# 3rd International Workshop on Empirical Methods in Energy Economics (EMEE2010)

**Surrey Energy Economics Centre (SEEC)**
**University of Surrey, UK**
**$24^{th}$ – $25^{th}$ June 2010**

***NOTE:***

**The following Abstract and/or Paper is *Work in Progress* for presentation and discussion at the EMEE2010 workshop. It therefore must not be referred to without the consent of the author(s).**

*Sponsored by:*

# Oil Prices: Heavy Tails, Mean Reversion and the Convenience Yield[1]

Jean-Thomas Bernard[2]
Université Laval

Lynda Khalaf[3]
Carleton University

Maral Kichian[4]
Bank of Canada

Sebastien McMahon[5]
Fédération des Caisses Desjardins

Revised April 16, 2010

[1]This work was supported in part by the Institut de Finance Mathematique de Montreal, the Canadian Network of Centres of Excellence [program on Mathematics of Information Technology and Complex Systems], the Social Sciences and Humanities Research Council of Canada, the Fonds FQRSC (Government of Quebec), the Chair on the Economics of Electrical Energy, and the Canada Research Chair in Environment. All remaining errors are our own and should not be attributed to the Bank of Canada.

[2]Département d'économique and Groupe de Recherche en économie de l'énergie de l'environement et des ressources naturelles (GREEN), Université Laval. Mailing address: Pavillon J.-A.-De Sève, St. Foy, Québec, Canada, G1K 7P4. Tel: (418) 656 5123; FAX: (418) 656 7412. Email: jtber@ecn.ulaval.ca.

[3]Economics Department, Carleton University, Centre interuniversitaire de recherche en économie quantitative (CIREQ), and GREEN. Mailing address: Economics Department, Carleton University, Loeb Building 1125 Colonel By Drive, Ottawa, Ontario, K1S 5B6 Canada. Tel: (613) 520-2600-8697; FAX: (613)-520-3906. Email: Lynda_Khalaf@carleton.ca.

[4]Canadian Economic Analysis Department, Bank of Canada, Ottawa, Canada. Tel: (613) 782 7460; FAX: (613) 782 7163. Email: mkichian@bankofcanada.ca.

[5]Division Risque de crédit pour les particuliers et petites entreprises, Fédération des Caisses Desjardins, Lévis, Québec, Canada. Tel: (418) 835-8444-3488. Email: sebastien.mc.mahon@desjardins.com.

**Abstract**

Empirical research on oil price dynamics reveals statistical support for various models of price paths, yet many of the competing models differ importantly with respect to their fundamental temporal properties. In this paper we examine alternative classes of models that allow for non-constancies in level or in volatility from a forecast-based perspective. Three specifications are considered: (i) random-walk models with GARCH and normal or student-$t$ innovations, (ii) Poisson-based jump-diffusion models with GARCH and normal or student-$t$ innovations, and (iii) mean-reverting models that allow for uncertainty in equilibrium price and for time-varying convenience yields. We compare forecasts in real time, for 1, 3 and 5 year horizons. Results based on future price data ranging from 1986 to 2007 strongly suggest that imposing the random walk for oil prices has pronounced costs for out-of-sample forecasting. Evidence in favor of price reversion to a continuously evolving mean underscores the importance of adequately modeling the convenience yield.



**Key words**: Oil Price; Convenience Yield; Mean Reversion.

**Journal of Economic Literature classification:** Q4, C52, C53, G1.

# 1 Introduction

The first oil shock in 1973 constitutes a milestone in the evolution of oil price in terms of both level and volatility. For a period of more than ninety years running from 1880 to early 1970, the oil price was less than 20$/barrel in today value most of the time and displayed little volatility. Since 1973, a new regime has been ushered in where the price of oil is on average more than twice higher than before and its volatility is fifty percent higher.[1] There is already a very large empirical literature on modeling oil price dynamics for various purposes including forecasting.[2] While volatility clusters, heavy tails and structural breaks are almost unanimously acknowledged as stylized facts, statistical support is claimed for many different types of oil price models and thus there is no agreement on a uniformly best-fit class of models. Of course it is seldom possible to single out a most-relevant model or specification that persists as a practical descriptive or forecasting device; yet in the case of oil markets, empirical conflicts have serious implications for our understanding of price dynamics because many of the competing models differ importantly with respect to their fundamental temporal properties.

Important examples of such disagreements include: (i) discerning deterministic from stochastic trends; (ii) disentangling structural change in fundamentals from inherent fluctuations; (iii) allowing for unexpected discontinuities in price levels or price changes; (iv) non-constancy of the variance of prices and of price changes; and (v) non-constancy of the convenience yield.[3] Such questions are founded on theory and are typically taken

[1] See Smith (2009). BP (2009) presents a graph of the oil price from 1860 to 2008.

[2] For recent references, statistical results and critical discussions, the reader may refer to Abosedra and Laopodis (1996), Berck and Roberts (1996), Wilson, Aggarwal and Inclan (1996), Ahrens and Sharma (1997), Schwartz (1997), Pindyck (1999), Schwartz and Smith (2000), Pindyck (2001), Cortazar and Schwartz (2003), Cortazar and Naranjo (2006), Lee, List and Strazicich (2006), Moshiri and Foroutan (2006), Postali and Picchetti (2006), Sadorsky (2006), Regnier (2007), Tabak and Cajueiro (2007), Alquist and Kilian (2009), Crespo-Cuaresma, Jumah and Karbuz (2009), Hamilton (2009), and Smith (2009).

[3] The convenience yield will be defined explicitly later on.

to the data by applying state of the art econometric methods. Well known examples of underlying equilibrium models include dynamic Hotelling-type models for non-renewable resource markets, storage and inventory models for commodity markets, and financial-theory based risk and hedging models for options and futures markets.[4] Commonly used econometric tools include a plethora of unit root and break tests, a battery of generalized autoregressive conditional heteroskedasticity [GARCH] based inference methods, Kalman-filter based time-varying-parameter or neural-network based estimations, or jump-diffusion based procedures.

