Master thesis for the Master of Philosophy in Economics degree

Crude oil prices and petroleum inventories

Remedies for a broken oil price forecasting model

Dan Grimstad

November 2007

Department of Economics
University of Oslo

Preface

I would like to express my gratitude to my supervisor, Roger Bjørnstad, for his invaluable and most effective guidance. I also want to thank Alice, my dearest, for her immense patience and love. All errors are mine.

Dan Grimstad

Oslo, 12th November 2007

Contents

1 Introduction

The empirical relationship between crude oil prices and petroleum inventories has been exploited in a number of short-term oil price forecasting models, e.g. Amano (1987), Kaufmann (1995), Kaufmann, Dees et al. (2004), Zamani (2004) and Ye, Zyren et al. (2002, 2005a, 2006a, 2006c).

One popular class of oil price forecasting models is based on the perception that an unexpected inventory level indicates an imminent price change. The underlying assumption is that imbalances between crude oil supply and demand affect inventories before price.

The last years, these “relative inventory level” models have failed, in the sense that they have consistently under-predicted the oil price. Some relate the failure to an apparent inversion in the relationship between price and inventories, as illustrated by Figure 1 below. Ye, Zyren et al. (2006b:556) find that the “unusual positive relationship” may indicate that the oil market is in a state of transition. Merino and Ortiz (2005) find that futures market activity may explain parts of the forecast failure. Tchilinguirian (2006:11) observes that the price-inventory relationship has moved “out of line”, and suggests that the reason may be a shift in the storage supply curve.

This thesis asks if there is a break in the relative inventory level model that can be attributed to the simplistic representation of the price-inventory relationship, or if the break is rather in the oil price process, unrelated to the price-inventory relationship.

The main find is a slow shift in the oil price process between 2002 and 2007, towards a new regime with a higher mean, higher variance and lower persistence. The investigated relative inventory models fail because they assume that the oil price reverts to a constant mean. Furthermore, it is found that whilst the price-inventory relationship is not broken, the statistical properties of the linear relative inventory model may be considerably improved by a quadratic log-log respecification.

The findings have some interesting consequences. Firstly, some of the lost forecasting power may be restored by introducing a model component that represents a slow shift in the long-run oil price. Although this fix does not explain the causes of the shift, it absorbs some of the distortions from the shift, and enables estimation of the parameters by the use of simple regression methods. Secondly, the proposed quadratic log-log specification allows for a more proper statistical treatment of uncertainty, and is well-specified enough to be used for hypothesis testing and computation of confidence bands. Finally, the identified regime shift resembles the shifts that Videgaray-Caso (1998) found to come about every 11th year or so, before reverting back after another 4 or 5 years.

A simple approach is followed to answer the research questions: The point of departure is a dissection of the most parsimonious amongst the reportedly successful relative inventory models. An explicit formulation of the embedded economic model justifies a slightly modified economic model that is more aligned with theory, and presumably more robust in the face of data. The two economic models lead to two econometric models: The baseline model, basically equal to the dissected model, and the competing quadratic log-log model. Both econometric models are then estimated using ordinary least squares and evaluated on the original dataset (1992-2003) and a more complete dataset (1992-2007), in order to investigate if the simple respecification remedies any break. Next, a version of the proposed model that allows for a smooth shift in the intercept and autoregressive parameters is specified and estimated using maximum likelihood and nonlinear least squares on the complete dataset, primarily in order to investigate the hypothesis that there is a break in the oil price process. The software tools PcGive 11.04, JMulTi 4.21 and Stata 9.2 are used in the estimations.

The thesis is organised as follows: The literature review justifies the research questions and hypotheses, by first explaining the rationale behind using relative inventory level for forecasting, and then arguing that price is nonlinear in inventories and that price shifts are common. The next section derives the two theoretical models that the empirical analysis will build on. The section that follows outlines the econometric methodology used in the empirical study, and justifies some of the choices. The empirical methodology section reports the results from the empirical study, and concludes with a discussion of the findings, in the context of the research questions and hypotheses. The last section contains some concluding remarks.

Figure 1 Oil price and the relative inventory level

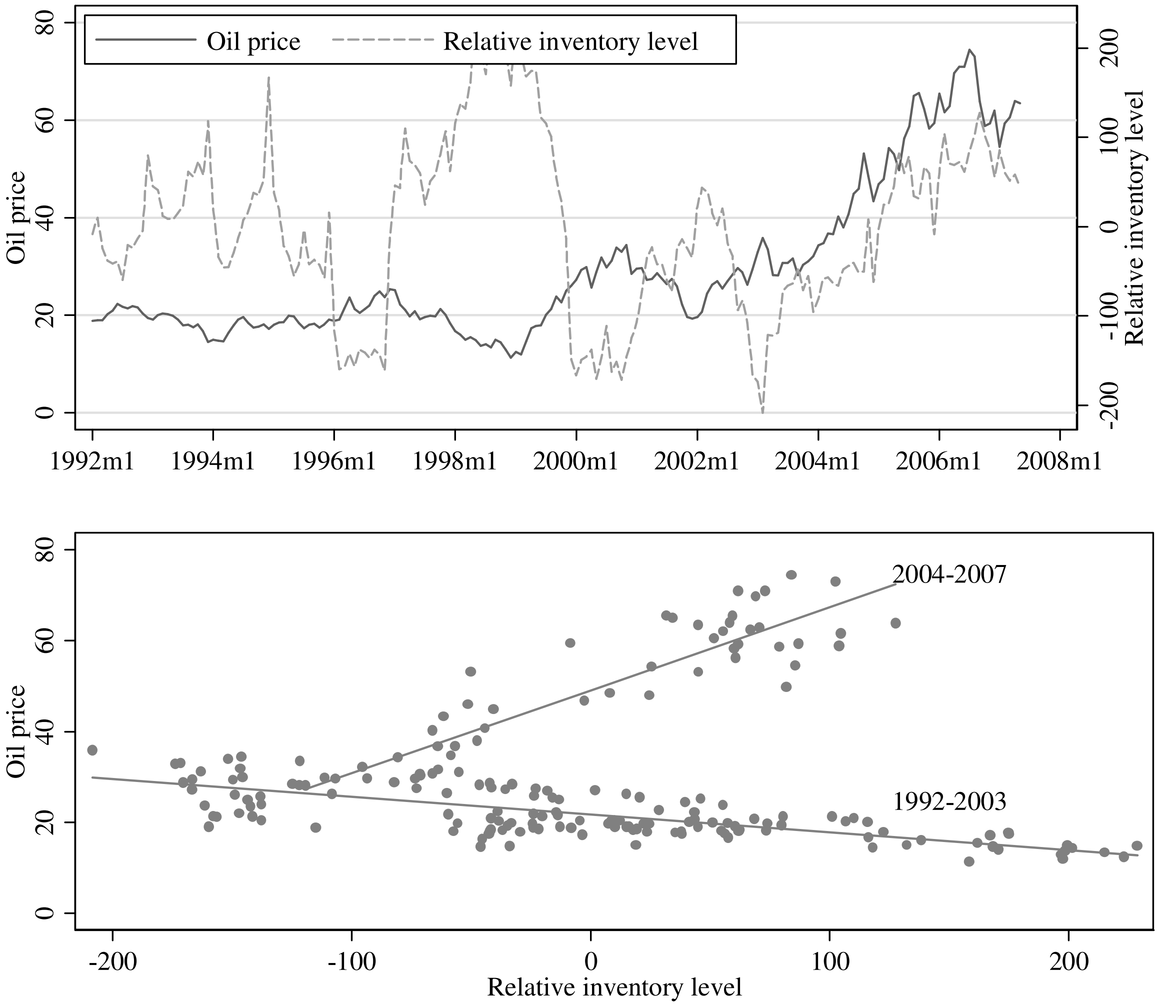

Note: Oil price is the nominal WTI spot price in USD per stock tank barrel. The relative inventory level is a measure of deviation of total OECD inventories away from “normal”, in million stock tank barrels, as defined in Ye, Zyren et al. (2002). The unlabeled x-axis in the upper graph shows time (year and month). The regression lines in the scatter plot are from linear regression of the labelled dataset segments.

2 Literature review

This section first describes a simple relative inventory level model and its theoretical content, and briefly reports the findings of Merino and Ortiz (2005) and Ye, Zyren et al. (2006b). Both suggest plausible explanations of the failure of the relative inventory level models. Then, the two issues of nonlinearity in the price-inventory relationship and shifts in the oil price process are explained and related to the research questions and hypotheses.

2.1 Forecasting oil price using linear relative inventory level

2.1.1 A simple relative inventory level model

Ye, Zyren et al. (2005) find that this simple relative inventory model outperforms a number other forecasting models:

$$OP_t = \alpha_0 + \alpha_1 OP_{t-1} + \sum_{i=0}^{3} \beta_i RIN_{t-i} + \sum_{j=0}^{5} \gamma_j D_j 911 + \zeta LAPR99 + \varepsilon_t$$

Here, OP_t is the oil price in month t and RIN_t is the relative inventory level in month t. The relative inventory level is defined as the difference between current inventories IN_t in month t and a normal inventory level IN_t^* for that month, i.e. $RIN_t = IN_t - IN_t^*$. The normal inventory level is the deseasonalised, linearly detrended inventory level. $D_j 911$, $j = 0,1,2,...,5$ are impulse dummy variables for the unstable period from October 2001 to March 2002, and $LAPR99$ is a level shift dummy representing an OPEC inventory policy shift.

2.1.2 The structural content of relative inventory level models

Ye, Zyren et al. (2002, 2006a) give an account of the rationale behind the relative inventory level models. Basically, these models consider an unexpected inventory change as an indicator of an imminent price change. The explanations seem to rely on three main arguments:

Firstly, shifts in supply or demand have an immediate impact on inventories, whilst the oil price takes some time to adjust. The rationale for this is that inventory holders form expectations concerning future price, demand or supply that influence current inventories; A refining company will build inventory if they expect a price increase or a supply shortage, and draw down inventories if they see slowing demand for refined products.

Secondly, the market price of crude oil is determined by commercial inventories, because these inventories are the marginal source of supply.

Thirdly, there is an expected inventory level that can be derived from previously realised inventory levels, taking into account seasonal influences and “long-term inventory trends”

(Ye, Zyren et al., 2005:491). This expected inventory level, denoted "normal" or "desired", is supposed to represent what the market actors plan for, and thus what they expect.

The fundamental idea is then: If current inventories for a month deviates from what it should be, according to inventory trends and seasonal swings, this is a signal that inventory holders have already adjusted to an anticipated change in price, supply or demand.

2.1.3 The failure of relative inventory level models, and some possible reasons

Relative inventory level models now fail, in the sense that they consistently under-predict the short-term oil price.

Merino and Ortiz (2005) observe that from mid-2003, there has been an increasing discrepancy between realised prices and price predictions from inventory level models. Their hypothesis is that the relative inventory level models should be able to explain the dramatic increase in the oil price, but that the models do not include enough information. Their approach is to extend a slightly modified relative inventory level model with a broad set of plausible explanatory variables: Two types of inventories, different futures prices, production and refining capacity, three indicators of futures market activity, bond market variables, foreign exchange market variables, and a commodity market index. They find that only futures market activity (proxied by non-commercials' long positions in futures markets) can explain a large part of the discrepancy between predicted prices and realised prices from mid-2003 to mid-2004.

Ye, Zyren et al. (2006b:556) briefly report that an unpublished study of theirs confirms that since 2004, both relative inventories models and excess production capacity models have under-predicted crude oil price. Their hypothesis is that because demand has outgrown capacity, excess production capacity is reduced, causing an upward price pressure, resulting in the "unusual positive relationship" between crude oil price and inventories. They "identify three crude oil market regimes since 1992: 1) periods of business as usual, 2) OPEC policy changes; and 3) reductions in excess capacity", and conclude that "either crucial market variables are missing from existing short-term price forecast models, or the market is currently in a transitional period".

2.2 Oil price is probably nonlinear in inventories

One challenge for the mostly linear relative inventory level models is that the oil price is probably nonlinear in inventories.

The theory of storage, represented by Kaldor (1939), Working (1949), Brennan (1958), and Telser (1958), postulates that the benefit of holding the next unit of inventory – the marginal convenience yield – falls as inventories build, i.e. an extra unit of inventory is more valuable if stocks are low. In the storage theory, the spread between the spot and futures prices of a commodity is explained by interest rates, storage costs and the convenience yield.

Since crude oil is a storable commodity, and storing oil above ground is expensive, the theory has implications that are relevant for oil price prediction: The (spot) price of oil contains a convenience yield component, and the marginal convenience yield on inventory declines at a decreasing rate as inventories build.

The marginal convenience yield is the basis of Pindycks (1993) present value model of rational commodity pricing, where the spot price of a storable commodity is given by the discounted value of expected future convenience yields. The discount factor consists of the risk-free interest rate plus a commodity-specific risk premium. A change in the price of a commodity is due to a change in the expected future convenience yields, a change which in turn may be caused by changes in current or expected demand or supply.

In the present value model, an inventory increase leads to a decrease in the convenience yield, which reduces the spot price, other things equal. The reduction in spot price is large if inventories increase from a low level, and the reduction in spot price is low if the inventories increase from an already high level.

The relative inventory level models build on a quite different explanation of the relationship between price and inventory, as evident from section 2.1 above. With the well-developed storage theory and the models derived from it in mind, it is reasonable to question the specification of the price-inventory relationship in the relative inventory level models: If the price-inventory correlation found by empirical relative inventory level studies in part reflects the convenience yield relationship above, the linear specification of most relative inventory level models may be wrong.

Ye, Zyren et al. (2002) recognise that the oil price responds differently at low inventories, and include a low-inventory variable in order to represent the asymmetric response in oil price of inventory changes at inventories below normal. This implies a price-inventory relationship that is kinked at zero relative inventory level, but otherwise linear. Thus, even if they dismiss storage theory as too demanding in terms of “expertise and specific data” (Ye, Zyren et al., 2002:325) , they have included a variable that at least in part may be able to capture some of the nonlinearity implied by the convenience yield.

