).

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).

3. What is the Best Predictor of the Spot Price of Oil?

The preceding section demonstrated that simple no-change forecasts of the price of oil tend to more accurate in the MSPE sense than forecasts based oil futures prices. This does not mean that the no-change forecast is necessarily the best predictor of the spot price. The first part of this section broadens the scope of forecasting methods to include other predictors. One alternative approach to measuring market expectations 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). Yet another approach is the use of more sophisticated econometric forecasting models of the spot price of crude oil.

3.1. Other Candidate Forecasting Models

3.1.1 Survey Forecasts

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:

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

3.1.2. Econometric Forecasts

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:

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

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

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

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 approach 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:

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

where $i_{t,h}$ refers to the interest rate at the relevant maturity *h*.[7] Although the Hotelling model

[7] 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.

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.[8]

3.2. Real-Time Forecast Accuracy of Other Forecasting Approaches

In this subsection, we compare the real time forecast accuracy of models (7)-(10) to that of the no-change forecast in (2). Section 2.3 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.

3.2.1. Hotelling Model

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 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.

3.2.2. 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

[8] 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

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.

3.2.3. 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.

3.3. Why the No-Change Forecast is a Plausible Measure of the Expected Spot Price

The central result of section 3.2 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.[9] 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

[9] 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 (also see Pagano and Pisani 2006). 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.

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 likely to 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.

4. Why Is the MSPE of Oil-Futures Prices so Large Relative to the No-Change Forecast?

The preceding section demonstrated that under the MSPE metric the best predictor of the nominal price of oil is the no-change forecast. This section examines in greater detail the differences between the no-change forecast and the forecast based on oil futures prices. A formal analysis of what precisely goes wrong with the oil futures forecast will help motivate the theoretical analysis of the oil spot and futures markets in the next section. For this purpose it is convenient to express the deviation of the futures price from the no-change forecast in percentage terms as $f_t^{(h)} - s_t$.

There are two possible reasons for the higher MSPE of $F_t^{(h)}$ relative to S_t. One is higher forecast bias; the other is a higher forecast variance. In Table 7, we first evaluate the possibility that $F_t^{(h)}$ is different on average from S_t. For expository purposes, we focus on the 3-month and 12-month horizons. Our sample period is 1989.1-2007.2, as a contiguous time series for the 12-month spread becomes available only starting in 1989.1. Using heteroskedasticity and autocorrelation consistent (HAC) standard errors, on average both the 3-month and 12-month spread are statistically different from zero at the 1% level. Although the rejection is decisive, Table 7 shows that the mean deviation is comparatively small in magnitude (about 1% at the 3-month horizon and below 5% at the 12 month horizon).

Whereas the bias of futures prices relative to the no-change forecast may seem small, the variability about the no-change forecast is not. As Table 7 shows, at any point in time, the discrepancy between the futures price and the spot price may be very large and go in either direction. It is this variability of the deviation of futures prices from spot prices rather than the differences in mean that drives the larger MSPE of futures-based forecasts and that makes the use of such oil price forecasts inadvisable. The time-variation in the spread is not only large, but highly persistent. In Table 7, we measure this persistence by modeling the spread 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 estimated persistence for the 3-month spread in the first column is 0.74, whereas that for the 12-month spread is 0.81.

The evidence in Table 7 is important because it suggests that the key to understanding the poor predictive accuracy of oil futures prices relative to the no-change forecast is to understand the causes of the excess variability of oil futures prices relative to the spot price of oil. The existence of such large fluctuations in the oil futures spread may seem surprising at first, considering the much lower variability and persistence of the futures spread in the widely studied foreign exchange futures market. The spread of foreign exchange futures prices over the spot exchange rate is well explained by the bilateral interest rate differential because the spread 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 spread in major foreign exchange markets tends to be small. This point is illustrated in Figure 3. The oil futures spread is far more variable than the U.S.-U.K. foreign exchange futures spread.

In the next section we propose a theoretical explanation of this discrepancy. We observe that the difference between the oil 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 propose a theoretical model that explains the persistent and large fluctuations in the spread 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 spread over time.

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

5.1. Model Description

The model in this section can be viewed as a generalization of the analysis in Pindyck (1994, 2001). 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 2008). The existence of capacity constraints implies that extracting less oil today does not permit more oil to be extracted in the future. Each period, Saudi Arabia receives a random oil endowment $\tilde{\omega}_t$. The oil endowment in period t is $\tilde{\omega}_t = \omega + \varepsilon_t$ with probability θ; and $\tilde{\omega}_t = \omega - \hat{\varepsilon}_t$ with probability $1-\theta$ and $\hat{\varepsilon}_t = \theta\varepsilon_t / (1-\theta)$ such that $E(\tilde{\omega}_t) = \omega.$ The variance of the oil endowment is σ_ε^2.

In each period, the United States chooses: (1) next period's above-ground inventory holdings of oil (I_t); (2) the number of oil futures contracts that deliver one barrel of oil next period; (3) the number of risk-free one-period bonds that yield $(1+r_{t,t+1})$, 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 price of the consumption good in period t is P_t and the spot price of oil is S_t. The price of the consumption good is the numeraire.

5.2. The United States' Demand for Oil

The United States chooses the amount of oil to use in the production of the consumption good; and the amount of oil to store as above-ground inventory. Imported oil can be transformed into the consumption good using the production function $F(Z_t)$, where Z_t is the quantity of oil the United States uses in producing the consumption good. We postulate that $F'(Z_t) > 0$, $F''(Z_t) < 0$, $F'''(Z_t) > 0$, and $\lim_{Z_t \to 0} F'(Z_t) = \infty$. The United States chooses Z_t such that the marginal product of oil equals the real price of oil in terms of the consumption good

(11) $$S_t/P_t = F'(Z_t),$$

which implies the demand schedule:

$$Z(S_t, P_t) \equiv F'^{-1}(S_t/P_t).$$

The resource constraint for crude oil is given by the identity

$$\Delta I_t \equiv \tilde{\omega}_t - Z\left(S_t, P_t\right).$$

Re-interpreting equation (11) as a demand function in ΔI_t, we obtain the inverse net demand function expressed as a function of the random Saudi oil endowment and the change in inventories:

$$\frac{S_t}{P_t} = F'(\tilde{\omega}_t - \Delta I_t) \equiv D(\tilde{\omega}_t, \Delta I_t).$$

If S_t/P_t is drawn on the vertical axis and ΔI_t on the horizontal axis, $D(\tilde{\omega}_t, \Delta I_t)$ is upward - sloping in ΔI_t.

5.3. No-Arbitrage Condition 1: The Oil Futures Market

If we are willing to impose, in addition, that both the United States and Saudi Arabia are risk neutral, as Bernanke (2004) explicitly assumed, then by the no-arbitrage condition that the expected return from holding inventories must equal the real price of oil today, it follows that

$$E_t\left[F_t/P_{t+1}\right] = E_t\left[S_{t+1}/P_{t+1}\right].$$

Using a linear Taylor series approximation, we obtain that

$$F_t \approx E_t\left[S_{t+1}\right]$$

Thus, the futures price will be an approximately unbiased predictor of the spot price.