In this paper we examine the issue of mean-reversion in oil price from a forecast-based perspective. We consider alternative classes of models that allow for non-constancies in level or in volatility. Our forecast-based analysis is partly motivated by Pindyck (1999) who argues that unit roots tests are inconclusive in the analysis of real oil prices observed on yearly basis. Formally, Pindyck expresses the minimum sample size for which the Dickey-Fuller test is significant when data are generated by a stationary auto-regressive process in terms of the autocorrelation coefficient. Given the persistence characterizing oil price, he concludes that a very long and practically unavailable series is required to perform reliable tests. Of course, this reasoning does not account for the added complications arising from structural shifts which are more likely the longer the series. Recent work applying random-walk tests which correct for breaks provides evidence against the unit-root hypothesis, yet debate is still on-going [see *e.g.* Lee et al. (2006), Postali and Picchetti (2006), and Hamilton (2009)].

From the econometric perspective, robustness of the unit-root-with-breaks class of tests to nonlinearity such as GARCH or random jumps-in-the-mean and/or to conditional non-normality may be questioned. In this paper, we do not aim to take a stand on the worth of such tests for the problem at hand, but we do believe that available evidence on

[4] For a recent exposition, see Hamilton (2009) and Smith (2009). See also Wirl (2008) and Alquist and Kilian (2009).

the importance of parameter non-constancies, whether for refuting or for justifying the unit-root hypothesis, should be taken seriously. Formally, one of the main conclusions to be drawn from this literature, at least as we interpret it, is that structural discontinuities should be accounted for in examining stochastic models for oil prices whether one adopts a unit-root or a mean-reverting model.

There are theoretical as well as practical reasons to suspect that unit-root or Geometric Brownian Motion models are not appropriate to model natural resources or commodity prices. In particular, demand and supply pressures and non-constant convenience yields in commodity markets suggest mean reversion to long-run equilibrium prices that may change over time due to resource depletion, technological change or product innovation. Demand and supply pressures can be intuitively seen to work as follows. When prices are higher (or lower) than some equilibrium level, high-cost producers will enter (or exit) the market, which pushes prices downward (or upward). Alternatively, the mean reversion question may be approached by analyzing the relationship between futures prices at different maturities and the spot price. Indeed, the studies cited above which support mean reversion rely on models which reflect the information contained in the futures price series about the spot price series, or in other words, focus on the convenience yield.

The convenience yield can be defined as the flow of goods and services that accrues to the owner of a commodity (a physical inventory) but not to the owner of a futures contract (a contract for future delivery).[5] The random-walk hypothesis is consistent with a constant convenience yield. In contrast, mean reversion and the positive correlation between spot price and convenience yield changes are consistent with the theory of storage with random shocks: when inventories decrease (or increase), the spot price will increase (or decrease) and the convenience yield will also increase (or decrease), however futures prices will not increase (or decrease) as much as the spot prices. Studies such as Schwartz

[5] For a cogent exposition in the context of oil price, see Pindyck (2001).

(1997), Schwartz and Smith (2000) or Pindyck (1999, 2001) refute the hypothesis of a constant convenience yield and their results suggest mean reversion to a long-run equilibrium that itself can change randomly over time.[6]

In response to these findings, this paper analyzes a mean-reverting class of models from Schwartz and Smith (2000) and Schwartz (1997), relative to various random-walk based alternatives, with focus on forecast performance.[7] The mean reverting structural form considered presumes a stochastic convenience yield, derives from the joint behaviour of spot and different future prices, and allows one to disentangle the persistent or long-run from the transitory or short-run component of oil price.

The non-mean reverting class of models considered integrates the major features documented in the literature that were found to be relevant for our purpose. In particular, as one of the most prominent stylized facts of oil price returns is that their volatility changes over-time, we consider various GARCH specifications. Given the importance of accounting for structural discontinuities, we also consider processes with random jumps.[8] Such jumps can be seen as an integral part of the empirical price process leading to relatively rare adjustments that can be distinguished from frequent and relatively "small" ordinary price fluctuations. These adjustments can result from an accident that shuts down the production of a large oil field for some time, or the unexpected decision of a producer to boost or cut down its production, or more generally from the arrival of unexpected information.[9] On breaks versus ARCH in oil prices, the reader may refer to Wilson et al.

[6] Note that Pindyck (1999) models long-run prices, whereas Schwartz (1997) or Schwartz and Smith (2000) propose models that incorporate both short- and long-run considerations.

[7] Crespo-Cuaresma et al. (2009) use also forecasting performance as the model selection criterion in their analysis of quarterly oil prices. They assume that the oil price trend has a unit root. See also Alquist and Kilian (2009).

[8] The relevant literature is vast; see for instance Merton (1976), Ball and Torous (1985), Jorion (1988), Bates (2000), Bakshi, Cao and Chen (2000), Das (2002), and Chernov, Gallant, Ghysels and Tauchen (2003).

[9] For a technical discussion on the relationship between jumps, breaks and GARCH processes, see

(1996).[10] In this paper, we aim to account for non-constancies in the mean and conditional variance via jump-GARCH processes as an alternative to dating and fitting breaks in addressing the question of mean reversion.

Our empirical analysis focuses on futures prices, ranging from 1986 to 2007. We compare the various model classes based on the mean-square forecast errors for daily, weekly and monthly frequencies, and for various forecast horizons, where we consider short-term and long-term forecasting. We use one-step-ahead out-of-sample forecasts, where parameter estimates are updated at every step of the procedure. Forecasting in real time has various practical advantages particularly given our focus on time-varying parameter models; practically, the unit-root model is given a fair chance since the drift parameter estimates adjust to additional observations. In addition, in models with jumps, where analytical formulae are not readily available for obtaining conditional expected forecast errors, we devise a simple simulation-based procedure to approximate these errors.[11] Our results support the mean-reverting model over all forecast horizons considered; indeed, observed forecast errors with the specification from Schwartz and Smith (2000) dominate, by far, the forecast errors of all non-mean reverting models considered. Our results underscore the importance of considering varying convenience yield in the analysis of oil price dynamics.

This paper is organized as follows. In section 2, we present the various models under consideration. Data and forecasting results are discussed in section 3. We conclude in section 4.

---

Drost, Nijman and Werker (1998) and Aït-Sahalia (2004).

[10] These authors study whether ARCH effects are wrongly included because the oil series present breaks-in-volatility. They consider daily 1-month futures from January 1984 to December 1992. Using a pre-test approach to identify multiple breaks, they find significant but less persistent ARCH effects in the presence of breaks.

[11] See Khalaf, Saphores and Bilodeau (2003) and Bernard, Khalaf, Kichian and McMahon (2008).