Another exception is Ye, Zyren et al. (2006a), where two variables representing the high-inventory state and the low-inventory state of the market are found to improve the otherwise linear model. The state variables are intended to represent that oil price responds differently to inventory changes when inventories are particularly high or low. The high-state variable is defined as the square of the relative inventory level that exceeds one standard deviation, and correspondingly for the low-state variable. In this model, the price-inventory relationship is linear in a range around zero relative inventory level, but quadratic outside the range.

2.3 Shifts in the oil price process are common

Both theoretical and empirical studies show that oil price shifts happen. This represents a problem for the linear regression relative inventory level models, often mandating the use of dummy variables in order to ensure parameter constancy.

2.3.1 Oil price shifts in theory

Hotellings (1931) theory of exhaustible resources implies that if crude oil production is competitive, the price of oil less the total marginal extraction costs will increase at the rate of interest until it eventually reaches a choke price, where demand is zero.

Using a basic Hotelling model, Pindyck (1999:12ff) shows that changes in the drivers of long-run oil price - demand, extraction costs and reserves - will cause level and slope shifts in the long-run oil price trend. Whether the drivers follow a nonstationary or stationary process does not matter, the long-run price will revert to the long-run oil price trend line that reflects the long-run total marginal cost. Since long-run forecasting of these drivers is difficult, Pindyck instead proposes to use a class of time series econometrics models with simple structure and a time-varying trend for long-run oil price forecasting. While no tests of statistical significance or stability are offered, the estimated model shows good forecasting performance, compared to a model with mean reversion to a constant trend. However, Bernard, Khalaf et al. (2004) test the statistical significance of the class of models proposed by Pindyck on the same dataset, and find no statistically significant parameter instability for the oil price, suggesting that a constant-parameter model would do fine.

Videgaray-Caso (1998:32) finds that the oil price switches between two regimes. Most of the time, the oil price follows a process with low mean, low volatility and high persistence. Every 11th year in average, this regime is temporarily replaced by another regime, characterized by a considerably higher mean, very high volatility and low persistence. After 4-5 years in average, the regime switches back again.

More recently, Radchenko (2005) finds that combining a shifting trend model with autoregressive and random walk models considerably decreases the forecast error compared to the error from the individual models.

Although the studies of Pindyck (1999), Videgaray-Caso (1998) and Radchenko (2005) all use the same dataset, consisting of yearly oil prices from 1870 to 1996, they suggest that shifts in the oil price process are common, and that there are theoretical reasons for them.

2.3.2 Shifts and events in linear models

Dummy variables are frequently used to ensure a satisfactory data fit of linear forecasting models.

The comprehensive world market forecasting model of Amano (1987) employs dummies for the years 1974, 1975, 1976, 1977, 1980 and 1985, scattered around its different equations, and documented as “other dummies”. Kaufmann (1995) uses dummies to represents the shocks in 1973, 1974 and 1986, and dummies for the strategic behaviour of OPEC in 1983-1985, 1987 and 1989. Kaufmann, Dees et al. (2004) and Zamani (2004) use dummies to capture the considerable, but short-term effect the Iraq/Persian Gulf had on the oil price. Ye, Zyren et al. (2005) find it necessary to include dummies for the 9/11-2001 events in New York, and a dummy for a significant change in OPEC behaviour.

The frequent use of dummies and the often sparse documentation suggest that including corrections for unpredictable events and shifts are a common exercise in linear oil price modelling and forecasting.

2.4 Research questions

Storage theory suggests that a part of the oil price is due to the convenience yield, and that the convenience yield component decreases at a declining rate as inventories build. Most relative inventory level models, however, specify a linear price-inventory relationship. Two exceptions are Ye, Zyren et al. (2002), where the price-inventory relationship is kinked at the normal inventory level, and Ye, Zyren et al. (2006a), where the relationship is quadratic outside a predefined range around the normal inventory level. The promising results achieved by these rather atheoretical exceptions justify asking:

Q1: Is there a structural break in the relative inventory level model, and is this related to the simplistic modelling of the price-inventory relationship?

Pindyck (1999) shows that oil price shifts are theoretically plausible. Videgaray-Caso (1998:32) finds that they occur frequently in practice, which may be one of the reasons why oil price shift dummies are quite common in empirical models. This justifies asking:

Q2: Is the break unrelated to the price-inventory relationship, and rather a consequence of a shift in the oil price process?

2.5 Hypotheses

The questions above lead to the following hypotheses.

H1: There is a break is in the price-inventory relationship, and it can be mitigated by introducing theoretically motivated price-inventory nonlinearities into the model.

H2: There is a break in the oil price process, unrelated to the price-inventory relationship. A model that allows for a shift in the long-run oil price does not break down.

The first hypothesis is primarily investigated by comparing the results from estimating a previously successful relative inventory level model on its original data period (1992-2003) versus an extended data period (1992-2007), and then carrying out the same estimation and comparison using a presumably better model. The second hypothesis is examined by estimating and evaluating a model that allows for a regime switch and a smooth transition in the oil price.

3 Two economic models of oil price and relative inventory level

This section derives two alternative models of the relationship between price and relative inventory level: The baseline model and a competitor denoted the quadratic log-log model. First, the baseline model is extracted from an empirical model in Ye, Zyren et al. (2005). Then, the competing quadratic log-log model is derived in order to solve two elasticity issues with the baseline model.

3.1 The baseline linear relative inventory level model

The empirical baseline model to be analysed later is based on the empirical model denoted RSTK in Ye, Zyren et al. (2005:494), described in section 2.1.1 above. Here, a modified version is briefly restated, before the static and dynamic formulations of the embedded economic model are extracted. Lastly, two issues with the elasticity of oil price with respect to relative inventory level are raised.

3.1.1 The empirical model RSTK from Ye, Zyren et al. (2005)

Disregarding the dummy variables, and with a parameterised number of lags, the empirical RSTK model in Ye, Zyren et al. (2005:494) may be stated as

$$OP_t = \alpha_0 + \alpha_1 OP_{t-1} + \sum\nolimits_{i=0}^{k} \beta_i RIN_{t-i} + \varepsilon_t ,$$

where OP_t is the oil price in month t and RIN_t is the relative inventory level in month t. The relative inventory level is defined as the difference between current inventories IN_t in month t and the normal inventory level IN_t^* for that month, i.e. $RIN_t = IN_t - IN_t^*$. The normal inventory level is the deseasonalised, detrended inventory level. α_0, α_1 and $\beta_0,...,\beta_k$ are parameters to be estimated, ε_t is the error term, and k is the number of lags.

3.1.2 Static formulation

The static economic model hidden in the empirical RSTK model may be expressed as

$$OP = \alpha + \beta RIN ,$$

i.e. the oil price OP depends linearly on the relative inventory level RIN. At normal inventories, $RIN = 0$, so α may be interpreted as the “normal” oil price.

The relative inventory level theory states that a decrease in relative inventories is accompanied by an oil price increase, so the parameter β is negative: OP is decreasing in RIN.

The elasticity of oil price with respect to relative inventory level expresses the percent change in oil price that results from a one percent change in the relative inventory level. In the static model above, the elasticity of oil price with respect to relative inventory level is

$$\frac{dOP}{dRIN}\frac{RIN}{OP}=\beta\frac{RIN}{OP}$$

i.e. the (point) elasticity depends on the oil price level.

3.1.3 Dynamic formulation

The dynamic economic model embedded in the RSTK model is the partial adjustment model

$$OP_t=\alpha_0+\alpha_1 OP_{t-1}+\sum\nolimits_{i=0}^{k}\beta_{t-i}RIN_{t-i}\ ,$$

which may be considered as a difference equation with the general solution

$$OP_t=\left(OP_0-OP_0^*\right)\alpha_1^t+\frac{\alpha_0}{1-\alpha_1}+\frac{\sum_{i=0}^{k}\beta_{t-i}}{1-\alpha_1}RIN_t\,.$$

Here, $OP_0-OP_0^*$ represents the initial disequilibrium. Assuming that the process has existed for some time and is not unstable ($-1<\alpha<1$ and the RIN DGP is stable), the OP_t will follow the time path

$$OP_t=\frac{\alpha_0}{1-\alpha_1}+\frac{\sum_{i=0}^{k}\beta_i}{1-\alpha_1}RIN_t\ .$$

Since RIN_t, by construction, fluctuates around a zero mean, the oil price will in this model end up fluctuating around a price of $\alpha_0/(1-\alpha_1)$, given that the stability conditions hold. In a long-run steady state $IN_t=IN_{t-1}=IN^*$, so $RIN_t=0$, hence the model implicitly states that the **long-run oil price** is

$$OP^*=\alpha_0/(1-\alpha_1)\,.$$

The long-run relative inventory level multiplier, denoted the ***RIN*-multiplier** is

$$\frac{dOP}{dRIN}=d\left(\frac{\alpha_0}{1-\alpha_1}+\frac{\sum_{i=0}^{k}\beta_i}{1-\alpha_1}RIN\right)=\frac{\sum_{i=0}^{k}\beta_i}{1-\alpha_1}$$

and represents the accumulated effect on the oil price of a change in the relative inventory level.

3.1.4 Two elasticity issues

With the elasticity $\beta(RIN/OP)$, the oil price response to a change in RIN varies with the oil price level OP. More precisely, the marginal elasticity

$$\frac{d}{dRIN}\left(\beta\frac{RIN}{OP}\right)=\beta\frac{1}{OP}$$

suggests that oil price responds more moderately at higher oil price levels. This dependency raises two issues:

Firstly, the dependency states that the oil price is generally less sensitive to inventory changes at higher price levels. It is difficult to find theoretical support for this behaviour.

Secondly, with a negative β, as the model postulates, an increase in RIN is accompanied by a lower oil price. This is consistent with storage theory. However, the marginal elasticity expresses that a percent change in RIN will lead to an increasingly larger percentage change in the OP as the RIN increases and the OP decreases. In the other direction, inventory changes mean less the closer inventories are to stock-out. This is contradictory to the behaviour implied by the storage theory above, where inventory changes become increasingly more important as stocks fall.

3.2 The quadratic log-log relative inventory level model

This section derives an alternative to the linear baseline model. The alternative model is denoted the quadratic log-log model, after its specification, and targets the two elasticity issues in the baseline model. The quadratic log-log model has an elasticity of oil price that is independent of the oil price, and depends nonlinearly on (a redefined) relative inventory level.

3.2.1 Static formulation

First, note that in the linear relative inventory level model, the “relative inventory level” is in fact the absolute difference, measured in million stock tank barrels of oil, between the current inventories IN and the assumed normal level IN^* for a given month.

If the relative inventory level is redefined to instead denote the proportion of the normal level, i.e. $RIN = IN/IN^*$, the RIN will always be a positive figure. Assuming that the relation between price and relative inventories is not linear, but

$$OP=\alpha RIN^{\beta}\text{, or alternatively, that }\ln OP=\ln\alpha+\beta\ln RIN\text{,}$$

the model of oil price now states that β, the elasticity of the oil price (OP) with respect to the proportion of normal inventories (RIN), is constant and equal to β,

$$\frac{dOP}{dRIN}\frac{RIN}{OP}=\frac{d}{dRIN}\left(\ln\alpha+\beta\ln RIN\right)\frac{RIN}{OP}=\beta$$

This mends issue 1 above, in that the elasticity no longer depends on the oil price level.

However, theory suggests that oil price responds more violently as inventories fall. If the new model is modified further, to

$$OP=\alpha RIN^{\beta+\eta\ln RIN},$$

or alternatively,

$$\ln OP=\ln\alpha+\beta\ln RIN+\eta\left(\ln RIN\right)^2,$$

implicit differentiation (assuming $OP=f(RIN)$) with respect RIN to yields

$$dOP/OP=\left(\beta/RIN\right)dRIN+\left(\left(2\eta\ln RIN\right)/RIN\right)dRIN=\left(\beta+2\eta\ln RIN\right)dRIN/RIN,$$

which reordered expresses that the **elasticity** of oil price with respect to relative inventory level is:

$$\frac{dOP}{dRIN}\frac{RIN}{OP}=\beta+2\eta\ln RIN.$$

Here, the additional term $2\eta\ln RIN$ may be interpreted as a nonlinear correction to the constant elasticity β, i.e. a modification of the response of oil price to an inventory change. The magnitude of the correction depends on the level of the RIN and the η parameter.

To be consistent with storage theory, the parameter β should be negative, and the parameter η should be positive. A negative β represents that the oil price increases when inventories decrease from a normal inventory level. A positive η ensures that the positive price effect of an inventory draw becomes larger as inventories decrease down towards stock-out. On the upside, a positive η ensures that inventory changes are not accompanied by large oil price swings when inventories are abundant:

At normal inventory levels, $RIN=1$, so the elasticity of OP with respect to RIN is just β. At inventories less than normal, $RIN<1$ and $\ln RIN<0$. The nonlinearity in the logarithm ensures that the negative contribution to the elasticity increases considerably at levels near stock-out, as RIN goes towards zero. On the upside; as inventories build, the correction term $2\eta\ln RIN$ reduces the elasticity, since $RIN>1\Rightarrow\ln RIN>0\Rightarrow 2\eta\ln RIN>0$. This represents that inventory changes are expected to have less impact at higher inventory levels.

Thus, the introduction of the quadratic term $\eta\left(\ln RIN\right)^2$ ensures that inventory changes become more important at lower inventory levels, which may solve the second elasticity issue found in the baseline model.

Recognising that the variables in the model are all in logs, and using the common convention of using lower case for log variables names, **the static quadratic log-log model** may be written as

$$op = \alpha + \beta rin + \eta rin^2 ,$$

where $op = \ln OP$, $rin = \ln RIN = \ln\left(IN / IN^*\right)$, β represents the constant part of the elasticity of OP with respect to RIN and η represents the nonlinear correction to the constant elasticity, as described above. The α parameter may be interpreted as the oil price when the inventory level is normal: At normal inventory levels, $IN = IN^*$, which gives $RIN = 1$ and $rin = 0$.