5.4. No-Arbitrage Condition 2: The Bond Market

Under risk neutrality, the real value of a bond today must equal the discounted real present value of a bond tomorrow:

(12) $$\frac{1}{P_t} = \beta(1 + r_{t,t+1})E_t\left[\frac{1}{P_{t+1}}\right] \quad \Leftrightarrow \quad \frac{1}{\beta(1 + r_{t,t+1})P_t} = E_t\left[\frac{1}{P_{t+1}}\right].$$

A linear Taylor series approximation implies that:

$$(12') \qquad 1/\beta(1+r_{t,t+1}) \approx 1 \qquad \Leftrightarrow \qquad r_{t,t+1} \approx 1/\beta - 1$$

5.5. No-Arbitrage Condition 3: The Market for Storage

The distinguishing feature of our model is the existence of a market for storage. Storage takes the form of holding above-ground oil inventories. 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. The convenience yield is a commonly used modeling device (see, e.g., Brennan 1991; Fama and French 1988; Gibson and Schwartz 1990; Pindyck 1994; Routledge et al. 2000; Schwartz 1997). Its microeconomic foundations have been discussed in Williams (1987), Ramey (1989), Litzenberger and Rabinowitz (1995), and Considine (1997), among others. We denote the convenience yield by $g = g(I_t, \sigma_\varepsilon^2)$. Let $g_1 = g_1(I_t, \sigma_\varepsilon^2)$ denote the marginal convenience yield associated with holding additional above-ground inventories between t and $t+1$. Following the commodity pricing literature, we impose that $g_1 > 0,\ g_{11} < 0,$ and $g_{12} > 0,$ where g_i denotes the derivative of g with respect to its i^{th} argument and g_{ij} the cross-partial derivative of g with respect to the arguments i and j. As increases in the variance make production shortfalls more likely, the marginal convenience yield from holding inventories is increasing in the variance. Throughout the paper we also postulate that the Inada condition

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

holds, which ensures that the U.S. holds strictly positive inventories. With $g_1(I_t, \sigma_\varepsilon^2)$ on the vertical axis and above-ground inventory holdings on the horizontal axis, the intersection of the $g_1(I_t, \sigma_\varepsilon^2)$ curve and inventory holdings I_t describes the equilibrium in the market for storage.

Abstracting from costs of storage, no arbitrage implies that storing a barrel of oil above ground for one period and simultaneously selling short a one-period futures contract is a risk-free strategy:

$$\underbrace{(1+r_{t,t+1})g_1 + E_t\left[\frac{S_{t+1}}{P_{t+1}}\right] - \frac{S_t}{P_t}}_{\text{Return from storing oil}} + \underbrace{E_t\left[\frac{F_t}{P_{t+1}}\right] - E_t\left[\frac{S_{t+1}}{P_{t+1}}\right]}_{\text{Return from selling futures contract}} = (1+r_{t,t+1})g_1 - \frac{S_t}{P_t} + E_t\left[\frac{F_t}{P_{t+1}}\right].$$

By no arbitrage, the returns to this investment must equal the return on investing the same amount at the risk-free rate:

$$r_{t,t+1}\frac{S_t}{P_t} = (1+r_{t,t+1})g_1 - \frac{S_t}{P_t} + E_t\left[\frac{F_t}{P_{t+1}}\right].$$

Since $E_t\left[F_t/P_{t+1}\right] \approx F_t/P_t$ given equation (12), we obtain:

$$(1+r_{t,t+1})\frac{S_t}{P_t} - \frac{F_t}{P_t} \approx (1+r_{t,t+1})g_1 \tag{13}$$

Equation (13) shows that the difference between the capitalized real spot price and the real futures price is equal to the capitalized marginal convenience yield.

5.6. A Permanent Mean-Preserving Spread of Oil Endowments

In this subsection, we derive two comparative statics results under risk neutrality. The first result is that an increase in uncertainty about the future oil supply shortfalls immediately raises the real spot price of oil; the second result is that under plausible assumptions this increase in uncertainty lowers the oil futures spread. We model the increase in uncertainty as a mean-preserving increase in the spread of the oil endowment shock. The thought experiment is an increase in ε_t.

The mean preserving spread helps us abstracts 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 (2008). To keep the model tractable, we model this increased uncertainty as an increase in the conditional variance of oil supplies, implicitly abstracting from the global business cycle or any other change in the conditional mean.

5.6.1. Result 1: An Increase in Uncertainty Increases the Real Spot Price

We solve the no-arbitrage condition (13) for S_t/P_t and substitute for $1\big/\left((1+r_{t,t+1})P_t\right)$ from equation (12) to obtain

$$\frac{S_t}{P_t} = \beta E_t\left[\frac{F_t}{P_{t+1}}\right] + g_1(I_t, \sigma_\varepsilon^2)$$

$E_t\left[F_t/P_{t+1}\right] = E_t\left[S_{t+1}/P_t\right]$ by the no-arbitrage condition in the futures market. Using equation (12) to substitute for the real price of oil in terms of the marginal product, we arrive at:

$$F'(\tilde{\omega}_t - \Delta I_t) = \beta E_t\left[F'(\tilde{\omega}_{t+1} - \Delta I_{t+1})\right] + g_1(I_t, \sigma_\varepsilon^2),$$

implying that the United States equates the marginal benefits and marginal costs of these inventory holdings. The mean-preserving spread drives a wedge between the left-hand and right-hand side of this intertemporal marginal efficiency condition. Because $F'(.)$ is convex, the mean-preserving spread increases $E_t\left[F'(\tilde{\omega}_{t+1} - \Delta I_{t+1})\right]$ by Jensen's inequality (Hirshleifer and Riley 1992). It also increases the marginal willingness to pay for inventories, given by $g_1(I_t, \sigma_\varepsilon^2)$. To re-establish intertemporal marginal efficiency, the United States must increase its inventory holdings such that equality is re-established.

Figure 4 illustrates the dynamic adjustment process of the real price of oil and of U.S. oil inventories in response to an exogenous increase in uncertainty about future oil supply shortfalls. Figure 4a plots the marginal convenience yield. Figure 4b shows the corresponding inverse U.S. demand function for oil. In the model, date t inventory holdings are determined by the quantity of inventories the U.S. decided to hold at time $t-1$. Suppose that we are at point A in Figure 4a at the beginning of the period. When there is a mean-preserving increase in the endowment spread, the marginal convenience yield schedule shifts upwards instantaneously, because the U.S. values each unit of inventory more than it did prior to the increase in uncertainty. We move along the inventory schedule from point A to point B. Consequently, by the concavity of its production function, the United States finds it optimal to increase its future inventory holdings relative to last period's inventory holdings. Thus $I_{t-1} \neq I_t^*$ and $\Delta I_t = I_t^* - I_{t-1} > 0$. This implies a decrease in the real price of oil over time, starting from point B, as the United States moves along the marginal convenience yield schedule towards point C. The accumulation of additional inventories is associated with a decline in the real price of oil, as the marginal convenience yield falls. The real price of oil in the new long-run equilibrium will be higher than its level at $t-1$, but lower than its impact level. To summarize, we expect the real price of oil to overshoot in response to increased uncertainty about future oil supply shortfalls, whereas inventories will be accumulated only gradually over time. The overshooting result for the real price of oil is analogous to the overshooting of the exchange rate in the Dornbusch (1976) model. It is driven by the assumption that inventories are predetermined and will not adjust fully to an increase in uncertainty on impact.