## 2 The Competing Models

Three alternative econometric specifications are considered which cover the recent and popular models applied to commodity prices: (i) random-walk models with GARCH effects, and with normal or Student-$t$ innovations, (ii) models with Poisson jumps and GARCH effects, and with normal or Student-$t$ innovations, and (iii) mean-reverting models that allow for uncertainty in the equilibrium to which prices revert. In what follows, specifications (i) and (ii) will be referred to as the random-walk-based class of models, and specification (iii) will be referred to as the mean-reverting class. For further reference, let

$$y_t = ln(Y_t) - ln(Y_{t-1}) \tag{2.1}$$

where $Y_t$ is the nominal price level at time $t$, $t = 1, ...T$.

GARCH models are among the most popular volatility models in practice, because they are capable of describing the well known volatility clustering feature as well as observed heavy tails, and are computationally and analytically tractable. The models that we consider in this class include the GARCH-in-mean model that is motivated by theory on storage[12], the exponential GARCH model that allows for asymmetric effects, and a standard GARCH model with Student-$t$ shocks that account for conditional heavy-tailed fundamentals.

Formally, we first consider the GARCH-in-mean model denoted GARCH-M(1,1):

$$\begin{aligned} y_t &= v_t + \beta h_t, & (2.2)\\ v_t &= \mu + \sqrt{h_t} z_t, & (2.3)\\ h_t &= \alpha_0 + \alpha_1 (y_{t-1} - \mu)^2 + \phi h_{t-1} & (2.4) \end{aligned}$$

where the fundamental shocks $z_t$, $t = 1, ...T$ are independently and identically distributed (i.i.d.) as standard normal or Gaussian variable

$$z_t \overset{i.i.d.}{\sim} N(0, 1). \tag{2.5}$$

[12] This was studied in particular by Beck (2001) in the context of commodity prices.

This model directly captures the relationship between level of price changes and volatility, via the $\phi$ coefficient in (2.4) and it is designed to allow for volatility clustering.

By setting $\beta = 0$, the model defined by (2.2) - (2.5) nests the standard GARCH(1,1) specification which we also consider. In this case, we also adopt an alternative specification for the fundamental shocks $z_t$, $t = 1, ...T$, namely the i.i.d. Student-$t$ distribution with $\tau$ degrees of freedom denoted as $t(\tau)$

$$z_t \overset{i.i.d.}{\sim} Student - t(\tau) \tag{2.6}$$

where $\tau$ is unknown. It is worth recalling that even under (2.5), the implied unconditional distribution of $v_t$ is non-normal, and in particular, the unconditional kurtosis exceeds 3 that is the Gaussian value. Hence even with normal shocks, GARCH processes would capture fat-tails. However, in applications of such models to financial high frequency data, it was observed that the unconditional kurtosis compatible with Gaussian-based GARCH models still understates the observed kurtosis in the data. Processes of the (2.6) form have thus been proposed to possibly address this problem.

For completion, we also consider the unit root (with drift) case, which corresponds to the above models setting the GARCH parameters to zero.

We next analyze the asymmetric specification denoted EGARCH(1,1) consisting of (2.2) - (2.3) with $\beta = 0$ and with

$$ln(h_t) = \alpha_0 + \phi ln(h_{t-1}) + \gamma \left( \frac{y_{t-1} - \mu}{\sqrt{h_{t-1}}} \right) + \eta \left[ \frac{|y_{t-1} - \mu|}{\sqrt{h_{t-1}}} - \frac{\sqrt{2}}{\sqrt{\pi}} \right]. \tag{2.7}$$

With this configuration, the sign of shocks is relevant. This property is particularly interesting for oil price, since negative shocks or news may conceivably affect volatility quite differently than positive ones.

Basically, the above models capture time-varying volatility as a function of the magnitude and/or the sign of lagged fundamental shocks. Except for the GARCH-M(1,1) model, these specifications restrict focus to non-constancies in the volatility, and not in the mean. There are alternative GARCH frameworks to model non-constancies in both mean and variance, for example by allowing for random jumps. From this class of models, we consider the following specification with Poisson jumps:

$$y_t = v_t + \sum_{i=1}^{n_t} lnP_{it}, \tag{2.8}$$

where $v_t$ is as defined in (2.3)-(2.4), $n_t$ is the number of jumps that occur between $t-1$ and $t$, and $P_{it}$ $(i = 1, ..., n_t)$ is the size of the $i$th jump over this time interval. Jumps follow a Poisson process with arrival rate $\lambda$, in other words, a jump occurs on average, every $1/\lambda$ periods, and the jump sizes $P_{it}$ are i.i.d. according to a lognormal distribution with mean $\theta$ and variance $\delta^2$. This definition implies that $n_t$ is an integer random variable and (2.8) also nests if $n_t = 0$, $t = 1, ...T$, or if $\lambda = 0$ the Gaussian and/or Student-$t$ GARCH(1,1) specification. We also consider a conditionally normal ARCH(1) with jumps, given its preponderance in the literature. This specification corresponds to (2.8) where $v_t$ is as defined in (2.3)-(2.4) and where the $\phi$ parameter in (2.4) set to zero.

The above random-walk based models and their extensions with jumps seem, however, more adapted to financial than to commodity markets. Indeed, while changes in stock prices are arguably unpredictable in an efficient market and lend support to the random-walk hypothesis, we can expect commodity prices to revert to some long-run trend. Intuitively, periods of repeated price increases should be followed by price decreases because price increases will induce new supplies and substitutes will become more attractive. Since these adjustments usually take time, the price of a commodity may overshoot its long-term marginal cost before eventually reverting to it. The converse is true for price decreases. A class of models that may be traced back to Gibson and Schwartz (1990) and Schwartz (1997) satisfies this requirement. These models revolve around the convenience yield concept. Here is a summary of the theoretical underpinning of such models.[13]

[13] Our summary borrows heavily from Pindyck (2001).

Mean-reversion stems from the crucial role that inventories play in the case of commodities. Numerous intuitive arguments may be envisaged to explain holding inventories. These include hedging in view of demand shifts against adjustment, marketing, scheduling or delivery costs, with obvious implications on price volatility. Indeed, in response to shifts in demand for commodities, producers make joint decisions on production and inventory levels, accounting for: (i) a spot price for the commodity, and (ii) a storage price determined from the so-called "marginal convenience yield". The latter is the flow of benefits that accrue to inventory holders from holding a marginal unit of inventory. In equilibrium it is equal to the cost of storage plus the interest loss on the purchase of the commodity (spot price) and minus the difference between the spot price and futures price. Two markets thus interact in the case of commodities, so equilibria in both markets are relevant, and interactions between these markets may be captured by analyzing the relationship between spot and futures markets. The following scenario may illustrate this dual market characteristic: in response to an unexpected temporary demand shock, spot prices are expected to increase and inventory holding to decrease while the futures prices stay constant. The price of inventory holding *i.e.* the convenience yield goes up; however it falls back as the situation returns to normal. Changes in spot price volatility may also trigger similar effects. Inventory adjustments are accompanied with adjustments in both spot and convenience yields.