3.2.2 Dynamic formulation

The **dynamic quadratic log-log model** is

$$op_t = \alpha_0 + \alpha_1 op_{t-1} + \sum_{i=0}^{k} \beta_i rin_{t-i} + \sum_{i=0}^{k} \eta_i rin^2{}_{t-i}$$

where $op_t = \ln OP_t$ denotes the natural logarithm of the oil price at month t, α_0 and α_1 are parameters that define the long-run oil price $OP^* = \alpha_0 / \left(1-\alpha_1\right)$, $rin_t = \ln RIN_t = \ln\left(IN_t / IN^*\right)$ is the natural logarithm of the ratio of current inventories IN_t to normal inventory level for the IN^* month, k is the number of lags, and $\beta_0,...,\beta_k$ and $\eta_0,...,\eta_k$ are elasticity parameters.

As in the baseline model, the **long-run oil price** is defined by

$$op^* = \alpha_0 / \left(1-\alpha_1\right),$$

with the exception that the oil price is in logs, not in levels.

In the quadratic log-log model, the long-run effect on oil price of a change in relative inventory level consists of the sum of two effects, which later are to be treated separately. The first effect is defined by the constant part of the elasticity, and denoted the **long-run *rin*-multiplier**:

$$\frac{dop}{drin} = d\left(\frac{\alpha_0}{1-\alpha_1} + \frac{\sum_{i=0}^{k}\beta_i}{1-\alpha_1} rin\right) = \frac{\sum_{i=0}^{k}\beta_i}{1-\alpha_1}$$

The second effect is defined by the nonlinear correction term, and in the lack of a better name denoted the **long-run *rin2*-multiplier** (or rin^2-multiplier):

$$\frac{dop}{drin^2} = d\left(\frac{\alpha_0}{1-\alpha_1} + \frac{\sum_{i=0}^{k}\eta_i}{1-\alpha_1} rin\right) = \frac{\sum_{i=0}^{k}\eta_i}{1-\alpha_1}$$

4 Econometric methodology

This section outlines the methodology that the empirical study is to follow, and defines some of the terms and abbreviations that later are used without further reference in the reported results and discussions.

4.1 Data

4.1.1 Use of two different datasets

The study uses two different datasets. The first dataset is chosen to match the data in the original study of Ye, Zyren et al. (2005), where the a relative inventory level model showed good forecasting performance. This dataset serves two main purposes: Firstly, it is necessary to confirm that the original model is adequate. Secondly, since an improvement to the baseline model is a possible conclusion (hypothesis 1), it is fair to use a dataset where the original model reportedly performs well. The second dataset is just an extension of the first dataset, that also includes the most recent period, where Ye, Zyren et al. (2006b) report that the relative inventory level model performs badly. This dataset enables a confirmation of the breakdown of the original model and an investigation of both hypothesis 1 and hypothesis 2.

4.1.2 Descriptive statistics

In the description of the data series, **Mean** and **Variance** denote the arithmetic mean and variance, respectively. **Normality** refers to the results from the D'Agostino, Belanger et al. (1990) test of normality against the null of non-normality.

4.1.3 Unit-root testing

Regressing variables that are non-stationary may indicate meaningful economic relationships in cases where the relationship is coincidental. Unit root tests are used for determining the order of integration of the variables. This is crucial for the models estimated here, since all variables are specified in levels.

The variables are tested for stationarity using various specifications of the augmented Dickey-Fuller (**ADF**) test, where the null is non-stationarity, and two variants of the test of Kwiatkowski, Phillips et al. (1992) (**KPSS**), where the null is stationarity.

The raw inventory level variable is tested for stationarity, even though it does not enter the regression models, because the relative inventory level variables are based on the linearly detrended and deseasonalised inventory level. This detrending is questionable if unit root tests show that inventory level is difference-stationary. The raw inventory level is tested using ADF tests that allow for a time trend, seasonal dummies and a known structural break. This is warranted by the fact that inventories show seasonal variations, and the assumptions of a time trend and a known shift due to an OPEC policy change (Ye, Zyren et al., 2005a:500). The results from the ADF tests are complemented by KPSS test results.

The oil price is tested for stationarity using a number of ADF test variants, included the one proposed by Leybourne, Newbold et al. (1998), which allow for a smooth transition in the alternative, instead of an abrupt shift.

4.2 Specification

Since the empirical analysis basically consists of evaluating the effect of specification changes to a baseline model that fails when it faces a new dataset, the functional form and the set of variables and lags for the estimated models are determined by the choice of baseline model.

The baseline model is a slightly modified version of the linear relative inventory level model denoted RSTK in Ye, Zyren et al. (2005:494). This particular model is chosen as a baseline because it is parsimonious and relevant: Besides dummy variables, the model uses only (relative) inventory level as explanatory variable. Nevertheless, the model reportedly performs better than more developed single-equation forecasting models and is used by the Energy Information Administration, USA (Ye, Zyren et al., 2005a:497).

The second model to estimate – the quadratic log-log model – represents a respecified relationship between inventories and oil price, compared to the baseline model. Since the purpose of this model is to evaluate the effect of this single change, the number of explanatory variables and lags is the same as in the baseline model.

The third model to estimate represents a gradual shift in the long-run oil price. This model is specified as a smooth transition regression model (**STR**), where the parameters related to the long-run oil price are allowed to vary with time. An STR model is preferred before a more common abrupt regime-switch model. The motivation for this is the assumption that the oil price process has changed gradually, not suddenly. The assumption of a smooth change is justified by the absence of factors that historically have been known to cause sudden changes, as e.g. OPEC policy changes (1973-1974), war (1991) or terrorist attacks (2001). The specific type of STR model chosen has only the time trend as transfer variable, as in Lin and Teräsvirta (1994). The use of this type of STR model is justified by the simplicity of hypothesis H2, which translates into an assertion about time-variation in one or two of the parameters. Following the STR modelling strategy devised by Teräsvirta (1998), the actual shape of the transition function is determined in the initial specification stage, where different nonlinear alternatives are tested against the null of linearity.

4.3 Estimation

The baseline model is primarily estimated by OLS using PcGive. For heteroskedasticity-robust statistics and coefficient estimates, Stata is used. The quadratic log-log model is estimated by OLS using PcGive. As for the STR models, JMulTi is initially used for determining the transition function shape, search for initial values and estimation of parameters, including the timing and slope parameters of the transition function. Later, two related versions of the STR model are estimated nonlinear least squares (NLS) and maximum likelihood (ML) estimation in PcGive and Stata. Parameter constancy is largely evaluated using recursive OLS and NLS estimations carried out in PcGive and Stata.

4.4 Evaluation

4.4.1 Residuals and misspecification

The residuals are tested in order to verify some of the least squares assumptions: Normality is preferred, but not required, because the sample size here is large enough to ensure normal sampling distribution without the additional requirement of normally distributed disturbances. No autocorrelation is preferred, because serial correlation may cause biased OLS estimates. Homoskedasticity is preferred, because heteroskedasticity invalidates the standard errors, t-statistics and F-statistics.

In the context of residual analysis, **Normality** refers to the results from the Doornik and Hansen (1994) normality test implemented in PcGive 11, which tests whether the residual distribution have skewness and kurtosis that similar to a normal distribution (Doornik, Hendry et al. (2006:225). **AR** refers to the results from the Lagrange-multiplier test for Nth order autocorrelation implemented in PcGive 11, as described by Doornik, Hendry et al. (2006:257). **ARCH** refers to the Engle (1982) test for autoregressive conditional heteroskedasticity implemented in PcGive 11, as described by Doornik, Hendry et al. (2006:259). **Hetero** refers to the results from the generalised heteroskedasticity test of White (1980), which is a test of the null that residual variances are equal (i.e. homoskedastic) against the alternative that residual variance can be explained by levels and squares of the regressors. **Hetero-X** refers to the results from the generalised heteroskedasticity test of White (1980), with cross products included, i.e. it is a test of the null that residual variances are equal against the alternative that residual variance can be explained by levels, squares and cross products of the regressors. **RESET** refers to the results from the Ramsey's regression error specification test. This is a general test for functional form misspecification, in that it is only a simple test for some kinds of omitted nonlinearities. It tests the null of correct specification against the alternative that powers of the fitted value from the (OLS) estimation have been neglected. The version of the test that is used here is an F-test of the joint significance of the (potentially neglected) powers of the fitted value. **RRESET** refers to the results from a heteroskedasticity-robust version of the Ramsey's regression error specification test in Stata 9.2.

4.4.2 Goodness-of-fit

Two simple measures are used for evaluation and comparison of goodness-of-fit. R^2 refers to the coefficient of determination, which measures of the proportion of the total variation that is explained by the regression model, and has the preferable property that it is unaffected by heteroskedasticity. **AIC** refers to the Akaike information criterion, proposed by Akaike (1974).

4.4.3 Break in the price-inventory relationship and in the oil price process

The possibility of breaks in the price-inventory relationship and the oil price process is investigated by exploring the evolution of coefficient estimates, long-run oil price estimates and long-run multiplier estimates, as generated by recursive estimations.

4.4.4 Long-run statements

Tests of statements about long-run oil price and multipliers are carried out using Wald-type tests in Stata.

5 Empirical methodology and findings

After an initial data description, the data series are tested for unit roots. Then, the baseline linear model and the proposed quadratic log-log model are specified, estimated and evaluated on the two datasets, in order to investigate the hypothesis that the price-inventory relationship is broken, and that a quadratic log-log specification will remedy the problem. Subsequently, two smooth transition versions of the quadratic log-log model are specified, estimated and evaluated, with the purpose of investigating the hypothesis that it is the oil price process that has changed. Finally, the findings are summarised and discussed in the context of the research questions and hypotheses.

5.1 Data

The data series used in the empirical analysis are monthly West Texas Intermediate (WTI) crude oil price and OECD total petroleum inventories from January 1992 to May 2007, both retrieved from Datastream[1].

The oil price time series consists of monthly averages of nominal end of day spot price in USD of one stock tank barrel of oil, delivered F.O.B. at Cushing, Oklahoma.

The inventory level time series consists of end of month levels of stocks of the total of governmental and commercial crude oil and petroleum products in the OECD countries, measured in million stock tank barrels.

The data deviate in two respects, compared with the study that reports the results of the baseline model (Ye, Zyren et al., 2005a). Firstly, the dataset is extended from April 2003 to May 2007, since the purpose here is to investigate the failure in the extended period. Secondly, the inventory data series used here includes governmental stocks and petroleum products, since Ye, Zyren et al. (2002) find this series to give better results than the series consisting of industrial crude oil stocks or commercial stocks alone.

The original dataset extends from January 1992 to April 2003, and is hereafter denoted “1992-2003”. The complete dataset extends from January 1992 to May 2007, and is hereafter denoted “1992-2007”.

---

[1] The oil price and the inventory data series have Datastream codes USPCOWTIA and OCINPP..P, respectively.

The variables used in the models are as follows:

OP_t The average WTI oil price in month t, in US dollars per stock tank barrel.

IN_t The total OECD petroleum inventory level at end of month t, in million stock tank barrels.

IN_t^* The normal inventory level at end of month t, defined as the residual from regressing IN_t on a constant, a linear time trend and seasonal monthly dummies.

RIN_t $= IN_t - IN_t^*$, the linear relative inventory level at end of month t.

op_t $= \ln OP_t$, the natural logarithm of WTI oil price at end of month t.

rin_t $= \ln\left(IN / IN^*\right)$, the nonlinear relative inventory level at end of month t.

Table 1 below characterises the variables that enter the regression models in terms of some simple, descriptive statistics. Figure 2, Figure 3 and Figure 4 on the next page illustrate the evolution of the variables over the 1992-2007 dataset.

Table 1 Descriptive statistics per dataset for the *OP*, *RIN*, *op* and *rin* variables.

Variable	1992-2003	1992-2007
OP		
Mean	21.609	29.045
Variance	28.842	230.529
Normality	0.036 *	0.032 *
RIN		
Normality	0.084	0.075
Skewness	0.263	0.412
op		
Mean	3.043	3.258
Variance	0.061	0.202
Normality	0.493	0.313
rin		
Normality	0.063	0.059
Skewness	0.372	0.570

Note: 1992-2003 denotes the dataset that spans from January 1992 to April 2003. 1992-2007 denotes the dataset that spans from January 1992 to May 2007. *OP* and *RIN* are the oil price in USD and the original relative inventory level, used in the BAS model, whilst *op* and *rin* is the oil price in logs (of the USD price) and the redefined relative inventory level, used in the QLL, STR0, STR1 and STR2 models. Normality and skewness are p-values from the test of D'Agostino, Belanger et al. (1990), where the alternative is non-normality. * and ** denote significance at the 5% and 10% level, respectively.

Figure 2 Evolution of total and "normal" inventories

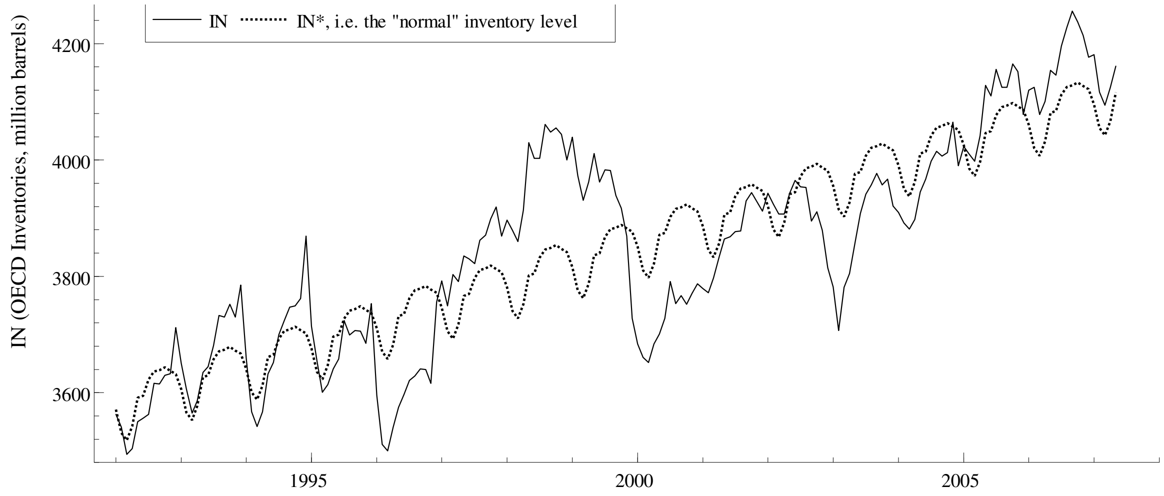

Note: The dotted line illustrates the seasonal and linear trend component that in the relative inventory level models are assumed to represent the "normal" inventory level for a given month. In the original relative inventory level forecasting model, the relative inventory variable is the difference between IN and IN*.