5.6.2. Result 2: An Increase in Uncertainty Decreases the Oil Futures Spread

By rearranging equation (13), we obtain an expression for the spread:

$$\frac{F_t - S_t}{S_t} = r_{t,t+1} - (1+r_{t,t+1})\frac{g_1(I_t,\sigma_\varepsilon^2)}{S_t/P_t}. \tag{14}$$

A sufficient condition for the oil futures spread to decrease in response to a mean-preserving spread is that

$$\frac{dr_{t,t+1}}{d\varepsilon_t} - \frac{d(1+r_{t,t+1})}{d\varepsilon_t}\left[\frac{g_1(I_t,\sigma_\varepsilon^2)}{S_t/P_t}\right] - (1+r_{t,t+1})\left\{\frac{1}{F'}\left[g_{11}(I_t,\sigma_\varepsilon^2)\frac{d\Delta I_t}{d\varepsilon_t} + g_{12}(I_t,\sigma_\varepsilon^2)\frac{d\sigma_\varepsilon^2}{d\varepsilon_t}\right] + \frac{F''}{F'^2}g_1(I_t,\sigma_\varepsilon^2)\frac{d\Delta I_t}{d\varepsilon_t}\right\} < 0$$

Since $dr_{t,t+1}/d\varepsilon_t \approx 0$, the first two terms in this expression are zero. The sign of the expression depends on the relative magnitudes of (1) the decrease in the marginal convenience yield associated with the increase in inventories triggered by the shock to the endowment distribution; and (2) the increase in the marginal convenience yield associated with the increase in σ_ε^2 triggered by the same shock. The spread declines if and only if

$$\frac{d\sigma_\varepsilon^2}{d\varepsilon_t} > -\frac{1}{g_{12}}\left[g_{11} + g_1\frac{F''}{F'}\right]\frac{d\Delta I_t}{d\varepsilon_t}. \tag{15}$$

We can express both $d\sigma_\varepsilon^2/d\varepsilon_t$ and $d\Delta I_t/d\varepsilon_t$ in terms of the model's parameters and show that expression (15) is equivalent to:

$$g_{12} > \frac{\lambda(1-\theta)B}{2\theta\varepsilon_t(1-\lambda)}, \tag{15$'$}$$

where $\lambda \equiv -\left(g_{11}/g_1 + F''/F'\right)g_1/(A-g_{11})$, $0<\lambda<1$; and

$$A = -F''(\tilde{\omega}_t - \Delta I_t) - \beta E_t\left[F''(\tilde{\omega}_{t+1} - \Delta I_{t+1})\right] > 0$$
$$B = \beta\theta[F''(\omega_t + \varepsilon_t - \Delta I_{t+1}) - F''(\omega_t + \hat{\varepsilon}_t - \Delta I_{t+1})] > 0$$

Hence, for a given stock of inventories and increase in ε_t, the spread 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. It represents the sensitivity of the marginal value of inventories in response to an increase in uncertainty. The shift of g_1 reflects the fact that following an increase in uncertainty each unit of inventory has greater value as insurance against supply shortfalls. In other words, the oil futures spread will decline if agents' willingness to pay for an extra barrel of oil to be used as insurance against oil supply shortfalls increases sufficiently in response to an unanticipated shift in uncertainty. It is well-documented that during past uncertainty shocks in the crude oil market, 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. Uncertainty shocks driven by exogenous events provide an economic explanation for the large and persistent fluctuations in the spread that undermine the forecasting accuracy of oil futures prices.[10]

6. Model Evaluation

6.1. Test 1: Can the Model Explain the Poor Forecast Accuracy of Oil Futures Prices?

The theoretical model predicts that under risk neutrality $F_t^{(h)} \approx E_t\left[S_{t+h}\right]$, which is approximately the result asserted by Bernanke (2004). Given this result, it may seem puzzling at first that the forecast accuracy of oil futures prices is poor in practice. This result follows naturally from the model, however. Sections 2 and 3 established that the best proxy for $E_t\left[S_{t+h}\right]$ is the no-change forecast S_t. There is no presumption in the theoretical model that $E_t\left[S_{t+h}\right] = S_t$. In fact, equation (14) implies that in equilibrium $F_t^{(h)}$ may be larger than, smaller than or equal to S_t. Thus, the evidence in Table 7 that on average over our sample period $F_t^{(h)}$ is slightly smaller than S_t is fully consistent with the theoretical model. Nevertheless, on average the (approximate) model expectation $F_t^{(h)}$ is fairly close to the econometric expectation S_t.

[10] Earlier we documented that the oil futures spread is highly persistent, but mean reverting (see Table 7). We also documented that the no-change forecast is the best predictor of the nominal spot price of oil. The conclusion that under plausible conditions the mean-reverting spread is associated with changes in the precautionary demand component of the spot price may seem to contradict 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 real 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.

The reason that $F_t^{(h)}$ is a poor predictor is not so much that it is different on average from S_t, but that it fluctuates widely relative to S_t. At any point in time, the discrepancy between the futures price and the 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 spread implies that the futures price may differ from the best predictor by as much as $20 in one direction or by as much as $18 in the other (see Table 7). Thus, policymakers relying on oil futures prices are likely to overestimate or underestimate the expected price of oil substantially at any given point in time, and the fact that these mistakes largely average out in the long run is of little consolation. Put differently, it is not that Bernanke's (2004) assertion that oil futures prices can be viewed as expected spot prices is necessarily wrong, but that it of limited use in practice given the large fluctuations in the futures price relative to the best predictor. Our theoretical model provides an explanation of this excess variability. In the model, fluctuations in the spread arise naturally from shifts in uncertainty about future oil supply shortfalls and will be indicative of fluctuations in the spot price of oil driven by precautionary demand for crude oil, provided g_{12} is large enough.[11] Hence, the theoretical model helps us understand the poor forecast accuracy of oil futures prices. Whether this explanation is empirically plausible is a question that we turn to next.

6.2. Test 2: Does the Proposed Indicator Move as Expected During Known Episodes of Uncertainty Shifts?

One way of judging the empirical content of the model is to verify that the spread moves in the expected direction at times of major unforeseen events such as the outbreak of the wars. In Figure 5, 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 as the Persian Gulf War and the 2003 Iraq War (which should have caused the spread to fall) and the Asian financial crisis and 9/11 (which should have caused the spread 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 spread, but

[11] 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 (2008) documents that capacity constraints in world crude oil production have been binding since the early 1990s.

arguably they are the most important reason.

Figure 5 plots the negative of the spread 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 (2008) 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 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 the spike 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. Test 3: How does the Proposed Indicator Compare to Alternative Measures of Precautionary Demand Shifts?

The indicator 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 (including crude oil), 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 shifts in the demand for oil that are specific to the oil market such as shifts in the precautionary demand for crude oil. 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 changes in the uncertainty 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 oil futures-based measure is essentially 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 despite 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 spread. 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 spread seems implausibly high, as discussed above. 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.

Not only does the correlation weaken after 2003.12, but the spread data and the VAR-based measure of the precautionary demand component of the spot price of oil paint a somewhat different picture (see Figure 6). 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 spread as an indicator of fluctuations in the precautionary demand component of the spot price over the last three years of the sample. These observations suggest that a structural change may have occurred around 2003.12 that is beyond the scope of the theoretical model in section 5. 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. Data from the Commodity Futures Trading Commission (not shown to conserve space) shed light on the share of speculators among oil futures traders since 1989 and reveal an unprecedented increase in speculative activities after 2003.12. To the extent that increased speculative trading tends to raise the price of oil futures more than the spot price (and hence increases the spread), this fact might provide an explanation for the weakening of the correlations at the end of the sample. Establishing such a link is left for future research.

6.4. Test 4: Does the VAR Response of the Real Price of Oil Match the Model Predictions?

Another implication of the theoretical model is that the real price of oil will overshoot in response to a mean-preserving spread, while inventory holdings will increase only gradually. If the Kilian (2007a) measure of the precautionary demand shock is valid, the response of the real price of oil in that VAR should exhibit overshooting, as predicted by the theoretical model. Figure 7 confirms that the response to an oil-specific demand shock indeed displays overshooting, suggesting that the interpretation of this shock as a precautionary demand shock is justified and indirectly supporting the interpretation of the futures-based indicator as a measure of expectations shifts. There is no evidence of such a pattern in response to other oil demand or oil supply shocks.