Tractable time series models are available to capture such effects, and are typically characterized by mean reversion, but a time-varying convenience yield implies that the mean to which price reverts is, itself, time-varying. We consider the two-factor model of Schwartz and Smith (2000) that formally allows for a time-varying long-run mean and that integrates both short- and long-run movements by construction. More precisely, the long-run equilibrium component follows a Brownian motion, whereas the short-run deviations follow an Ornstein-Uhlenbeck process that reverts towards zero. The model can be written in continuous time for the log level of the spot price as:

$$\begin{aligned} \ln(Y_t) &= \chi_t + \xi_t, \\ d\chi_t &= -\kappa\chi_t dt + \sigma_\chi dz_\chi, \\ d\xi_t &= \mu_\xi dt + \sigma_\xi dz_\xi, \\ dz_\chi dz_\xi &= \rho_{\chi\xi} dt, \end{aligned} \tag{2.9}$$

where $\xi_t$ is the log equilibrium price of oil at time $t$, $\chi_t$ is the deviation of the log price at time $t$ with respect to the equilibrium price, and $dz_\chi$ and $dz_\xi$ are correlated increments of Brownian motions. The mean-reversion coefficient $\kappa$ represents the rate of speed at which the price reverts to its equilibrium *i.e.*, the rate at which short-run deviations disappear, $\mu_\xi$ is the growth rate of the log of equilibrium price, and $\sigma_\chi$ and $\sigma_\xi$ are the short-run and equilibrium volatilities of the process, respectively. For estimation purposes, (2.9) can be discretized as follows:

$$\ln(Y_t) = \chi_t + \xi_t, \tag{2.10}$$

$$\chi_t = e^{-\kappa}\chi_{t-1} + \epsilon_t^\chi, \tag{2.11}$$

$$\xi_t = \mu_\xi + \xi_{t-1} + \epsilon_t^\xi. \tag{2.12}$$

Such a decomposition allows one to disentangle the persistent from the transitory components of oil price.

To explain how a stochastic convenience yield intervenes in this model, Schwartz and Smith (2000) relate (2.9) to the following model from Schwartz (1997):

$$\begin{aligned} dX_t &= (\mu - \delta_t - \frac{1}{2}\sigma_1^2)dt + \sigma_1 dz_1, \\ d\delta_t &= \kappa(\alpha - \delta_t)dt + \sigma_2 dz_2, \\ dz_1 dz_2 &= \rho dt, \end{aligned} \tag{2.13}$$

where $X_t = \ln(Y_t)$ (in our notation, $X_t$ gives the log of the current spot price), $dz_1$ and $dz_2$ are correlated increments of Brownian motions, and $\delta_t$ is the convenience yield, which intervenes as a reduction in the drift term of (2.13). Formally, Schwartz and Smith (2000) show that processes (2.9) and (2.13) are equivalent, in the sense that factors of (2.9) can be written as a linear combination of the factors in (2.13). In particular, $\chi_t = \frac{1}{\kappa}(\delta_t - \alpha)$; note that $\kappa$ gives the short-term mean-reversion rate in both versions of the model, which justifies the overlap in notation.

Given the above specification, closed form solutions for the prices of futures can be obtained using standard valuation methods. For this purpose, Schwartz and Smith (2000) derive the corresponding risk-neutral process as:

$$\begin{aligned} d\chi_t &= (-\kappa\chi_t - \lambda_\chi)\, dt + \sigma_\chi dz_\chi^*, \\ d\xi_t &= \left(\mu_\xi - \lambda_\xi\right) dt + \sigma_\xi dz_\xi^*, \\ dz_\chi^* dz_\xi^* &= \rho_{\chi\xi} dt \end{aligned}$$

where $\lambda_\chi$ and $\lambda_\xi$ are constant reductions in the drifts of each process, and again, $dz_\chi^*$ and $dz_\xi^*$ are correlated increments of Brownian motions. Risk-adjustment now implies mean-reversion to -$\kappa/\lambda_\chi$ (rather than zero) for the short-run process. In the latter valuation framework, assuming that future prices are given by the expected future spot price leads to the following specification for future prices:

$$\ln(F_{n,t}) = e^{-kn}\chi_t + \xi_t + A(n), \tag{2.14}$$

where $F_{n,t}$ represents the market price, at time $t$, for a futures contract with time $n$ until maturity and

$$\begin{aligned} \xi_t &= \mu_\xi + \xi_{t-1} + \epsilon_t^\xi, \\ \chi_t &= e^{-\kappa}\chi_{t-1} + \epsilon_t^\chi, \\ A(n) &= \left(\mu_\xi - \lambda_\xi\right) n - \left(1 - e^{-\kappa n}\right) \frac{\lambda_\chi}{\kappa} \\ &\quad + \frac{1}{2}\left(\left(1 - e^{-2\kappa n}\right) \frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 n + 2\left(1 - e^{-\kappa n}\right) \frac{\rho_{\chi\xi}\sigma_\chi\sigma_\xi}{\kappa}\right). \end{aligned}$$

The system can then be written in a state-space form, and the latter is amenable to estimation via the Kalman filter.[14]

To conclude, two issues are worth noting. First, the GARCH models we consider are formulated on returns and explicitly assume that price is non-stationary. Schwartz and Smith's model treats price as the sum of two processes, one of them is a Brownian motion and non-stationary by construction. This model thus differs from standard random walks mainly via the Ornstein-Uhlenbeck process for the short term deviation. Secondly, volatility clustering is overlooked by Schwartz and Smith's model; nevertheless time varying coefficients models provide credible specifications to fit fat tails even when time varying volatility is not explicitly built-in.