Figure 3 Evolution of the OP and RIN variables used in the linear BAS model.

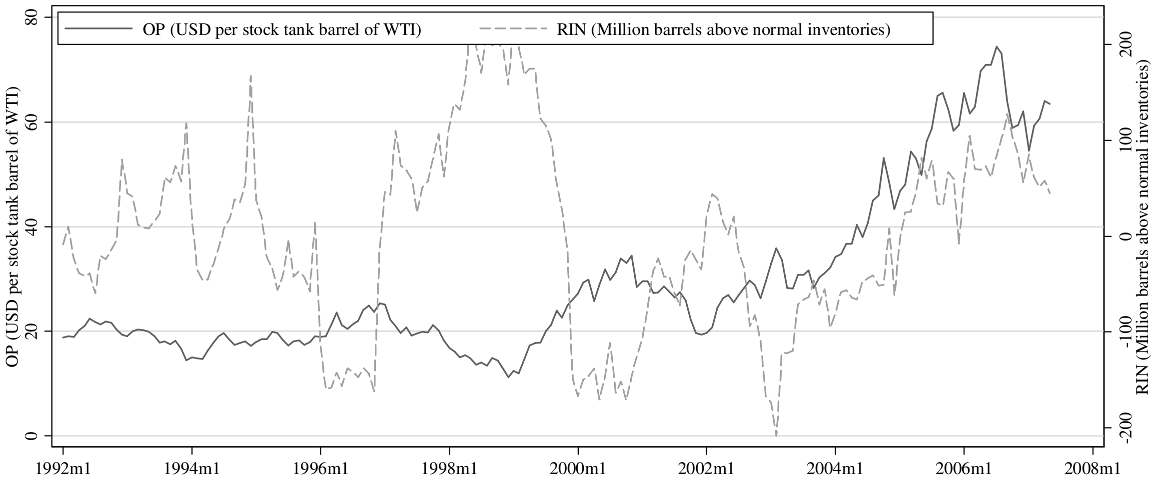

Figure 4 Evolution of the op and rin variables, used in the quadratic log-log specifications (QLL, STR0, STR1 and STR2).

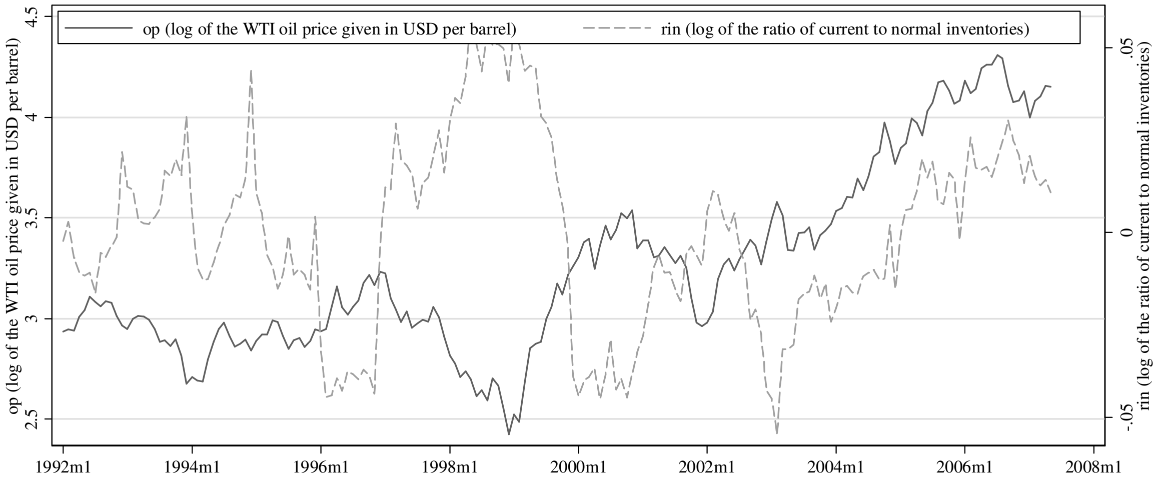

5.2 Unit root tests

The variables are tested for stationarity using various specifications of the augmented Dickey-Fuller test and two variants of the test of Kwiatkowski, Phillips et al. (1992). The results from the tests are reported in Table 2 on the next page.

OP_t and op_t are found to be $I(0)$, assuming that there is a nonlinear shift around 2004. For other specifications of the test, the results are ambiguous or conclude with nonstationarity. IN_t is found to be $I(0)$, assuming a dummy shift in 1999. The other test specifications give mixed results. RIN_t and rin_t are found to be $I(0)$.

The results suggest that if the 1999 shift in IN_t is questionable, the RIN_t and rin_t are possibly derived inappropriately, since they are based on the linearly detrended IN_t. Ye, Zyren et al. (2005:500) argue that there was a policy change by OPEC in 1999, where the organisation started restraining supply. Furthermore, stationarity of OP_t and op_t rests on the assumption of a nonlinear oil price shift around 2004. This is exactly the assumption underlying this thesis. If these two shift assumptions do not hold, the regression results from the relative inventory level models may be spurious.

The conclusion is that the variables are stationary, but that the stationarity requires shift assumptions that may be unrealistic.

Table 2 Unit root test results

Variable	Test statistic	T	Test specification	Result
OP	ADF 0.085	168	const	*I(1)*
	ADF -1.501	173	const+trend	*I(1)*
	ADF -3.675 **	173	const+trend+exponential shift from 2004:9	*I(0)*
	ADF 4.355 **	160	const+trend+logistic shift	*I(0)*
	KPSS 4.335 **	185	const	*I(1)*
	KPSS 1.103 **	185	const+trend	*I(1)*
op	ADF -1.170	172	const	*I(1)*
	ADF -3.222 *	173	const+trend	*I(0)*
	ADF -3.386 *	171	const+trend+exponential shift from 2004:7	*I(0)*
	ADF -4.651 **	160	const+trend+logistic shift	*I(0)*
	KPSS 4.542 **	185	const	*I(1)*
	KPSS 0.910 **	185	const+trend	*I(1)*
IN	ADF -3.471 *	172	const+trend+seasonals	*I(0)*
	ADF -3.189 *	171	as above with dummy shift from 1999:4	*I(0)*
	ADF -3.980 **	171	as above but dummy shift from 1999:12	*I(0)*
	KPSS 4.330 **	185	const	*I(1)*
	KPSS 0.228 **	185	const+trend	*I(1)*
RIN	ADF -3.622 **	172		*I(0)*
	KPSS 0.258	185	const	*I(0)*
rin	ADF -3.647 **	172		*I(0)*
	KPSS 0.249	185	const	*I(0)*

Note: ADF denotes Augmented Dickey-Fuller tests, with critical values from Davidson and MacKinnon (1993:708), from Lanne and Lütkepohl (2002:113) for the tests with logistic or dummy shift, and from Leybourne, Newbold et al. (1998:88) for the tests with logistic shift. For the latter, the critical values are -4.761 (1%), -4.161 (5%) and -3.851 (10%) for T<200. The number of lags determined by the minimum AIC is used if less than 20, in other cases the maximum of the number of lags from minimising final prediction error or the Hannan-Quinn criterion. Where this rule produced zero lags, the number of lags is found by maximising the ADF statistic. KPSS denotes the test of Kwiatkowski, Phillips et al. (1992). For $T = 185$ observations, the critical values are 0.739 (1%), 0.347 (5%) and 0.463 (10%). The test is based on two lags.

5.3 The baseline linear regression model BAS

In this section, a previously successful, but now broken linear relative inventory level model is estimated and evaluated on the two data sets. The estimation is expected to show good fit on the 1992-2003 dataset, as in the original study of Ye, Zyren et al. (2005). With the test results from the ADF (with intercept) unit root test in mind, the linear relative inventory model is expected to reveal a unit root in the oil price when estimated on the 1992-2007 dataset.

5.3.1 Specification

The baseline model, denoted BAS, is specified as

$$OP_t = \alpha_0 + \alpha_1 OP_{t-1} + \sum\nolimits_{i=0}^{5} \beta_i RIN_{t-i} + \varepsilon_t \,.$$

Compared to the original model from Ye, Zyren et al. (2005), there are some minor modifications: Firstly, the number of lags is increased from three to five, in order to make the model somewhat more general and consistent with the model in Ye, Zyren et al. (2002). Secondly, the dummy variables representing Twin Tower terrorist attacks and a OPEC policy shift are excluded from the model, because they may obscure the effects of a specification change.

5.3.2 Estimation

The model is estimated by OLS using PcGive 11.04 and Stata 9.2 on the two datasets 1992-2003 and 1992-2007.

5.3.3 Evaluation

1992-2003 dataset

For the 1992-2003 dataset, the residual statistics (in Table 5) show no signs of nonnormality or autocorrelation, but significant heteroskedasticity (also autoregressive conditional heteroskedasticity). RIN normality is not rejected (Table 1), so incorrect data transformation, wrong functional form or omitted variables stand out as possible causes for the heteroskedasticity. However, the R^2 shows a fairly good fit, and the insignificant robust RESET statistic suggests that no variables are omitted, leading to the conclusion that there may be something wrong with the data transformation or functional form of this particular model of the 1992-2003 data generating process.

The estimate of the autoregressive coefficient (a1=0.9 for BAS on 1992-2003 in Table 6) suggests that the OP process is mean-reverting, with a long-run oil price of 21.89 (BAS on

1992-2003 in Table 8). The long-run *RIN* -multiplier is -0.04 (Table 8), significant at the 1% level (Table 9). (The significance is from a heteroskedasticity-robust F-test).

1992-2007 dataset

For the 1992-2007 dataset, the residual statistics from the OLS estimation of the BAS model (in Table 5) reveal the additional problem of nonnormality. While the high coefficient of determination (0.97) seems to indicate a good fit, the partial R^2 (not reported here) reveals that virtually all of the variance is explained by the autoregressive term. The heteroskedasticity-robust autoregressive coefficient estimate (a1=1.003** for BAS on 1992-2003 in Table 6) suggests that the oil price process is non-stationary. The *RIN* coefficient estimates (Table 6) appear to be little affected by the extension of the dataset.

Parameter constancy

Comparing the two sets of long-run estimates for the BAS model in Table 8, it is evident that there has been some break in the parameter constancy.

A recursive estimation of a rolling 48-month window is carried out in order to study the evolution of the long-run oil price and *RIN*-multiplier, and the resulting Figure 5 indicates two instable periods: One starting late 2004, and another late in 2006.

The model is also estimated recursively with an increasing window size, with results shown in Figure 6, Figure 7 and Figure 8. The recursive residuals in Figure 6 suggest that the regression error increases, particularly around 2000 and 2004-2006, two periods that also show an increase in price. The evolution of the recursive statistics for the intercept and autoregressive coefficients (Figure 7) conveys that around 2004, the autoregressive coefficient estimate shifts from about 0.9 to unity, and the constant coefficient drops to zero.

While recursive statistics for the individual RIN coefficients (Figure 8) are inconclusive, the evolution of the sum of the RIN coefficients (Figure 9) indicates that the RIN coefficients tend towards zero and lose their explanatory power in 2004

5.3.4 Findings

The results indicate a break in the baseline linear regression model in late 2004, when the oil price process changed from mean-reverting to non-stationary. The resulting uncertainty in the model limits conclusions about a change in the price-inventory relationship. The finding of a regression error that increases with price lends some support to the argument that this model may have a problem representing the elasticity of oil price with respect to inventories: If this elasticity in reality is constant, residual error that increases with the price level is to be expected. In any case, this result suggests that there is some is some systematic variance that ends up in the residuals, and that the model is misspecified. This is supported by the findings of heteroskedasticity and indications of wrong functional form or data transformation.

Figure 5 Evolution of the long-run oil price and long-run RIN multiplier in BAS

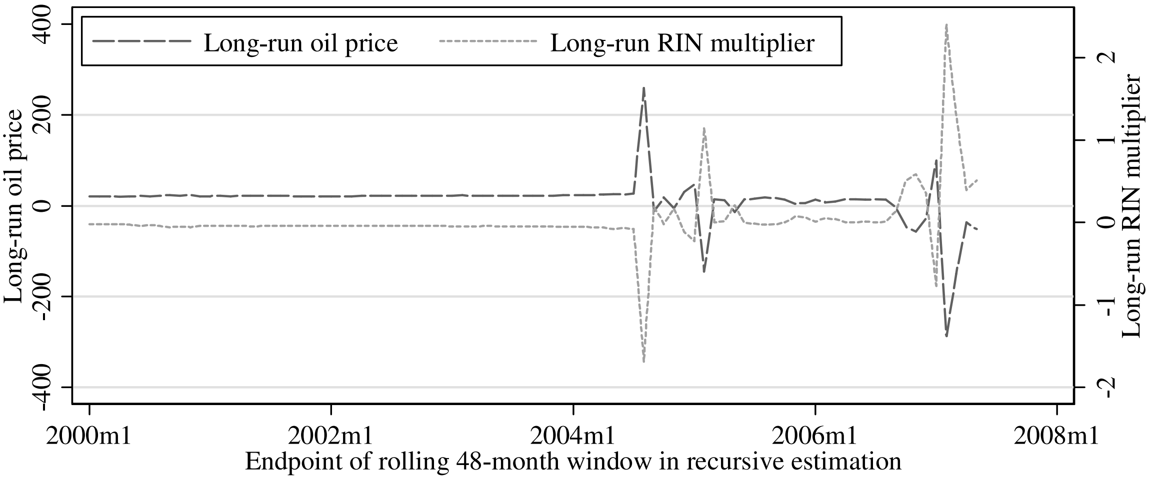

Note: Results are from a 48-month rolling window recursive OLS estimation of the BAS model. Initial period is excluded for improved readability.