6.5. Implications for Crude Oil Inventories

Whereas we have focused on the empirical relationship between increased concerns about future oil supply shortfalls and the precautionary demand component of the real 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 shifts in uncertainty about future oil supply shortfalls. First, whereas for the real price of oil we were able to use a VAR decomposition to focus specifically on the precautionary demand component of the real price, no similar measure of the precautionary demand component in oil inventories exists, making it impossible to identify the consequences of precautionary demand shocks for inventories. 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.

There is, however, anecdotal evidence from oil industry experts documenting that shifts in precautionary demand coincide with a strong motive for inventory accumulation. This situation has been aptly described by Terzian (1985) 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.” (p. 260)

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". These accounts are consistent with the implications of our theoretical model.

7. 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 oil futures spread. The key insights can be summarized as follows: First, it is widely believed that prices of oil futures are accurate predictors of forecast spot prices in the MSPE sense. Using observations up to February of 2007, we showed that the price of crude oil futures is not an accurate predictor of the spot price of crude oil. Many users of oil futures-based forecasts are aware of this caveat and understand that futures-based forecasts may be poor, but still believe that they 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.

Second, we showed that the large MSPE of oil futures-based forecasts is driven not by the bias, but by the variability of the futures price about the spot price. We documented large and time-varying deviations of oil futures prices from the spot price of oil, as measured by the oil futures spread. 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, which helps explain the poor forecasting accuracy of oil futures prices.

Third, our analysis demonstrates that fluctuations in the oil futures spread are larger and more persistent than fluctuations in the spread of foreign exchange futures. 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. We proposed a theoretical model of the oil spot market and oil futures market that incorporates this marginal convenience yield. The model implies that the oil futures spread is directly linked to shifts in oil market fundamentals. We showed that shifts in the uncertainty about future oil supply shortfalls cause fluctuations in the oil futures spread not found in models of the foreign exchange futures market. Our model explains the excess variability of oil futures prices relative to the no-change forecast and the resulting poor forecast accuracy of oil futures prices.

Fourth, we showed that, under plausible conditions, the oil futures spread will decline, as the precautionary demand component of the real spot price of crude oil increases. Thus, the negative of the spread may be viewed as an indicator of fluctuations in the real price of crude oil driven by precautionary demand for oil. The time path of the oil futures spread since 1989 suggested major shifts in precautionary demand for oil during the Persian Gulf War and following the Asian crisis, for example. These results provided 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 play an important role in explaining oil price fluctuations (see, e.g., Kilian 2008). In addition, we documented that our measure of oil price movements driven by uncertainty shifts matches up well with independently obtained VAR based measures, and that our model predicts the overshooting of the real price of oil found in VAR analysis. Our analysis also is consistent with anecdotal evidence of hoarding in oil inventory markets during times of crisis.

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. (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.

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.

Considine, T. J. (1997), "Inventories Under Joint Production: An Empirical Analysis of Petroleum Refining," *Review of Economics and Statistics*, 79(3), 493-502.

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

Hotelling, H. (1931), "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39(2), 137-175.

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 (invited by the *Journal of Economic Literature*).

Kilian, L. (2008), "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

Litzenberger, R. H. and N. Rabinowitz (1995), "Backwardation in Oil Futures Markets: Theory and Empirical Evidence," *Journal of Finance*, 50(5), 1517-1545.

Neuberger, A. (1999), "Hedging Long-Term Exposures with Multiple Short-Term Futures Contracts," *Review of Financial Studies*, 12(3), 429-459.

New York Mercantile Exchange (2007a), "Light Sweet Crude Oil: Market Data," http://www.nymex.com/lsco_fut_descri.aspx.

New York Mercantile Exchange (2007b), "Light Sweet Crude Oil: Contract Detail," http://www.nymex.com/WS_spec.aspx.

Pagano, P., and M. Pisani (2006), "Risk-Adjusted Forecasts of Oil Prices," Discussion Paper No. 585, Banca D'Italia.

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.

Ramey, V. (1994), "Inventories as Factors of Production and Economic Fluctuations," *American Economic Review*, 79(3), 338-354.

Routledge, B.R., Seppi, D.J., and C.S. Spatt (2000), "Equilibrium Forward Curves for Commodities," *Journal of Finance*, 55(3), 1297-1338.

Schwartz, E.S. (1997), "The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging," *Journal of Finance*, 52, 923-973.

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.

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

Williams, J. (1987), "Futures Markets: A Consequence of Risk Aversion or Transaction Costs?" *Journal of Political Economy*, 95(5), 1000-1023.

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 $f_t^{(h)} - s_t$
(Percent)

	3 Month	12 Month
Mean	-1.12	-4.88
(*p*-value)	(0.00)	(0.00)
Mean Abs. Deviation	2.72	8.89
Max	12.3	30.1
Min	-10.1	-27.7
Persistence	0.74	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 $s_t - f_t^{(h)}$ and the VAR Estimate of the Precautionary Demand Component of Real Spot Price of Crude Oil (Percent)