# 3 Empirical Analysis

We use daily crude oil prices obtained from the *U.S. Department of Energy, Energy Information Administration*, for 1, 2, 3 and 4 month futures, from January 2, 1986 to January 9, 2007. From these, we construct weekly and monthly prices, the former using Wednesday values and the latter using the price on the Wednesday that is closest to the 15th day of that month. For the few cases where the Wednesday value is not available, the Tuesday value closest to the 15th day of that month is used. The models described in the previous section are estimated by numerical maximization of the likelihood function. In the case of the mean-reverting model, the state-space form associated with model (2.14) is estimated *jointly* for the four maturities considered using the Kalman filter.[15]

We rely on a standard summary measure to describe forecast performance, for daily and weekly frequencies, and for various forecast horizons. Specifically we compare the

[14] For a detailed treatment of this method, see Kim and Nelson (1999).

[15] For parsimony purposes and following Schwartz and Smith (2000), a diagonal matrix is assumed for the measurement errors in prices.

models under consideration, using the mean-square prediction errors (MSPE) based on one-step-ahead forecasts for each different model and frequency.[16] Forecasts and corresponding prediction errors are calculated in real time.

We proceed as follows. Given a sample of size $T+K$, we first set apart $K$ observations at the end of the sample, that correspond to the forecast horizon considered. The model is then estimated on the remaining sample (i.e., until $T$); the dependent variable's value is forecast for period $T+1$ and denoted $ln(\hat{Y}_{T+1|T})$. The $T+1$ forecast error resulting from the comparison of $ln(\hat{Y}_{T+1|T})$ and $ln(Y_{T+1})$ is computed. Next, the $T+1$ observed value of the dependent variable is added to our sample, and the model is re-estimated. The $T+2$ observation is then forecast and denoted $ln(\hat{Y}_{T+2|T+1})$, the $T+2$ forecast error is computed, and so on, until all $K$ observations are covered. The MSPE criterion is then defined as:

$$MSPE = \frac{1}{K}\sum_{k=1}^{K}\left[ln(\hat{Y}_{T+k|T+k-1}) - ln(Y_{T+k})\right]^2.$$

For the models that include jump features, a closed-form analytical solution for the forecast error is unavailable. We propose, as in Bernard et al. (2008) and Khalaf et al. (2003), a simulation-based approximation described as follows.

For a given model with jumps, say (2.8), we first estimate the model parameters over the sample of size $T$; denote the latter estimates $\hat{\alpha}_0$, $\hat{\alpha}_1$, $\hat{\mu}$, $\hat{\lambda}$, $\hat{\theta}$, $\hat{\delta}^2$ and $\hat{\tau}$ when relevant. Then, drawing from a normal or $t(\hat{\tau})$-distribution for the residuals, a Poisson distribution with estimated mean $\hat{\lambda}$ for the arrivals of the jumps, and a normal distribution with mean $\hat{\theta}$ and variance $\hat{\delta}^2$ for the amplitude of each jump, we generate 1,000 simulated values of the dependent variable $\widetilde{Y}_{T+1}$. The forecast value of $Y_{T+1}$ is then taken to be the average value of these 1,000 $\widetilde{Y}_{T+1}$, and the $T+1$ forecast error is computed. At this point, the observed value of the dependent variable, $Y_{T+1}$, is added to the sample, the model is

[16] The mean absolute forecast errors yield almost identical model ranking and the results are not reported here.

re-estimated, and the entire simulation process is repeated. Thus, $\widetilde{Y}_{T+2}$ is obtained, as well as the forecast error for $T + 2$. The above steps are repeated until $T + K$ forecast errors are obtained, which are then used to construct the MSPE.

For each frequency, we conduct out-of-sample one-step-ahead dynamic forecasts for three forecast horizons: one, three, and five years. Forecasts in real time are particularly useful here, in order to give the GARCH models a fair chance relative to the time-varying mean case. Indeed, it is well known that, except for the GARCH-in-mean and the jump-GARCH case, typical GARCH forecasts will not adapt unless adaptive estimation and forecasting is considered, since the associated optimal forecast is equal to the conditional mean. We thus view the rolling estimation burden as a worthy exercise from our perspective. Results are reported in Tables 1 - 3.

As may be checked from these Tables, the mean-reverting model of Schwartz and Smith (2000), with stochastic convenience yield, emerges as the best model for all of the forecast horizons considered by a wide margin. Indeed, random-walk models with GARCH and with or without jumps or asymmetries are largely inferior to the mean-reverting model. As a robustness analysis, and to check whether such an outcome is totally driven by soaring markets in the last couple of years, we repeated the above exercise, using data from January 2, 1986 to January 9, 2005 only. Results available upon request are qualitatively similar for all future prices, all data frequencies and all forecasting horizons considered.

Our results may be viewed as a forecast-based evidence in favour of mean reversion in oil prices. While lending support to a different alternative hypothesis that is motivated by our focus on high frequency models and data, our findings are in line with rejections of the unit root in the presence of breaks as in Lee et al. (2006) and Postali and Picchetti (2006). This observation must be qualified since both the mean reversion and break concepts compatible with our findings differ, as formally specified above, from the standard unit-root test setting.

First, we focus on price reversion to a continuously-evolving rather than to an abruptlybreaking mean. The random walk model we consider also allows for discontinuities formulated as Poisson jumps that are inherent to the data generating process, rather than as abrupt breaks or outliers. Such specifications are founded on theory, as discussed by e.g. Schwartz (1997), Schwartz and Smith (2000) or Pindyck (2001), and are gaining popularity in economics and finance in general, when fitting medium-to-long time series. Indeed, while the consequences of ignoring breaks in means and variances are well understood, over-fitting problems are difficult to avoid when searching for potentially many change-points in practice. Such considerations have recently motivated time-varying-parameter specifications - as alternatives to break-point ones - that allow for continuous and possibly unpredictable shifts in slopes and trends through stochastic coefficients.[17]

Secondly, we integrate the convenience yield into the structure of the mean reverting model considered. The difference between spot prices and futures prices is a major component of the marginal convenience yield. We find that useful inference results from fitting the information contained in different futures price series about spot price, through the evolution of this yield.

Thirdly, relative to the typical unit-root test, both the random-walk and the alternative mean-reverting models considered are non-linear, given our focus on high frequency data. Indeed, in such contexts, non-linear models may better capture asymmetries and relatively extreme or atypical events, for it is possible to formulate non-linear processes which naturally imply such effects. It is well known that neglected atypical events can blur inference with linear models, leading to unreliable decisions on major issues, including mean-reversion, and to particularly inaccurate forecasts. Correcting linear models for outliers or breaks could address associated inference problems, on an ex-post basis. Yet for forecasting purposes, non-linear models designed to accommodate non-normal features of the data on an ex-ante basis, hold potentially more promise. The models we consider

[17] The popular ARCH and GARCH specifications may also be viewed as cases where volatility parameters are modelled as a stochastic process.

also control for non-constancies in both the mean and the conditional variance.