Figure 6 Recursive residuals in BAS

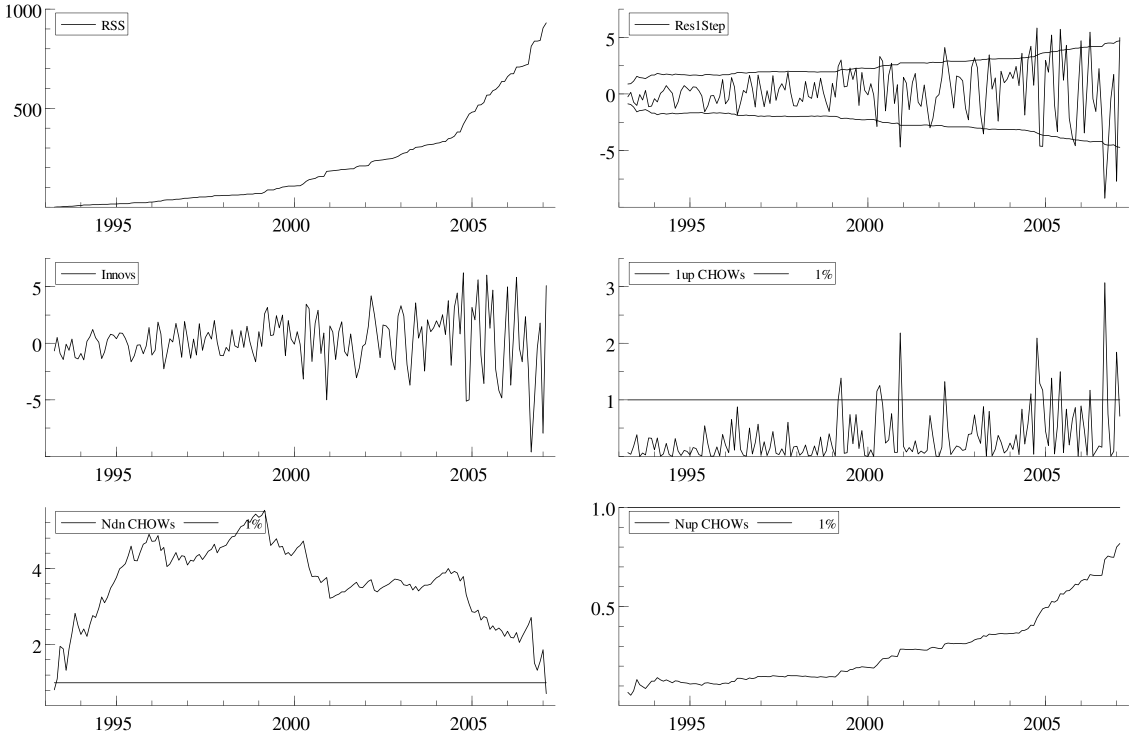

Note: Residuals are from a recursive expanding-window least squares regression of the BAS model on the 1992-2007 dataset. Confidence bands and test limits are invalid because of heteroskedasticity.

Figure 7 Evolution of the intercept and autoregressive coefficients in BAS.

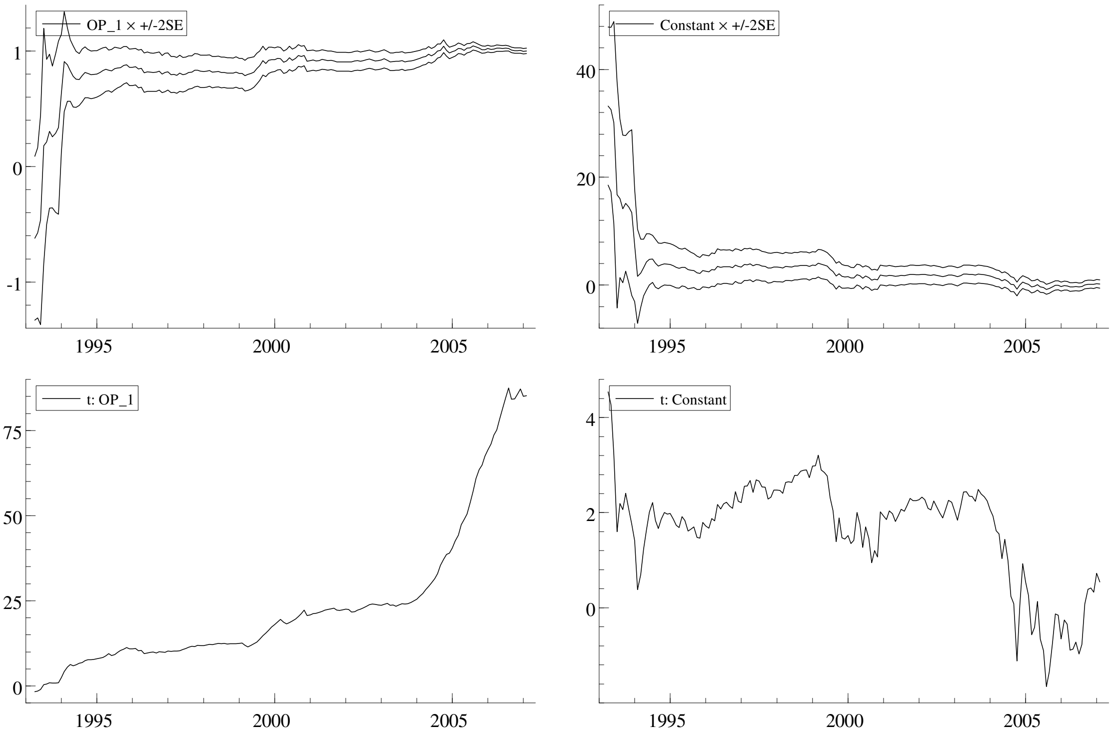

Note: Residuals are from a recursive expanding-window least squares regression of the BAS model on the 1992-2007 dataset. Confidence bands and test limits are invalid because of heteroskedasticity.

Figure 8 Evolution of individual RIN coefficients and corresponding t-values in BAS

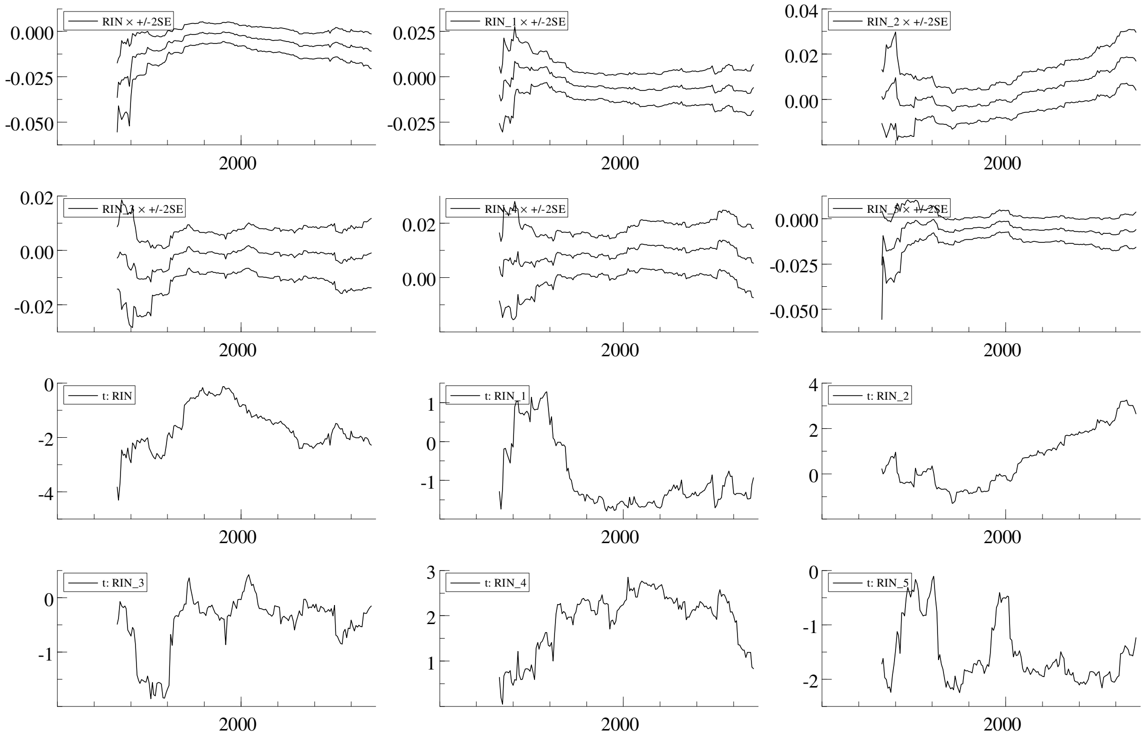

Note: Residuals are from a recursive increasing-window least squares regression of the BAS model on the 1992-2007 dataset. Confidence bands and test limits are invalid because of heteroskedasticity.

Figure 9 Evolution of the sum of RIN coefficients in BAS

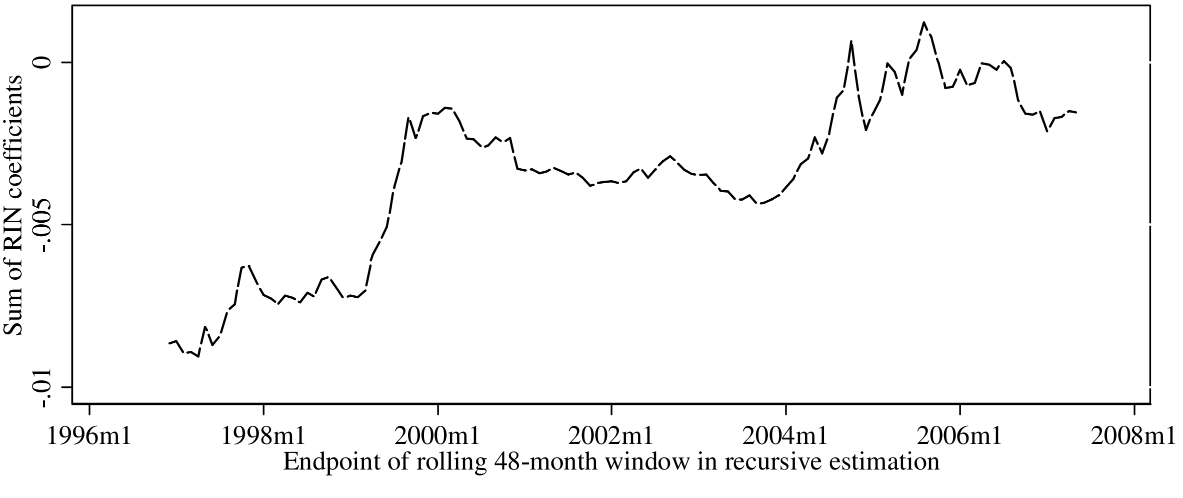

Note: Residuals are from a recursive rolling-window least squares regression of the BAS model on the 1992-2007 dataset.

5.4 The quadratic log-log regression model QLL

In this section, the proposed quadratic log-log alternative to the baseline model is specified, estimated and evaluated on the same two datasets. Since the proposed model presumably solves some issues with the baseline model, the results are expected to show a better data fit than the baseline linear model. Without a shift represented in the model, the estimations are expected to reveal a unit root when the complete dataset is employed.

5.4.1 Specification

A dynamic econometric model corresponding to the quadratic log-log economic model above, denoted the QLL model, is specified as

$$op_t = \alpha_0 + \alpha_1 op_{t-1} + \sum\nolimits_{i=0}^{5} \beta_i rin_{t-i} + \sum\nolimits_{i=0}^{5} \eta_i rin^2{}_{t-i} + \varepsilon_t ,$$

where $op_t = \ln OP_t$ denotes the natural logarithm of the oil price at month t, $rin_t = \ln RIN_t$ is the natural logarithm of the ratio of current inventories IN_t to normal inventories IN_t^*, and $\beta_0,...,\beta_k$ $\eta_0,...,\eta_k$ are elasticity parameters and α_0 and α_1 are parameters that define the long-run oil price $OP^* = \alpha_0 / (1 - \alpha_1)$, and ε_t is the error term. The number of lags is 5 in order to obtain results comparable to the baseline model.

5.4.2 Estimation

The model was estimated by OLS and NLS using PcGive 11 and Stata 9.2 on the two datasets 1992-2003 and 1992-2007.

5.4.3 Evaluation

1992-2003 dataset

For the 1992-2003 dataset, the diagnostics from the OLS estimation (Table 5) show signs of nonnormality or autocorrelation, and only insignificant autoregressive conditional heteroskedasticity. The R^2 shows a fairly good fit, and the insignificant RESET suggests no misspecification.

The long-run estimates (Table 8) suggest that the op process is mean-reverting, with an estimated long-price of 19.19, significant at the 1% level. The long-run rin-multiplier estimate is significant at the 5% level, and shows that rin has a negative impact on the oil price. The estimated rin^2- multiplier is estimated to 136.54. Thus, is it positive, as postulated by the theoretical model. However, this result is only significant at the 17% level (Table 9).