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

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 of Oil
1983.3-2007.2

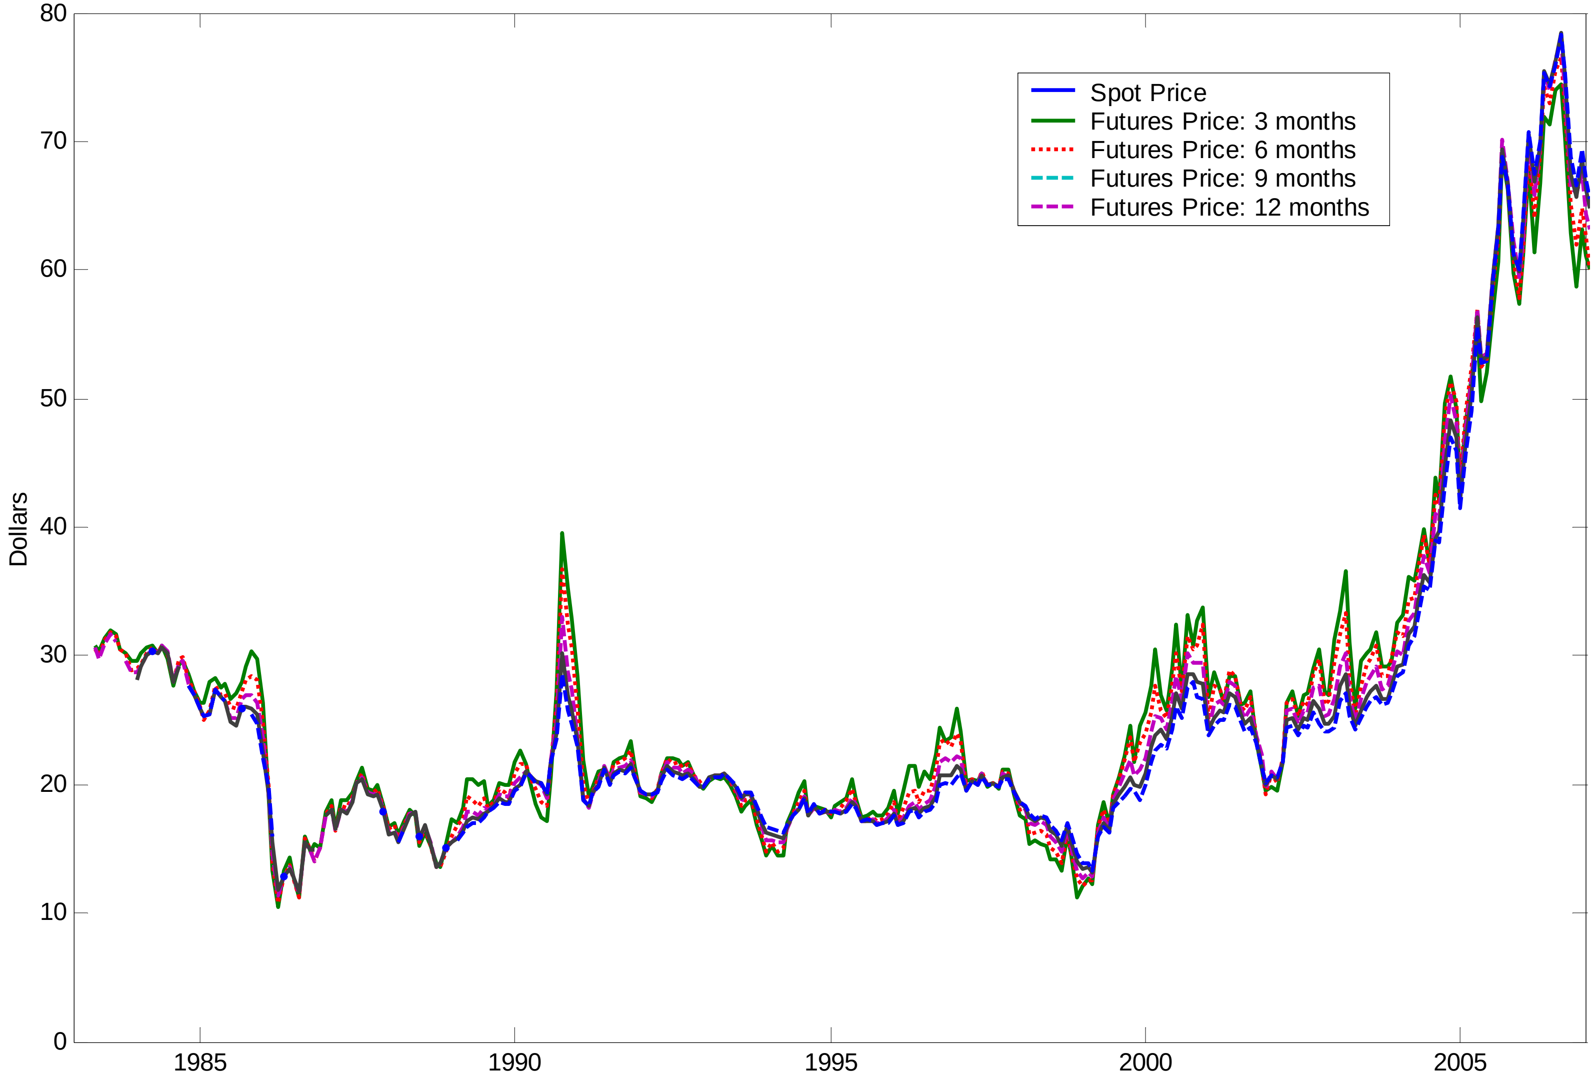

Source: Computed as described in the text based on daily NYMEX oil futures prices and the daily WTI spot price.