Arguments in favour of random walks with or without GARCH effects, or in other words, a focus on returns rather than on levels, may persist in the literature on oil prices for tractability reasons.[18] Time series models and even structural models for that matter are approximations, at best, and simple and parsimonious models are often preferred, at least empirically. So in setting up a framework to describe highly persistent data, reliance on unit roots is conceivable even when economic theory suggests mean reversion. Indeed, when the speed of mean reversion is particularly slow, or when considering low frequency data and/or when sample sizes are limited, random walks, as empirical time paths, may simply fit the facts.

Our findings strongly suggest that imposing the random walk for oil prices has pronounced costs for out-of-sample forecasting. It is worth recalling that Schwartz and Smith's model formulates the spot price as the sum of two processes, one of them is Brownian motion. The key difference relative to a standard random walk arises from the use of an Ornstein-Uhlenbeck process for the short term deviation.

Our conclusions must be interpreted in the context of the data frequency under consideration. The fact that monthly data based forecasts corroborate our findings with daily and weekly data, is noteworthy. Yet we do not claim generality; the mean reverting model we considered is - at least conceptually - not designed to fit the long run evolution of yearly oil prices. Sample size concern is one clearly relevant issue, yet more importantly, the underlying market movements it reflects, as described above, are rather shorter run ones. In this regard, it is worth citing the study from Pindyck (1999) which is more concerned with long run forecasts of oil prices, and thus relies on low frequency data. Specifically, Pindyck examines an annual oil price series, in real terms, defined as averages of producers' prices in the US over the period 1870-1996. He proposes an autore-

[18] For instance, see e.g. Postali and Picchetti (2006) or Sadorsky (2006).

gressive model (in logs) with time-varying intercept and trend coefficients. The forecasts from this model based on several data sub-samples are compared to (i) actual data when it is possible, and (ii) forecasts from a model reverting to a fixed linear trend. Pindyck's main results center, among others, on the following. First the forecasts (through the 21st century) seem not to be affected by the inclusion or exclusion of the 1974-1981 period, although these years have witnessed large price variations. Second, the forecasts using data till 1981 are nonetheless closer to the real data than the forecasts of the fixed-mean reverting model. This work relates to our results in the sense that mean reversion to a stochastic trend emerges as a model of choice.

Finally, on comparing forecasts based on the non-mean reverting models considered *i.e.* with reference to all results reported in Tables 1 - 3 except those pertaining to Schwartz & Smith's model, two observations are worth noting. First, conditionally non-normal models or models with discontinuous jumps produce marginally better forecasts to some extent, with daily data. Secondly, the GARCH-M(1,1) model performs generally well within this class of models. This result agrees with Beck (2001) who argues that for storable commodities, random walk models allowing for conditional heteroskedasticity (specifically ARCH-in-mean effects) are compatible with rational expectations and risk aversion. Observe that the analysis in Beck (2001) relies on a long-run theoretical model, suitable for relatively high frequency data. Finally, improvements relative to the no-GARCH unit root basis case are more marked the longer the maturity and the longer the forecast horizon.

Taken collectively, our comparative results underscore the merit of the model in Schwartz and Smith (2000) and motivate further improvements to it. As Schwartz and Smith (2000) outline, these include improving the long-run equation (incorporating, for example, formulations as in Pindyck 1999), or the short-run one, by adding discrete jumps.

More generally, we view our results as lending support to models for oil price which attempt to fit the inter-related equilibria in the spot and storage markets. In this context, our evidence on mean reversion has broad economic significance, because the mean-reverting structural form we consider captures both equilibria through the dynamics of the convenience yield.

# 4 Conclusion

In this paper, we compare several stochastic models for oil prices, focusing on forecast performance. The models differ regarding the assumptions related to mean reversion and structural discontinuities (time-varying first and second moments). We compare forecasts in real time, for 1, 3 and 5 year horizons. For the jump-based models, we rely on numerical methods to approximate forecast errors.

Our analysis with futures price data ranging from 1986 to 2007 strongly suggests that imposing the random walk for oil prices has pronounced costs for out-of-sample forecasting. Indeed, the model which produces the best forecasts for all horizons and for all data frequencies considered assumes price reversion to a long-run equilibrium that itself can change randomly over time. Random-walk based alternative models, corrected for jumps, conditional non-normality and time-varying volatility, produce markedly inferior forecasts. This evidence suggests mean reversion in oil markets, with prices reverting to a continuously evolving mean where the underlying evolution is linked to a continuously evolving convenience yield.

The mean reverting model we consider decomposes price into the sum of a Brownian motion for the long-run evolution, and an Ornstein-Uhlenbeck process for the short term deviations. The key difference relative to a standard random walk stems from the latter component. Such decomposition is based on the joint evolution of spot and future oil prices and on a built-in specification for the convenience yield, which allows to disentangle permanent from transitory price shocks. In this context, favorable inference regarding mean reversion has broad economic relevance and underscores the crucial role that inventories play in the case of commodities.

Reaching beyond the mean reversion issue, our results illustrate the importance of relying on a time series models for oil which allows, to some extent, and mainly via formulating a process for the dynamics of the convenience yield, to describe the interrelated equilibria in the spot and storage market. These models though popular in the finance literature have attracted much less interest in economics. Our findings strongly suggest that adopting such models yields huge pays off when forecasting oil prices.