1992-2007 dataset

For the 1992-2007 dataset, the residual statistics from the OLS estimation of the QLL model (in Table 5) show no signs of nonnormality, autocorrelation or heteroskedasticity. The RESET p-value indicates that the model is not misspecified. However, the autoregressive coefficient estimate (a1=1.001** for QLL on 1992-2007 in Table 6) indicates that the *op* process is non-stationary, which is corroborated by the non-significance of the corresponding long-run estimates in Table 8 and Table 9.

Parameter constancy

By comparing the long-run oil price and multiplier estimates from OLS estimation of the QLL model on the two datasets (Table 8), it is evident that there has been some parameter constancy break in the linear QLL model.

A recursive estimation of a rolling 48-month window is carried out, to study the time development of the long-run oil price and multipliers. The results (Figure 10) suggest that there is a major change in late 2004, with a jump in the long-run price and a sign change in the rin-multiplier, and then serious instability from late 2006.

The QLL model is also estimated recursively using OLS and with an increasing window size. The results are shown in Figure 13, Figure 14, Figure 15 and Figure 16 below. The recursive residuals (Figure 16) are rather stable. The RSS curve is fairly linear, but suggests some instability between 1996 and 1999. The instability is confirmed by the Chow test diagram, which shows that there was a break in this period. There is no apparent systematic development in the variance, as found in the corresponding BAS evaluation. The recursive statistics for the intercept and autoregressive coefficients (Figure 15) show results similar to the BAS model: In 2004, the autoregressive coefficient estimate shifts from about 0.9 to unity, and the intercept coefficient drops to zero. The recursive statistics for the individual coefficients of rin and rin^2 (Figure 13 and Figure 14, respectively) are inconclusive, as for the comparable results from the BAS model.

The graphs depicting the evolution of the total coefficients show that the rin-coefficients in sum change around 2004 (Figure 11), but that the rin^2-coefficients in sum do not. With possible nonstationarity in the oil price process, this result is not likely to indicate anything.

5.4.4 Findings

The results from the OLS estimations of the QLL model on the two datasets indicate that around 2004, there is a break in the QLL model. As in the BAS model, the oil price process changes from mean-reverting to non-stationary in 2004. The resulting uncertainty limits conclusions about a break in the price-inventory relationship. However, one interesting finding from the stable period 1992-2003 is that the misspecification evident from the BAS model is not present in the QLL model.

Figure 10 Evolution of the long-run oil price and long-run rin multiplier in QLL

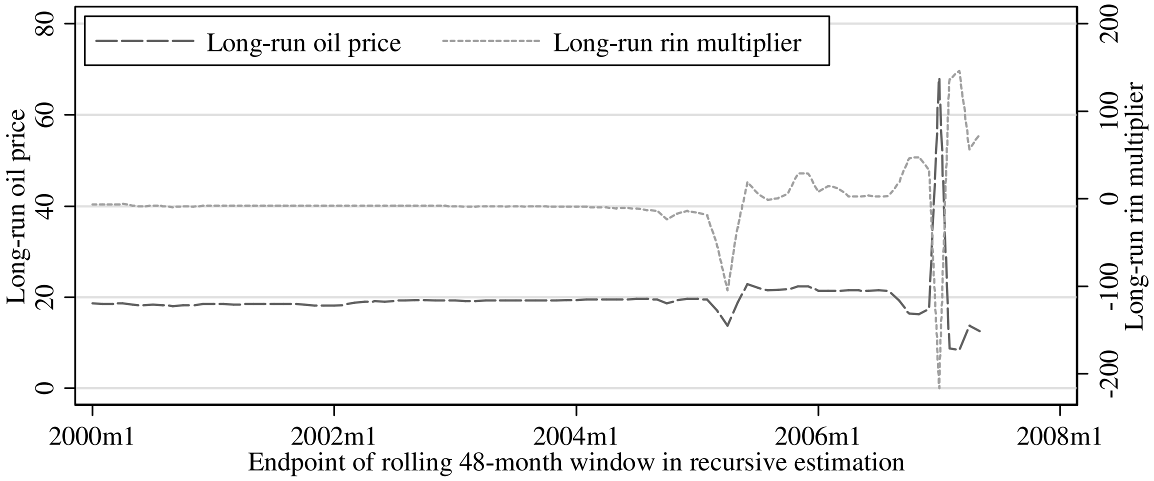

Note: Results are from a 48-month rolling window recursive OLS estimation of the QLL model. Initial period is excluded for improved readability.

Figure 11 Evolution of the sum of rin coefficients in QLL

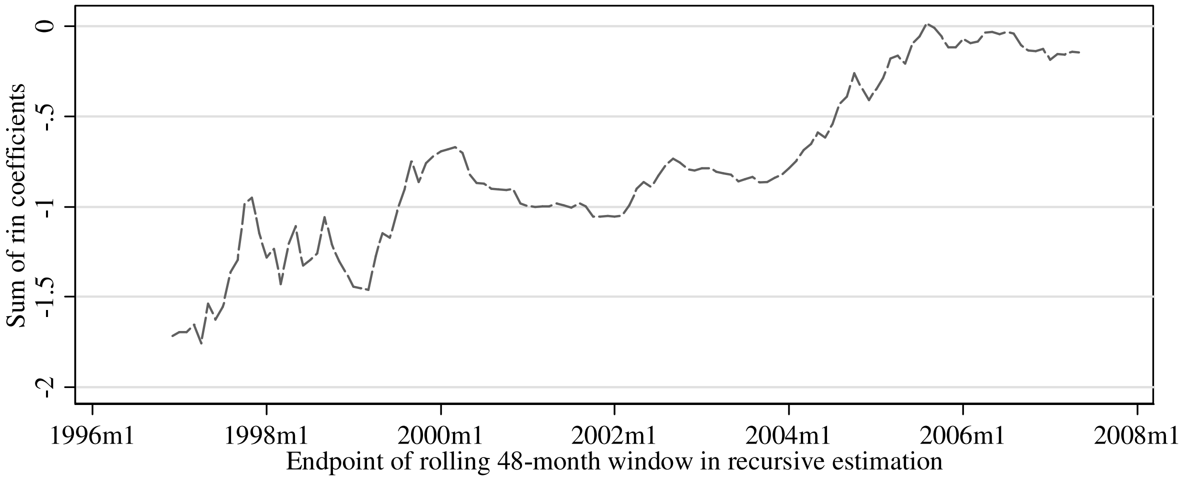

Figure 12 Evolution of sum of rin2 coefficients in QLL

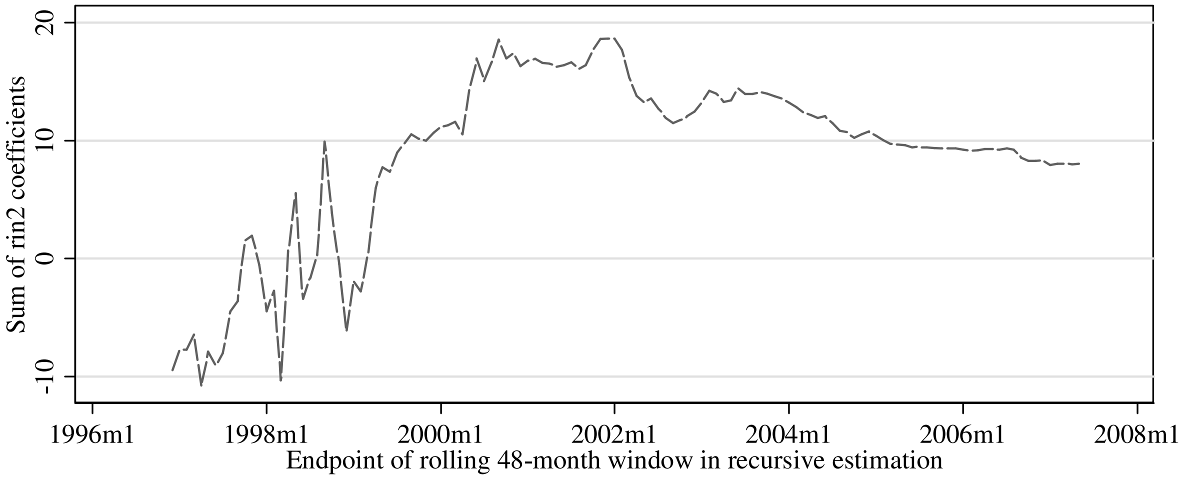

Note: Results are from a rolling-window recursive OLS estimation of QLL on the 1992-2007 dataset.

Figure 13 Stability of the individual rin coefficients in QLL.

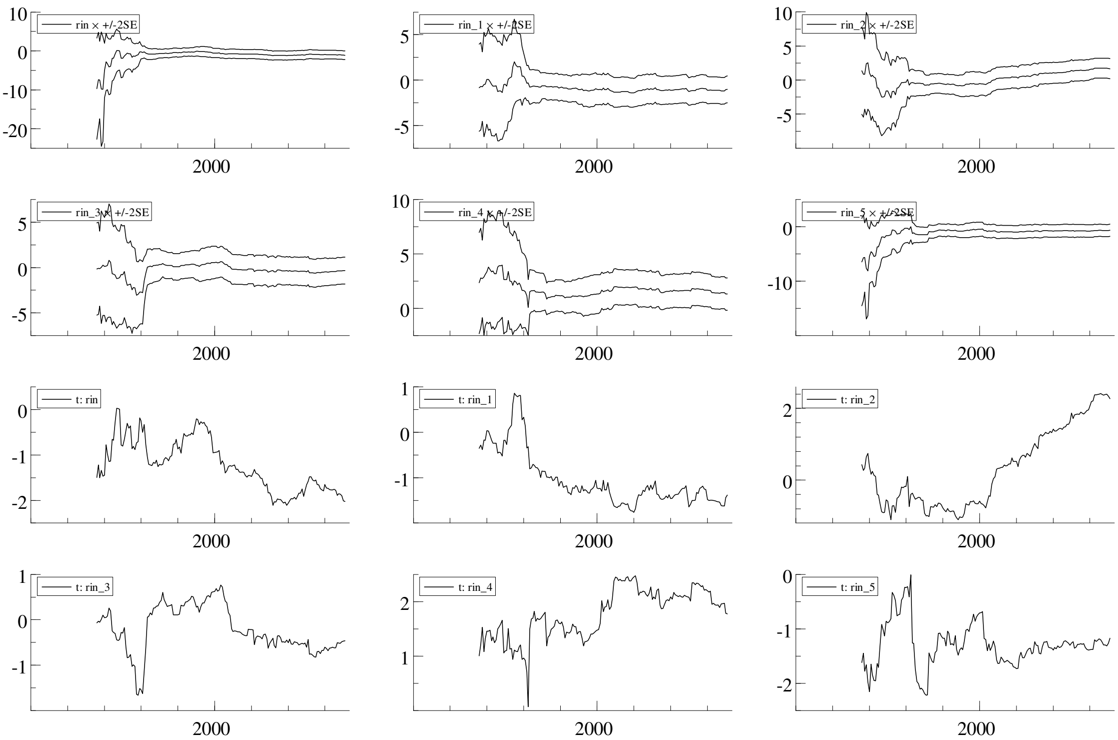

Note: Results are from an increasing-window recursive OLS estimation of QLL on the 1992-2007 dataset.

Figure 14 Stability of the individual rin2-coefficients in the QLL model.

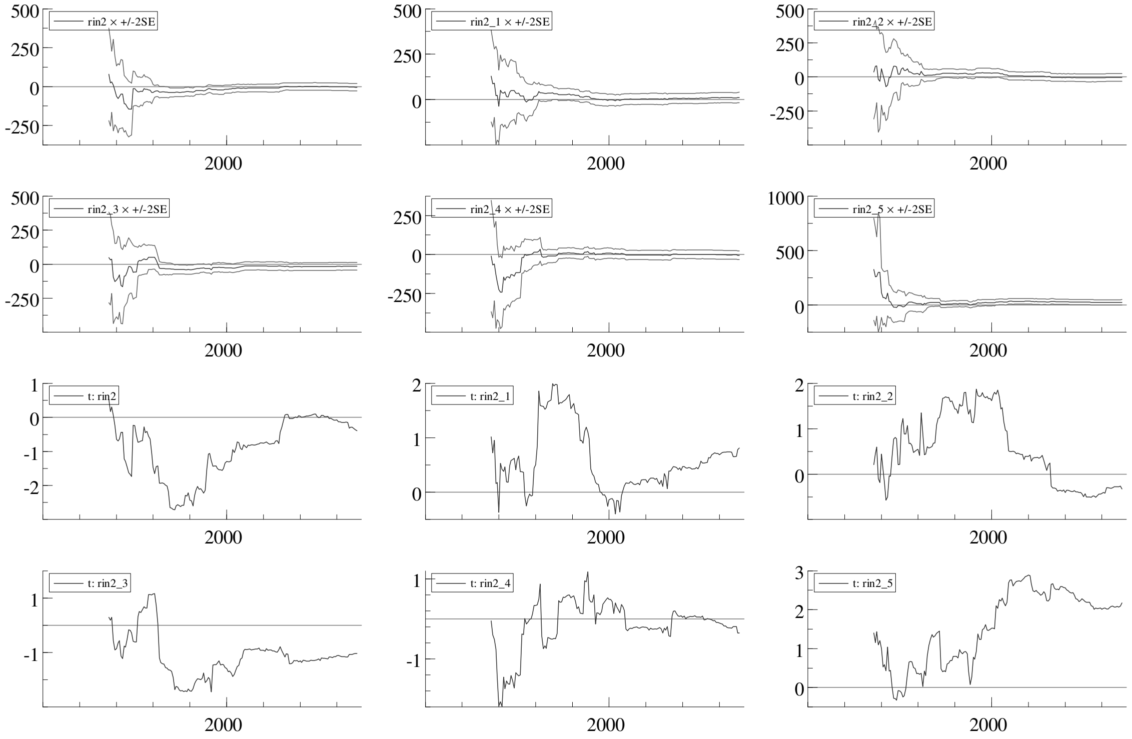

Note: Results are from an increasing-window recursive OLS estimation of QLL on the 1992-2007 dataset.