Figure 2: Volume of NYMEX Oil Futures Contracts

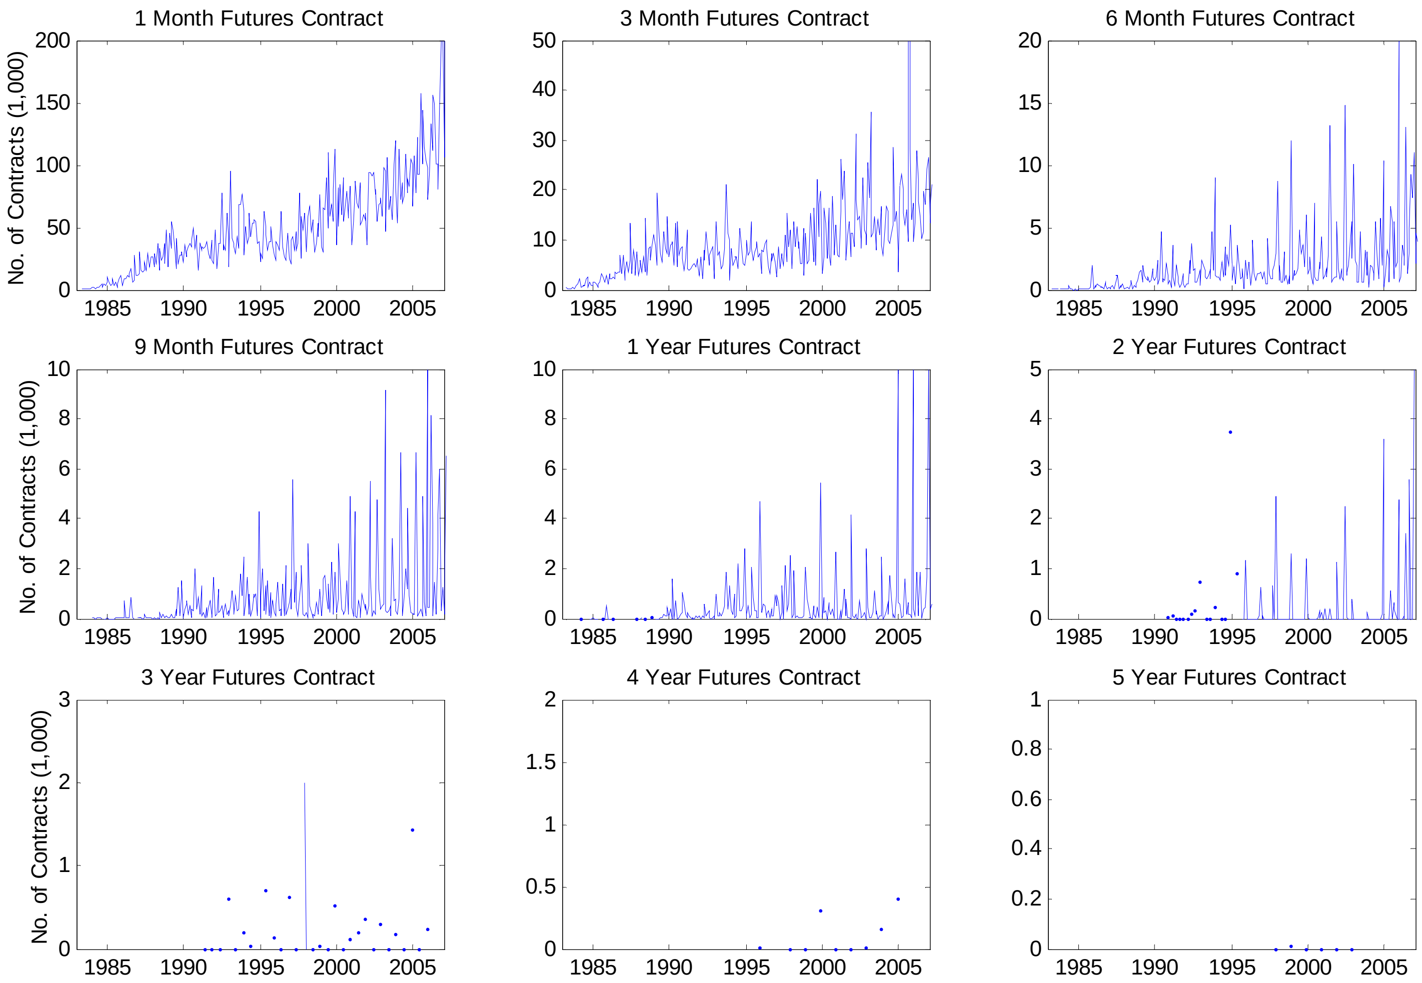

Source: *Price-data.com*

Figure 3: Oil Futures Spread and Foreign Exchange Futures Spread
3-Month Horizon

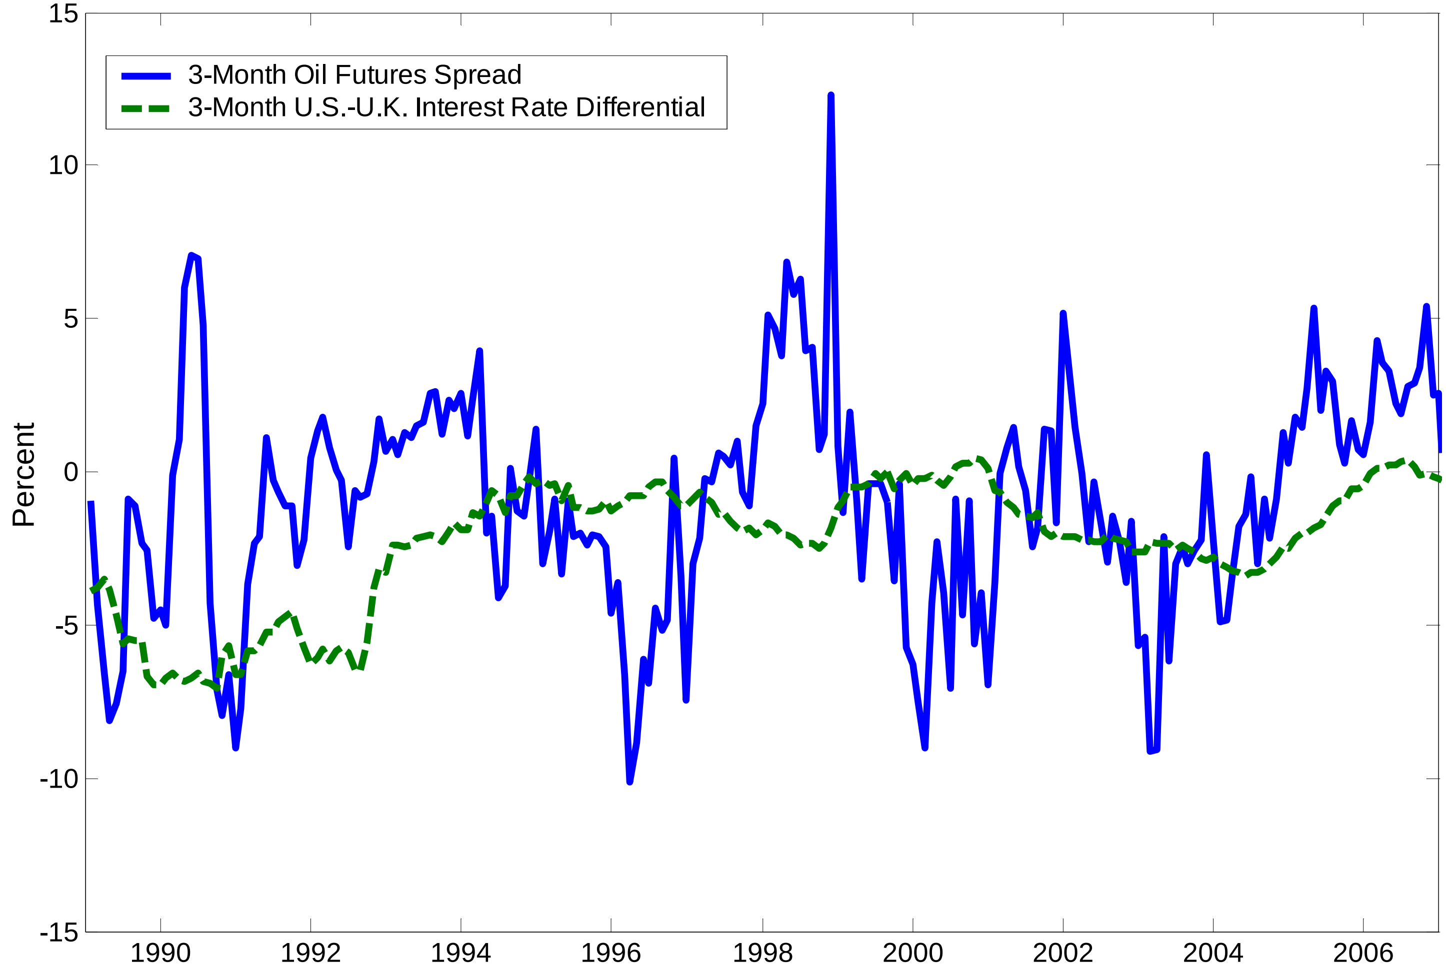

NOTES: The interest rates are end-of-month Treasury bill rates from the Bank of England and the Federal Reserve Board.

Figure 4a: The Effect of an Increase in Uncertainty on the Marginal Convenience Yield