Table 1: Oil Price Forecast Errors, Daily Data

| Daily Frequency | 1 month Futures | | | 3 month Futures | | |
|---|---|---|---|---|---|---|
| Forecast Horizon | 1 year | 3 years | 5 years | 1 year | 3 years | 5 years |
| Unit root | .0115 | .3558 | .8256 | .0105 | .3736 | .7323 |
| GARCH(1,1) | .0113 | .3324 | .8706 | .0099 | .3636 | .7377 |
| GARCH-M(1,1) | .0118 | .4259 | 1.711 | .0099 | .3912 | *.6306* |
| EGARCH(1,1) | .0104 | .3429 | *.7835* | .0102 | .3695 | .7672 |
| GARCH-T(1,1) | *.0101* | *.3318* | .7984 | *.0093* | .3524 | .7175 |
| ARCH(1) with jumps | .0106 | .3796 | .8583 | .0104 | *.3265* | .8884 |
| GARCH(1,1) with jumps | .0150 | .3822 | 1.2272 | .0124 | .4526 | 1.0065 |
| GARCH-$t$(1,1) with jumps | .0103 | .3444 | .8443 | .0096 | .3672 | .7536 |
| Schwartz & Smith | **.0006** | **.0016** | **.0003** | **.00001** | **.0014** | **.0004** |
| Daily Frequency | 2 months Futures | | | 4 months Futures | | |
| Forecast Horizon | 1 year | 3 years | 5 years | 1 year | 3 years | 5 years |
| Unit Root | .0102 | .3469 | .7456 | .0113 | .4203 | .7539 |
| GARCH(1,1) | .0097 | .3294 | .7232 | .0102 | .4083 | .7601 |
| GARCH-M(1,1) | *.0091* | *.2875* | *.7196* | *.0084* | *.3909* | .9071 |
| EGARCH(1,1) | .0098 | .3362 | .8413 | .0105 | .4200 | .8004 |
| GARCH-T(1,1) | .0095 | .3297 | .7249 | .0096 | .3998 | *.7504* |
| ARCH(1) with jumps | .0101 | .4024 | .9002 | .0109 | .4697 | .8995 |
| GARCH(1,1) with jumps | .0114 | .3971 | 1.0103 | .0126 | .4635 | .9787 |
| GARCH-$t$(1,1) with jumps | .0095 | .3379 | .7565 | .0101 | .1483 | .7901 |
| Schwartz & Smith | **.0001** | **.0015** | **.0005** | **.0001** | **.0250** | **.0001** |

Note: Numbers shown are MSPE; numbers in bold refer to the model which minimizes the corresponding forecast error measure; numbers in italics refer to the non-mean reverting model which minimizes the corresponding forecast error measure.

Table 2: Oil Price Forecast Errors, Weekly Data

| Weekly Frequency | 1 month Futures | | | 3 month Futures | | |
|---|---|---|---|---|---|---|
| Forecast Horizon | 1 year | 3 years | 5 years | 1 year | 3 years | 5 years |
| Unit Root | .0160 | .2949 | .7028 | .0184 | .3340 | .6582 |
| GARCH(1,1) | .0174 | .3273 | .6670 | .0191 | .3581 | .6224 |
| GARCH-M(1,1) | *.0071* | *.1745* | *.4237* | .0182 | *.1103* | *.3607* |
| EGARCH(1,1) | .0176 | .3346 | .6919 | .0191 | .3596 | .6226 |
| GARCH-T(1,1) | .5338 | 5.1347 | 10.7391 | .6621 | 7.0137 | 11.7519 |
| ARCH(1) with jumps | .0145 | .5491 | .7275 | *.0170* | .3871 | .6500 |
| GARCH(1,1) with jumps | .0239 | .4005 | .8603 | .0275 | .4719 | .8354 |
| GARCH-$t$(1,1) with jumps | .0174 | .3143 | .6453 | .0191 | .3497 | .6123 |
| Schwartz & Smith | **.0004** | **.0003** | **.0002** | **.0001** | **.0001** | **.0001** |
| **Weekly Frequency** | **2 months Futures** | | | **4 months Futures** | | |
| Forecast Horizon | 1 year | 3 years | 5 years | 1 year | 3 years | 5 years |
| Unit Root | .0174 | .3110 | .6659 | .0190 | .3572 | .6492 |
| GARCH(1,1) | .0182 | .3341 | .6260 | .0196 | .3827 | .6183 |
| GARCH-M(1,1) | *.0069* | *.0057* | *.0167* | *.0158* | *.0996* | *.2980* |
| EGARCH(1,1) | .0183 | .3354 | .6300 | .0195 | .3809 | .6118 |
| GARCH-T(1,1) | .6885 | 7.1246 | 13.1479 | .6181 | 6.7638 | 10.4958 |
| ARCH(1) with jumps | .0172 | .1834 | .6671 | .0161 | .3620 | .6620 |
| GARCH(1,1) with jumps | .0281 | .4464 | .8638 | .0270 | .5022 | .8310 |
| GARCH-$t$(1,1) with jumps | .0183 | .3255 | .6154 | .0197 | .3739 | .6081 |
| Schwartz & Smith | **.0001** | **.0001** | **.0001** | **.0001** | **.0001** | **.0001** |

Note: Numbers shown are MSPE; numbers in bold refer to the model which minimizes the corresponding forecast error measure; numbers in italics refer to the non-mean reverting model which minimizes the corresponding forecast error measure.

Table 3: Oil Price Forecast Errors, Monthly Data

| Monthly Frequency | 1 month Futures | | | 3 month Futures | | |
|---|---|---|---|---|---|---|
| Forecast Horizon | 1 year | 3 years | 5 years | 1 year | 3 years | 5 years |
| | MSPE | MSPE | MSPE | MSPE | MSPE | MSPE |
| Unit Root | .0015 | .2479 | .8461 | .0110 | .2917 | .7769 |
| GARCH(1,1) | *.0014* | .2967 | .8527 | .0108 | .3319 | .8045 |
| GARCH-M(1,1) | .0098 | *.2324* | *.7151* | *.0087* | *.2503* | *.6098* |
| EGARCH(1,1) | .0114 | .3051 | .8800 | .0108 | .3318 | .8035 |
| GARCH-T(1,1) | .0115 | .3138 | .8721 | .0109 | .3298 | .8039 |
| ARCH(1) with jumps | .0123 | .2932 | .8334 | .0142 | .5220 | 1.1605 |
| GARCH(1,1) with jumps | .0129 | .3140 | .9817 | .0132 | .3885 | .9811 |
| GARCH-$t$(1,1) with jumps | .0121 | .2946 | .8427 | .0112 | .3269 | .7782 |
| Schwartz & Smith | **.0004** | **.0002** | **.0002** | **.0001** | **.0001** | **.0001** |
| Monthly Frequency | 2 months Futures | | | 4 months Futures | | |
| Forecast Horizon | 1 year | 3 years | 5 years | 1 year | 3 years | 5 years |
| | MSPE | MSPE | MSPE | MSPE | MSPE | MSPE |
| Unit Root | .0018 | .2513 | .8468 | .5363 | .3196 | .7431 |
| GARCH(1,1) | *.0016* | .3049 | .8727 | .5319 | .3459 | .7594 |
| GARCH-M(1,1) | .0100 | *.2373* | *.7019* | *.4961* | *.2515* | *.5483* |
| EGARCH(1,1) | .0117 | .3207 | .9069 | .5320 | .3453 | .7556 |
| GARCH-T(1,1) | .0117 | .3212 | .8774 | .5319 | .3445 | .7593 |
| ARCH(1) with jumps | .0145 | .4285 | 1.2146 | .5379 | .4761 | 1.1162 |
| GARCH(1,1) with jumps | .0145 | .3632 | 1.0565 | .5470 | .3949 | .9374 |
| GARCH-$t$(1,1) with jumps | .0124 | .3026 | .8516 | .5330 | .3391 | .7325 |
| Schwartz & Smith | **.0008** | **.0005** | **.0006** | **.0001** | **.0001** | **.0001** |

Note: Numbers shown are MSPE; numbers in bold refer to the model which minimizes the corresponding forecast error measure; numbers in italics refer to the non-mean reverting model which minimizes the corresponding forecast error measure.