Figure 15 Stability of the intercept and autoregressive coefficients in the QLL model.

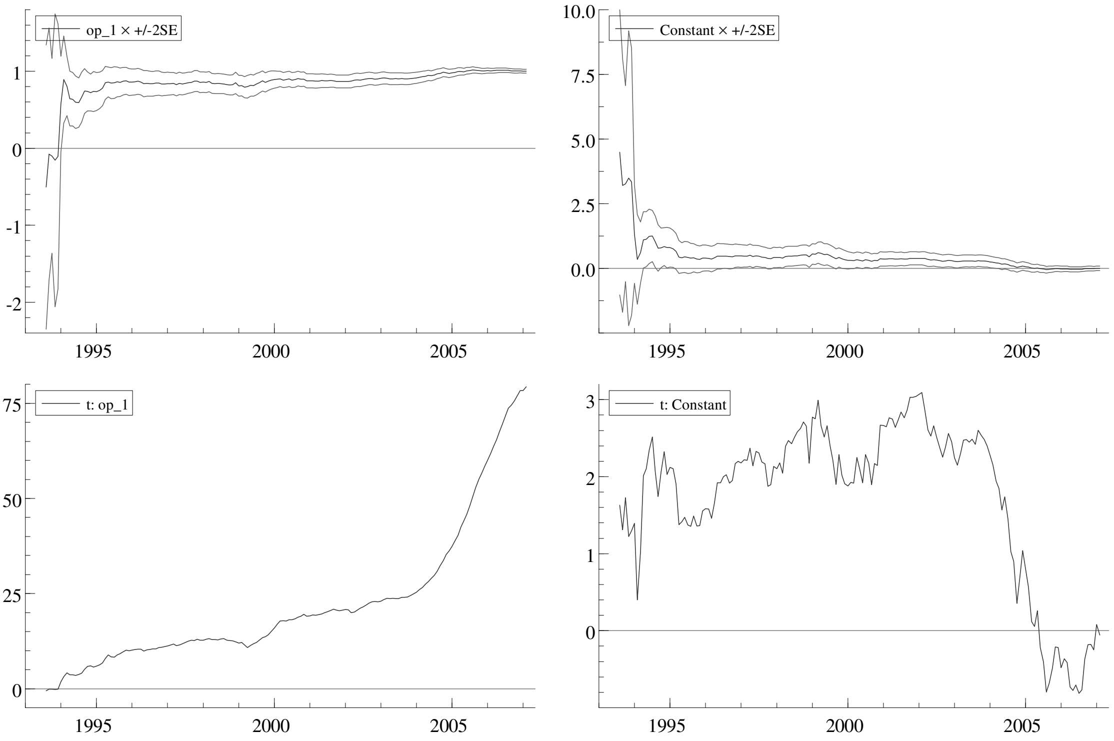

Note: Results are from an increasing-window recursive OLS estimation of QLL on the 1992-2007 dataset.

Figure 16 Recursive residuals in QLL.

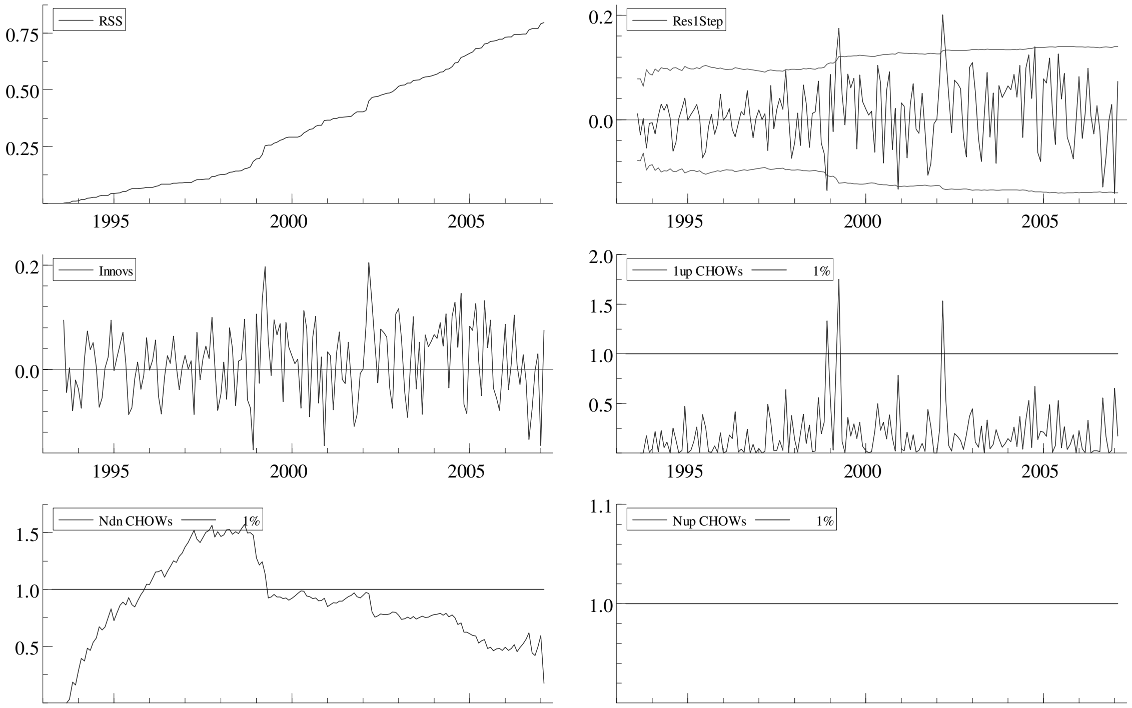

Note: Residuals are from an increasing-window recursive least squares regression of the QLL on the 1992-2007 dataset.

5.5 The smooth transition regression model STR0

In this section, a smooth transition regression version of the QLL model is specified, estimated and evaluated.

5.5.1 Specification

Formulation

Since the quadratic log-log specification of the QLL model above appears to have better statistical properties than the BAS model, the QLL is used as a starting point for the STR0 specification. Therefore, a quadratic log-log regression model with a smooth transition nonlinear in the intercept and autoregressive parameters, denoted STR0, is specified as

$$op_t = \alpha_0 + \alpha_1 op_{t-1} + \sum\nolimits_{i=0}^{5} \beta_i rin_{t-i} + \sum\nolimits_{i=0}^{5} \beta_i rin^2{}_{t-i} + (\theta_0 + \theta_1 op_{t-1}) F(s_t) + \varepsilon_t,$$

where s_t is the time trend, and $F(s_t)$ is a function of the time trend, bounded between zero and one, and ε_t is the error term. In this setup, the θ_0 and θ_1 coefficients represent time-varying additions to the intercept α_0 and the autoregressive coefficient α_1, respectively. The relative inventory level relationship is assumed consistent throughout the whole period, in accordance with hypothesis 2. This is reflected in the model in the sense that the rin and rin^2 variables do not enter the nonlinear part, and that the $\beta_0,...,\beta_5$ are fixed parameters.

Linearity test

The test of the null of linearity against the alternative of nonlinearity of the single logistic (LSTR1) type rejects linearity at the 2% level. Since the alternative with LSTR1 type of nonlinearity leads to stronger rejection than the two other types of nonlinearity, the linearity test suggests the logistic transition function

$$F(s_t;\gamma,c) = \left(1+\exp\left(-\gamma(s_t - c)\right)\right)^{-1}$$

where s_t is the time trend, γ controls the slope and c determines the location. The regression model to be estimated is therefore

$$op_t = \alpha_0 + \alpha_1 op_{t-1} + \sum\nolimits_{i=0}^{5} \beta_i rin_{t-i} + \sum\nolimits_{i=0}^{5} \beta_i rin^2{}_{t-i} + (\theta_0 + \theta_1 op_{t-1})\left(1+\exp\left(-\gamma(s_t - c)\right)\right)^{-1} + \varepsilon_t$$

5.5.2 Estimation

The model is estimated using maximum likelihood estimation in JMulTi 4. Coefficient estimates and diagnostics are reported in Table 3 below.

Table 3 Coefficient estimates and diagnostics for STR0

variable	start	estimate	SD	t-stat	p-value
----- linear part ------					
CONST	0.354	0.349	0.118	2.961	0.004
op(t-1)	0.879	0.882	0.039	22.686	0.000
rin(t)	-1.205	-1.283	0.534	-2.404	0.017
rin2(t)	-3.997	-3.637	11.239	-0.324	0.747
rin(t-1)	-1.229	-1.244	0.697	-1.786	0.076
rin2(t-1)	10.136	9.659	13.395	0.721	0.472
rin(t-2)	1.357	1.348	0.696	1.936	0.055
rin2(t-2)	-4.254	-4.402	13.462	-0.327	0.744
rin(t-3)	-0.313	-0.289	0.696	-0.415	0.679
rin2(t-3)	-13.922	-13.775	13.380	-1.030	0.305
rin(t-4)	1.224	1.237	0.695	1.781	0.077
rin2(t-4)	-5.547	-4.858	13.356	-0.364	0.717
rin(t-5)	-0.650	-0.686	0.530	-1.294	0.198
rin2(t-5)	28.694	29.238	11.143	2.624	0.010
---- nonlinear part ----					
CONST	0.970	1.156	0.677	1.706	0.090
op(t-1)	-0.189	-0.239	0.165	-1.452	0.149
Gamma	4.852	7.830	2.429	NaN	NaN
C1	149.138	148.181	4.144	NaN	NaN
---- diagnostics ----					
AIC: -5.330					
R2: 0.982					
sigma: 0.066					

Note: Start denotes the initial value. SD is the standard deviation. In the nonlinear part, gamma denotes the slope parameter of the transfer function, and C1 is the shift location parameter, representing the mid-point in time of the logistic shift. The estimated value of 148 corresponds to April 2004.

5.5.3 Evaluation

Residuals

The tests of the null of no residual autocorrelation do not reject at any lags, suggesting errors are not serially correlated. The ARCH-LM test finds no autoregressive heteroskedasticity, and the Jarque-Bera test does not reveal nonnormality.

Remaining nonlinearity

Tests of the null of no remaining nonlinearity against two alternatives with nonlinearity of the smooth transition type are carried out, using op, rin and rin^2 including lags as possible transfer variables, and excluding rin and rin^2 terms from the alternatives. The results suggest that there is some significant remaining nonlinearity associated with the rin_{t-5}^2 term, but this is not pursued further.

Parameter constancy

The parameter constancy tests of Lin and Teräsvirta (1994) tests the null of constant parameters against three different alternatives with smooth, continuous change in the parameters. Test results for different combinations of parameters are reported in Table 4 below.

Table 4 P-values from F-tests of parameter constancy of the quadratic log-log STR0 model against three different types of time-varying nonconstancy.

Null hypothesis	A1	A2	A3
a0, a1, rin and rin2 parameters are all constant	0.0199	0.0711	0.0953
a0 and a1 are constant	0.0329	0.0278	0.1117
a0 is constant	0.0793	0.0404	0.0553
a1 is constant	0.1081	0.0514	0.0529
All but rin parameters are constant	0.0681	0.0259	0.0604
All but rin2 parameters are constant	0.0188	0.0335	0.1017
All but a0 are constant	0.0793	0.0404	0.0553
All but a1 are constant	0.1081	0.0514	0.0529

Note: Remaining parameters not mentioned are assumed constant in the null hypotheses. A1 represents a monotonic change over time, A2 a nonmonotonic change symmetric around a point in time, and A3 a more complex STR-type time-varying change.

The null hypothesis that α_0 and α_1 are constant is rejected (for all alternatives, but most strongly against the A2 alternative). This suggests that the specified model does not capture all the time-variation in the α_0 and α_1 parameters.

The null hypothesis that all parameters are constant is rejected against all alternatives, but most strongly against the A1 alternative. The test of the null hypothesis that rin and rin^2 coefficients are constant failed to execute.

Although the parameter constancy tests reject the STR0 model, they may be interpreted as support of the two nonlinear time-varying alternatives A1 and A2. Since these alternatives are parameterised, they may be investigated further using nonlinear least squares estimation. The test results suggest that the A2 represents the time-dependent parameter change better than A1, in that it rejects most strongly the null hypothesis that α_0 and α_1 are constant. However, the difference in rejection probabilities is small (2.8% against 3.2%), and the A2 shape is not consistent with the hypothesis 2 of a gradual increase. Therefore, the A1 is selected as the most interesting candidate. However, attempts of estimating the A1 alternative fails. The estimation does not converge, possibly because of sensitivity to initial values. The estimation fails to produce statistics that can convey any meaningful information about the time-varying properties of the α_0 and α_1 parameters.

Graphics

The plots in Figure 17 and Figure 18 below show that some of the time-variation in the α_0 and α_1 may be captured by a logistic transfer function. The transfer function in Figure 18 indicates a smooth shift that starts somewhere between 2001 and 2003. However, a comparison of the explanation power from the linear part (left half of Figure 17) with the nonlinear part (right half of Figure 17) reveals that the LSTR1 nonlinearity captures only a fraction of the time-variation: At the end of the data period, the oil price is about 4.16, measured in logs, where 0.16 of this is explained by the nonlinear part of the model. Converted to USD, this means that the nonlinear part at the end of the data period explains 10 of the 65 USD price. Assuming the long-run price has shifted 45 USD from a mean around 20 USD to another around 65 USD, the logistic transition explains about a quarter of the shift in the oil price.

Figure 17 Fitted values from the linear (left) and logistic (right) parts of the STR0 model.