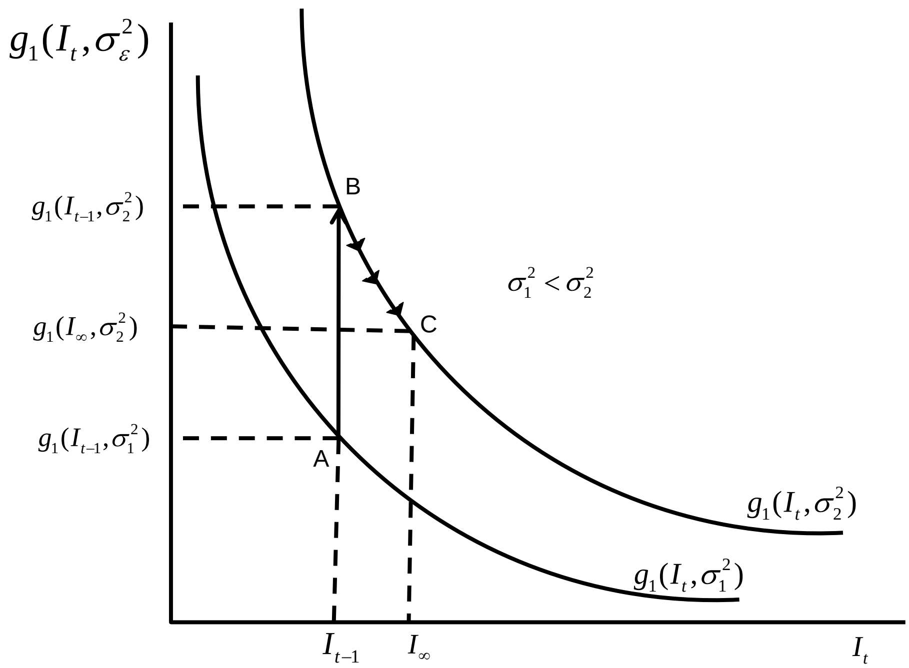

Figure 4b: The Effect of an Increase in Uncertainty on the Demand for Oil

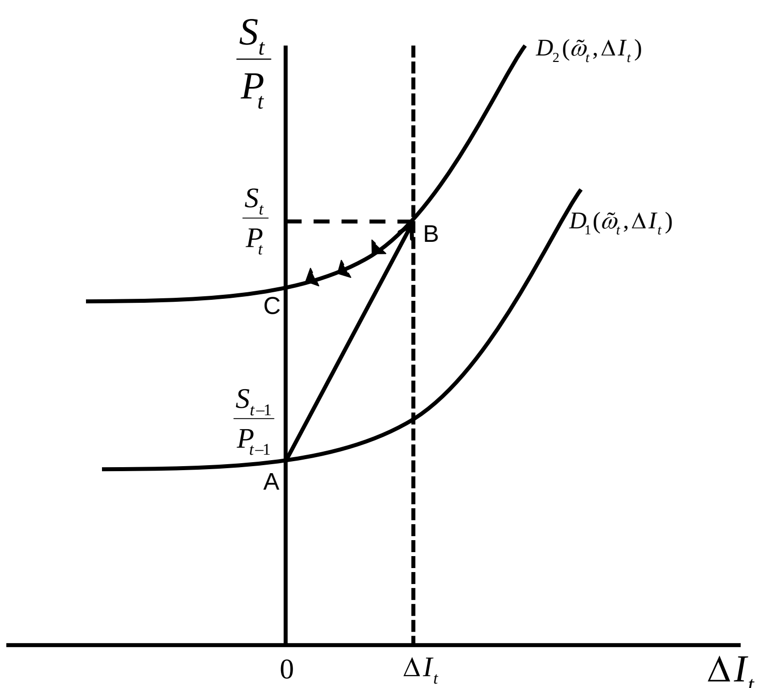

Figure 5: $s_t - f_t^{(h)}$ by Horizon
1989.1-2007.2

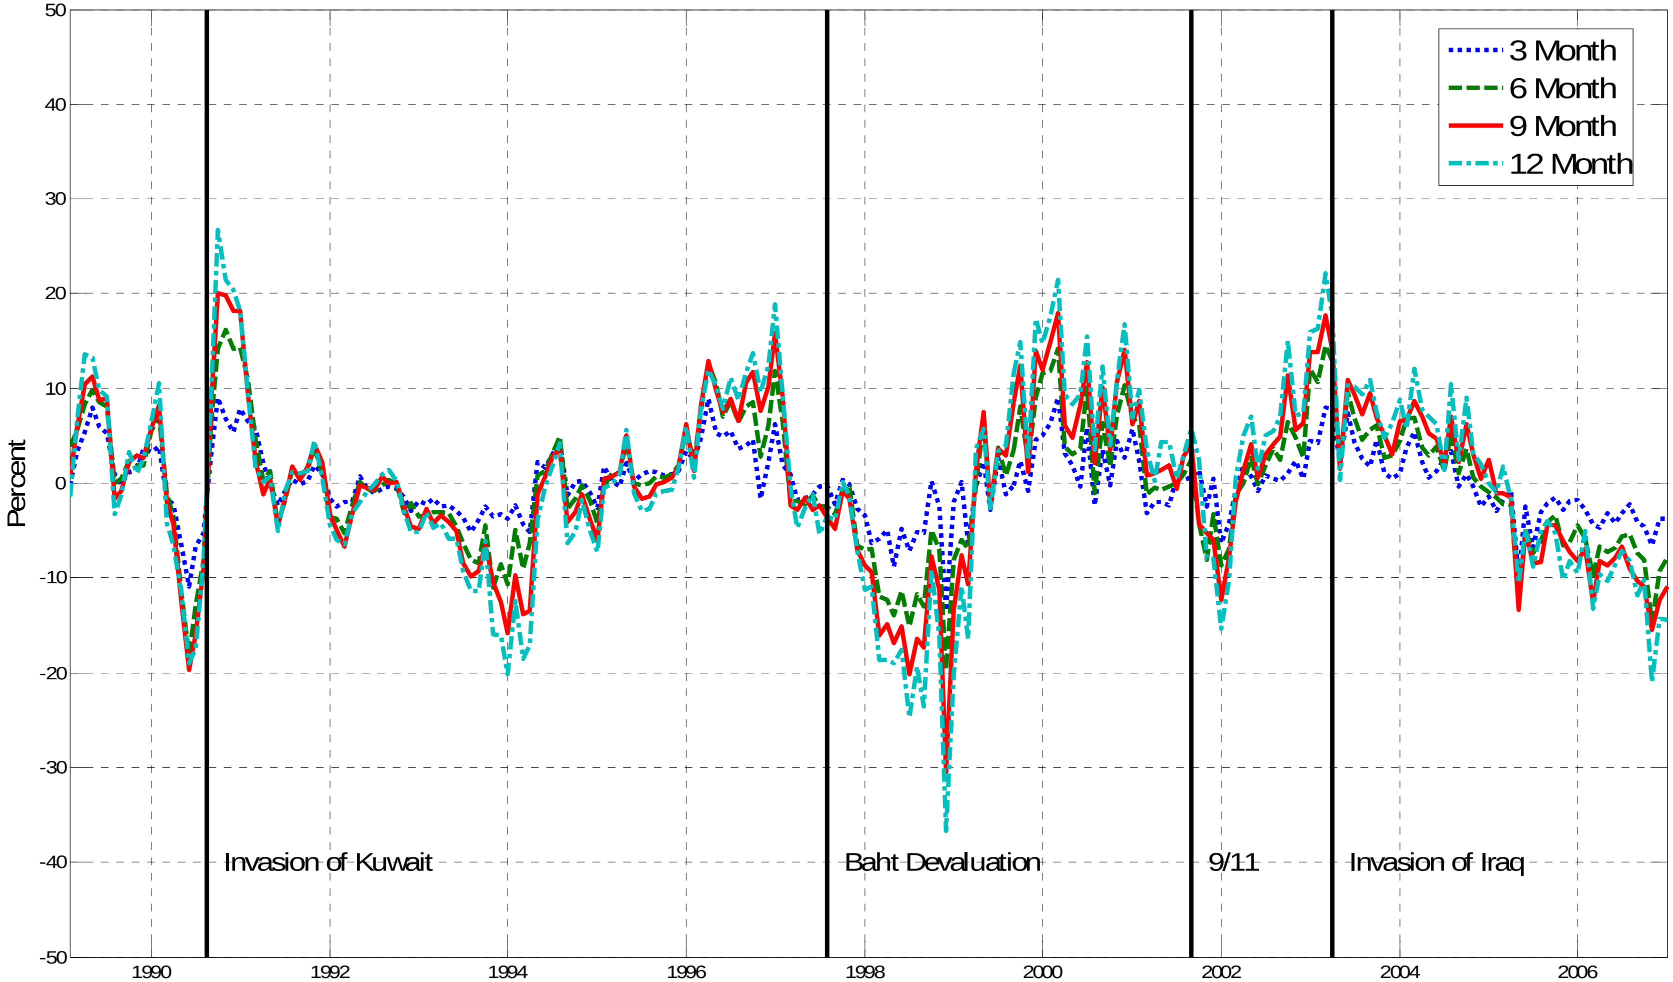

Figure 6: $s_t - f_t^{12}$ and VAR-Based Estimate of Precautionary Demand Component of Spot Price 1989.1-2006.12