# References

Abosedra, S. S. and Laopodis, N. T. (1996), 'Stochastic behavior of crude oil prices: A GARCH investigation', *Journal of Energy and Development* **21**(2), 283–91.

Ahrens, W. A. and Sharma, V. R. (1997), 'Trends in natural resource commodity prices: Deterministic or stochastic?', *Journal of Environmental Economics and Management* **33**, 59–74.

Aït-Sahalia, Y. (2004), 'Disentangling diffusion from jumps', *Journal of Financial Economics* **74**, 487–528.

Alquist, R. and Kilian, L. (2009), 'What do we learn from the price of crude oil futures?', *Journal of Applied Econometrics* **forthcoming**.

Bakshi, G., Cao, C. and Chen, Z. (2000), 'Pricing and hedging long-term options models', *Journal of Econometrics* **94**, 277–318.

Ball, C. and Torous, W. (1985), 'On jumps in common stock prices and their impact on call option pricing', *Journal of Finance* **40**, 155–73.

Bates, D. S. (2000), 'Post-87 crash fears in S&P 500 futures options', *Journal of Econometrics* **94**, 181–238.

Beck, S. (2001), 'Autoregressive conditional heteroscedasticity in commodity spot prices', *Journal of Applied Econometrics* **16**, 115–32.

Berck, P. and Roberts, M. (1996), 'Natural resource prices: Will they ever turn up', *Journal of Environmental Economics and Management* **9**, 122–137.

Bernard, J.-T., Khalaf, L., Kichian, M. and McMahon, S. (2008), 'Forecasting commodity prices: GARCH, jumps, and mean-reversion', *Journal of Forecasting* **27**, 279–291.

BP (2009), *BP Statistical Review of World Energy* .

Chernov, M., Gallant, A. R., Ghysels, E. and Tauchen, G. (2003), 'Alternative models for stock price dynamics', *Journal of Econometrics* **116**, 225–57.

Cortazar, G. and Naranjo, L. (2006), 'An N-factor gaussian model for oil futures prices', *The Journal of Future Markets* **26**, 243–268.

Cortazar, G. and Schwartz, E. S. (2003), 'Implementing a stochastic model for oil future prices', *Energy Economics* **25**, 215–238.

Crespo-Cuaresma, J., Jumah, A. and Karbuz, S. (2009), 'Modelling and forecasting oil prices: The role of asymmetric cycles', *The Energy Journal* **30**, 81–90.

Das, S. R. (2002), 'The surprise element: Jumps in interest rates', *Journal of Econometrics* **106**, 27–65.

Drost, F. C., Nijman, T. E. and Werker, B. J. M. (1998), 'Estimation and testing in models containing both jump and conditional heteroskedasticity', *Journal of Business and Economic Statistics* **16**, 237–43.

Gibson, R. and Schwartz, E. S. (1990), 'Stochastic convenience yield and the pricing of oil contingent claims', *Journal of Finance* **45**, 959–76.

Hamilton, J. D. (2009), 'Understanding crude oil prices', *The Energy Journal* **30**, 179Ű206.

Jorion, P. (1988), 'On jump processes in the foreign exchange and stock markets', *Review of Financial Studies* **1**, 427–45.

Khalaf, L., Saphores, J.-D. and Bilodeau, J.-F. (2003), 'Simulation-based exact jump tests in models with conditional heteroskedasticity', *Journal of Economic Dynamic and Control* **28**, 531–553.

Kim, C.-J. and Nelson, C. R. (1999), *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, The MIT Press, Cambridge, Massachusetts.

Lee, J., List, J. A. and Strazicich, M. C. (2006), 'Non-renewable resource prices: Deterministic or stochastic trends?', *Journal of Environmental Economics and Management* **51**, 354–370.

Merton, R. C. (1976), 'Option pricing when underlying stock returns are discontinuous', *Journal of Financial Economics* **3**, 125–44.

Moshiri, S. and Foroutan, F. (2006), 'Forecasting nonlinear crude oil future prices', *The Energy Journal* **27**, 81–85.

Pindyck, R. S. (1999), 'The long run evolution of energy prices', *The Energy Journal* **20**, 1–27.

Pindyck, R. S. (2001), 'The dynamics of commodity spot and futures markets: a primer', *The Energy Journal* **22**, 1–29.

Postali, F. A. S. and Picchetti, P. (2006), 'Geometric brownian motion and structural breaks in oil prices: A quantitative analysis', *Energy Economics* **28**, 506–522.

Regnier, E. (2007), 'Oil and energy price volatility', *Energy Economics* **29**, 405–427.

Sadorsky, P. (2006), 'Modeling and forecasting petroleum futures volatility', *Energy Economics* **28**, 467–488.

Schwartz, E. S. (1997), 'The stochastic behavior of commodity prices : Implications for valuation and hedging', *The Journal of Finance* **52**, 923–973.

Schwartz, E. S. and Smith, J. E. (2000), 'Short-term variations and long-term dynamics in commodity prices', *Management Science* **46**, 893–911.

Smith, J. L. (2009), ‘World oil: Market or mayhem’, *Journal of Economic Perspectives* **23**, 145Ű164.

Tabak, B. M. and Cajueiro, D. O. (2007), ‘Are the crude oil markets becoming weakly efficient over time? a test for time-varying long-range dependence in prices and volatility’, *Energy Economics* **29**, 28–36.

Wilson, B., Aggarwal, R. and Inclan, C. (1996), ‘Detecting volatility changes across the oil sector’, *The Journal of Future Markets* **16**(3), 313–330.

Wirl, F. (2008), ‘Why do oil prices jump (or fall)?’, *Energy Policy* **36**(3), 1029–1043.