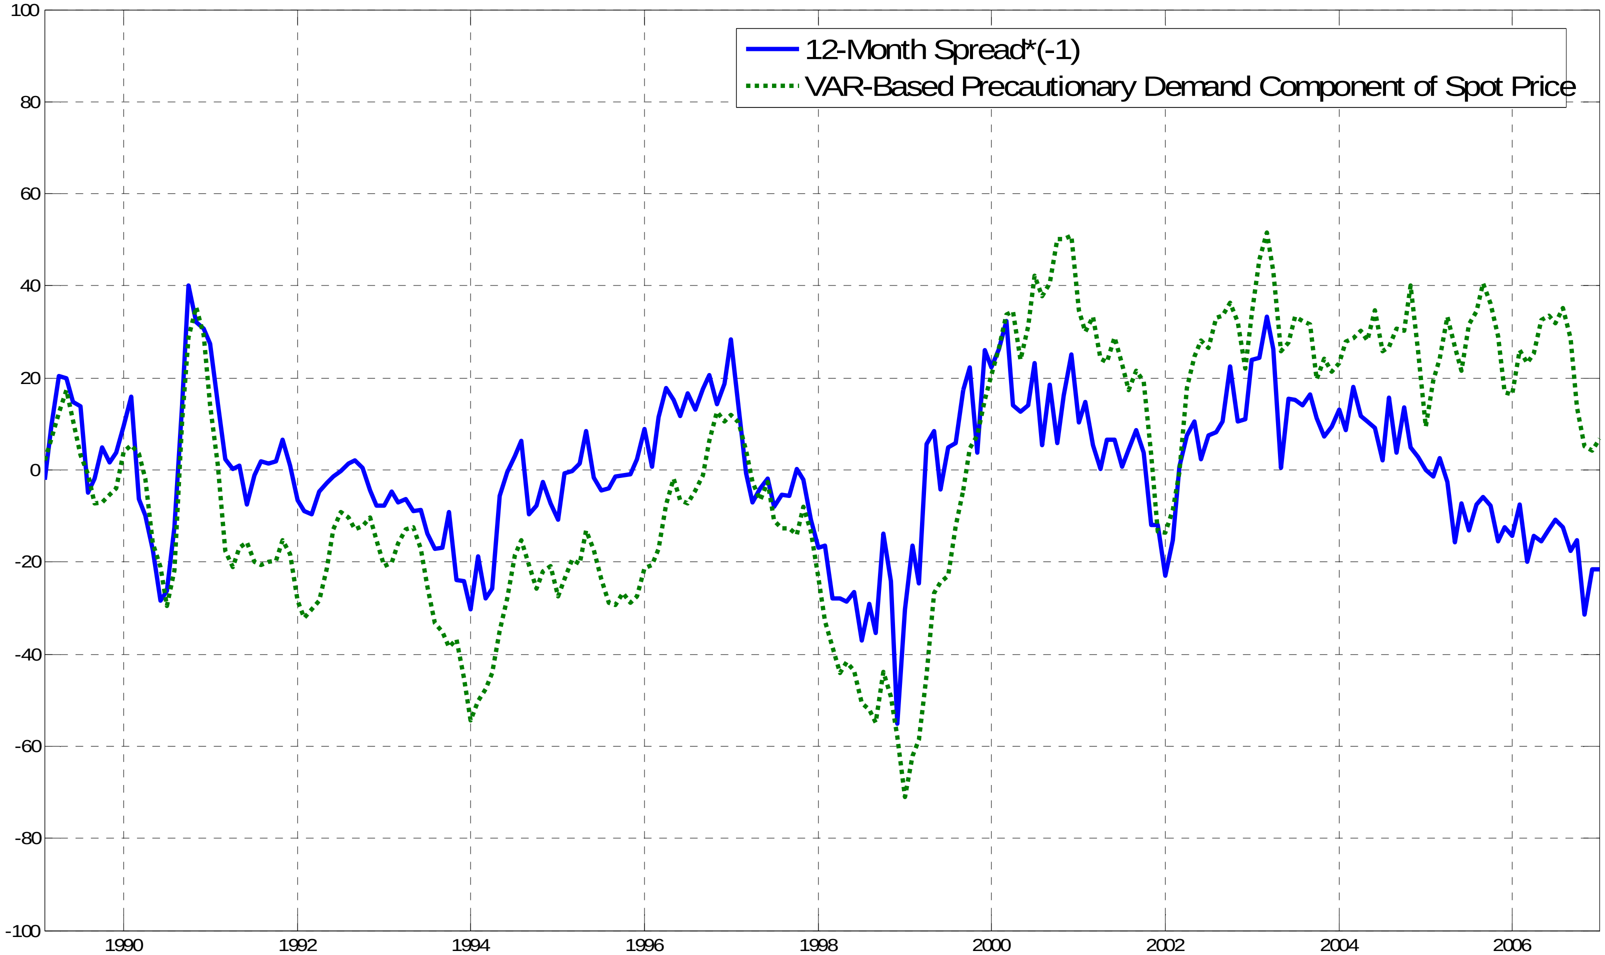

NOTES: The spread has been scaled by -1.5 to improve the readability of the graph. Since the spread is essentially an index that transformation does not involve any loss of generality.

Figure 7: Response of the Real Price of Oil to a Positive Precautionary Demand Shock
Point Estimate with One- and Two-Standard Error Bands

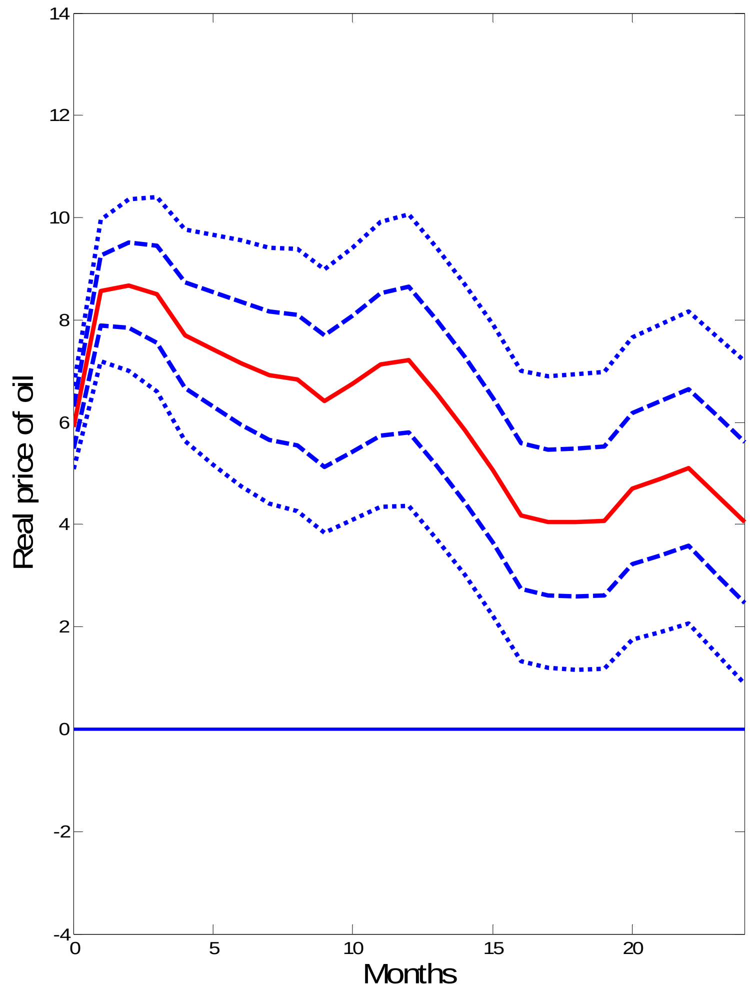

SOURCE: Kilian (2007a). The sample period is 1973.2-2006.12. The model is described in the text.