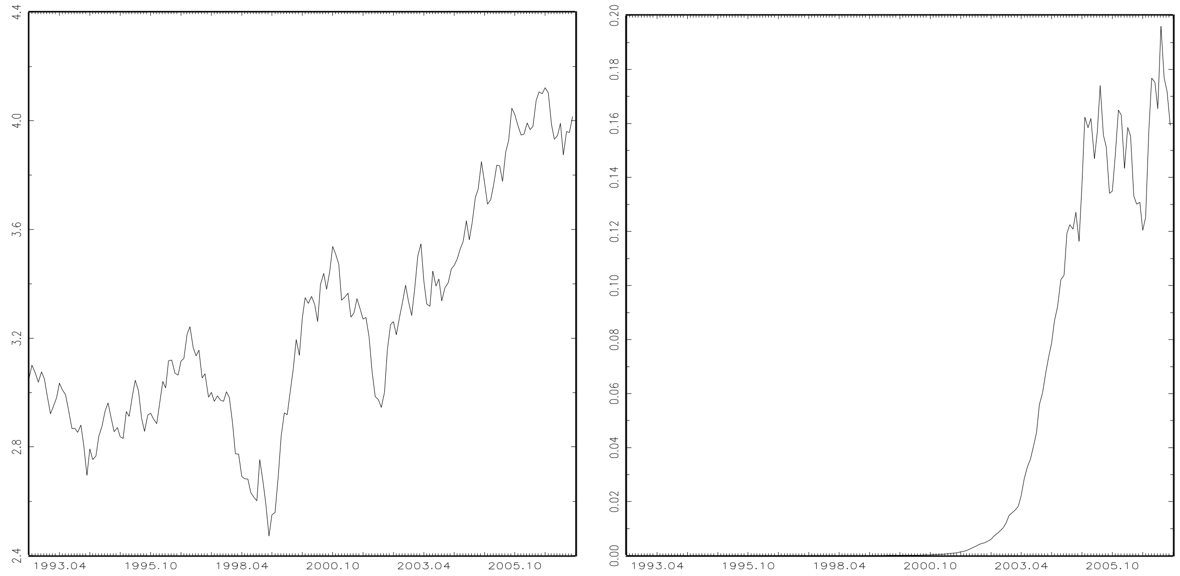

Note: Oil price is in log of nominal WTI price in USD per stock tank barrel.

Figure 18 Estimated transfer function F in the STR0 model.

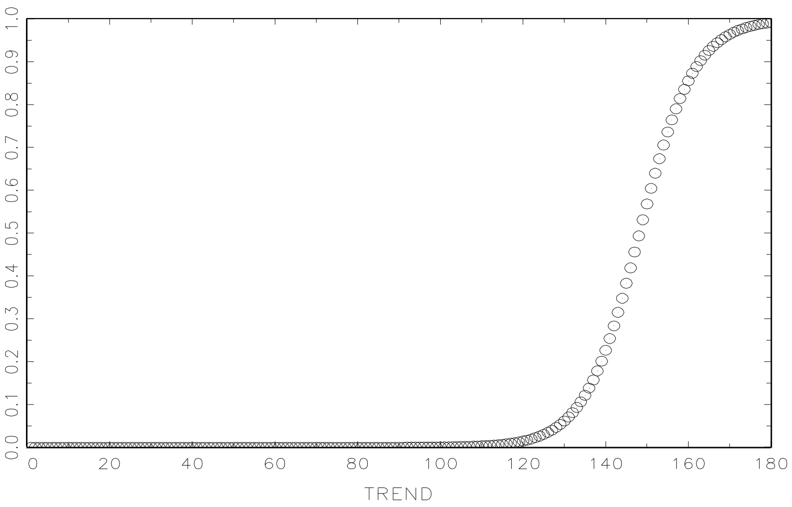

Note: Value of the transfer function on the Y axis. The X axis show time data points, where 0=1992m2, 100=2000m4, 140=2003m8, 150=2004m6 and 180=2006m12. The shift is estimated to have a midpoint in April 2004.

5.5.4 Findings

The main finding from estimation of the STR0 model is that a minor part of the time-variation in the α_0 and α_1 parameters may be represented by a logistic shift that starts around 2001-2003. The parameter constancy tests suggest that the model is misspecified, but lend support to time-varying alternatives. The estimation of the most relevant alternative fails.

5.6 The smooth transition regression models STR1 and STR2

Since the NLS estimation of the time-varying alternative suggested by the STR0 estimation above does not reveal any meaningful information, a plain NLS estimation of the smooth transition model suggested by the linearity tests is carried out, in order to generate results comparable to the previous OLS estimations of the BAS and QLL models.

5.6.1 Specification

A quadratic log-log smooth transition regression model that allows for a shift in the intercept and autoregressive coefficients is denoted STR2 and specified as

$$op_t = \alpha_0 + \alpha_1 op_{t-1} + \sum\nolimits_{i=0}^{5} \beta_i rin_{t-i} + \sum\nolimits_{i=0}^{5} \beta_i rin^2_{t-i} + \left(\alpha_{10} + \alpha_{11} op_{t-1}\right)\left(1 + \exp\left(-\gamma\left(s_t - c\right)\right)\right)^{-1} + \varepsilon_t,$$

where the linear part of the model is unchanged, compared to the QLL specification. As in the QLL model, $op_t = \ln OP_t$ denotes the natural logarithm of the oil price at time t, $rin_t = \ln RIN_t$ is the natural logarithm of the redefined relative inventory level $RIN_t = IN_t / IN_t^*$, and $\beta_0,...,\beta_k$ $\eta_0,...,\eta_k$ are elasticity parameters, and ε_t is the error term. α_0 and α_1 are the intercept and autoregressive parameters in the linear part, and α_{10} and α_{10} are parameters that represents the contribution to these from the nonlinear shift. I.e. before the shift, the implicit long-run oil price is $\alpha_0 / \left(1 - \alpha_1\right)$, and at the end of the shift, the implicit long-run price is $(\alpha_0 + \alpha_{10}) / \left(1 - (\alpha_1 + \alpha_{11})\right)$. The logistic shift function is parameterised by the time trend s_t, the slope γ, and the location parameter c represents the mid-point in time of the logistic shift.

5.6.2 Estimation

Two versions of the model are estimated by NLS and recursive NLS in PcGive 11.04 and Stata 9.2.

5.6.3 Evaluation

Residuals and goodness-of-fit

The residual statistics from the NLS estimations of the models on the complete 1992-2007 dataset (column STR1 and STR2 in Table 5) show no signs of nonnormality, autocorrelation or heteroskedasticity. The coefficient of determination is high (0.98) and RSS is low (0.70-0.71).

Estimates of coefficients, long-run oil price and long-run multipliers

In STR1 - the model with a shift only in the intercept – the long-run oil price estimates from the NLS regressions (Table 8) suggest that the *op* process is mean-reverting around 19.3 until a slow shift with midpoint in April 2004 (c=147.9 in Table 6) brings the long-run oil price estimate to an estimated level at 84.57 (Table 8). The long-run rin-multiplier is estimated to -7.22 (Table 8). These estimates are all significant at the 1% level.

STR2 – the model with a shift only in the intercept – results in a similar long-run price estimate before the shift (19.28 in Table 8). The shift is somewhat steeper (Gamma=0.15 in Table 6, compared to 0.10 for the STR1 model), and the shift has its midpoint a few months later, in September 2004 (c=153.5 in Table 6). The shift flattens out with a long-run price estimate of 67.9 (Table 8). These estimates are all significant at the 1% level, whilst the long-run rin-multiplier is estimated to -7.75 at the 5% level. The α_{11} parameter is estimated to -0.246 (a11 in Table 6), suggesting that at the end of the shift, the process reverts slower to its mean than before the shift. This last result, however, is only significant at the 20% level (Table 7).

As for the long-run rin^2- multiplier, it is estimated to around 81 in the STR1 model, and 103 in the STR2 model (Table 8). Thus, the multiplier is positive as the theoretical QLL model postulates, but the estimates are only significant at the 22% and 20% level, respectively (Table 9).

The slope coefficients are estimated to 0.11 in STR1, and 0.15 in STR2 (Table 6), significant at the 5% and 1% levels, respectively (Table 7) . This result suggests that the shift is slightly steeper in the STR2 model, where time-variation is allowed also in the autoregressive parameter.

Parameter constancy

Table 8 shows that the estimates of the long-run oil prices (before and after the shift) and the long-run rin -multiplier are all significant (Table 9). This suggests that the parameters are rather stable over the entire data period, in contrast to the results from the linear QLL and BAS models.

The STR2 model is subjected to a recursive estimation of a rolling 48-month window, in order to investigate the evolution of the long-run oil price and the multipliers further, with results shown in Figure 19 and Figure 20.

Figure 19 illustrates that the long-run oil price (which in this figure is the total of oil price before shift and the contribution during the shift) is stable around 20 before the price in 2002-2003 starts increasing to the 70-100 price region in 2006-2007.

Furthermore, Figure 19 illustrates that the rin -multiplier is now fairly constant over the entire data period, in contrast with the instability found in the QLL model. One interpretation of this result is that the introduction of time-variation in the long-run oil price has stabilised the price-inventory relationship in the model. Figure 20 suggests that the rin and rin^2-

multipliers in the STR2 model are stable in the period with severe instability in the BAS and QLL models. The long-run rin-multiplier is in the region of -3 to -10, and the long-run rin^2-multiplier is in the 50-150 region, at least for the last five years.

In order to compare the residuals from the STR2 NLS estimation with the BAS and QLL estimations, the STR2 model is also estimated recursively with an increasing window size. The results (Figure 21) suggest that the STR2 model reduces error compared to the QLL: The RSS curve appears even more linear, and the Chow test is within the 1% limit (with one minor exception).

5.6.4 Findings

The NLS estimations of STR1 and STR2 suggest that the relative inventory level parameters are fairly constant in the period 1992-2007, but that the intercept and autoregressive parameters are time-varying. Provided that a logistic time-variation is allowed both in the intercept and the autoregressive coefficient, the results show a smooth shift in the long-run oil price from a stable level around 20 to a new level around 68. The shift starts around 2002, is at its steepest in September 2004, and flattens out towards the end of the data period, May 2007. Furthermore, the oil price process appears to have become less persistent during the shift.

Figure 19 Evolution of long-run oil price and long-run rin-multiplier in STR2

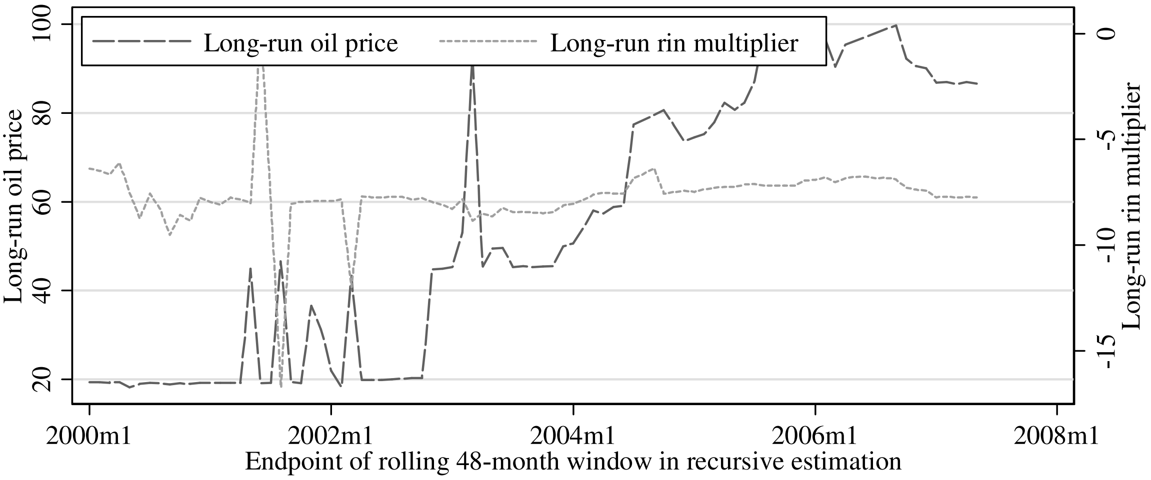

Note: Results are from a 48-month rolling window recursive NLS estimation of the STR2 model. Initial period is excluded to enhance readability. One outlier (due to non-convergence in the optimisation) is excluded.

Figure 20 Evolution of the long-run rin and rin2 multipliers in STR2

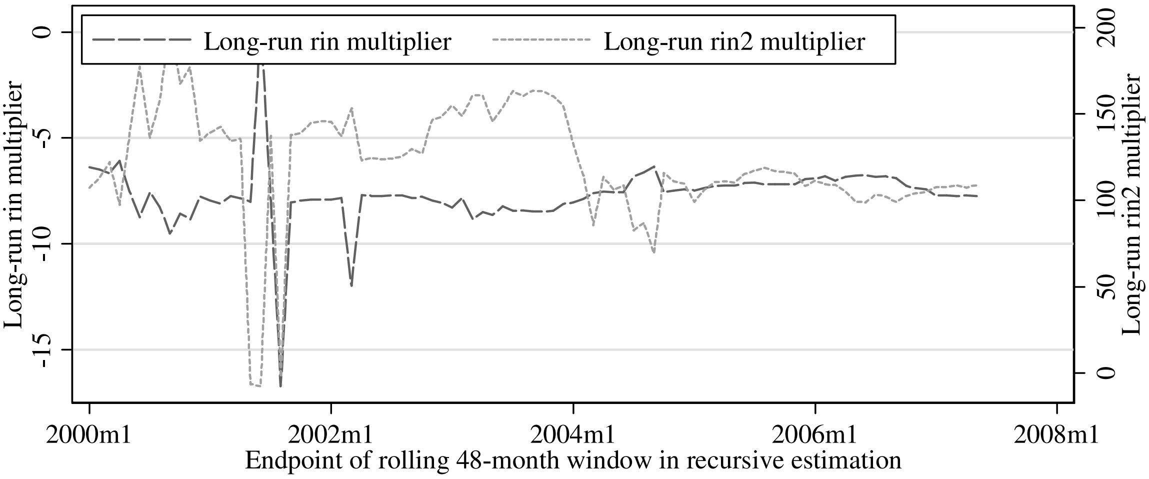

Note: Results are from a 48-month rolling window recursive NLS estimation of the STR2 model. Initial years are excluded to enhance readability.

Figure 21 Residual stability in the STR2 model

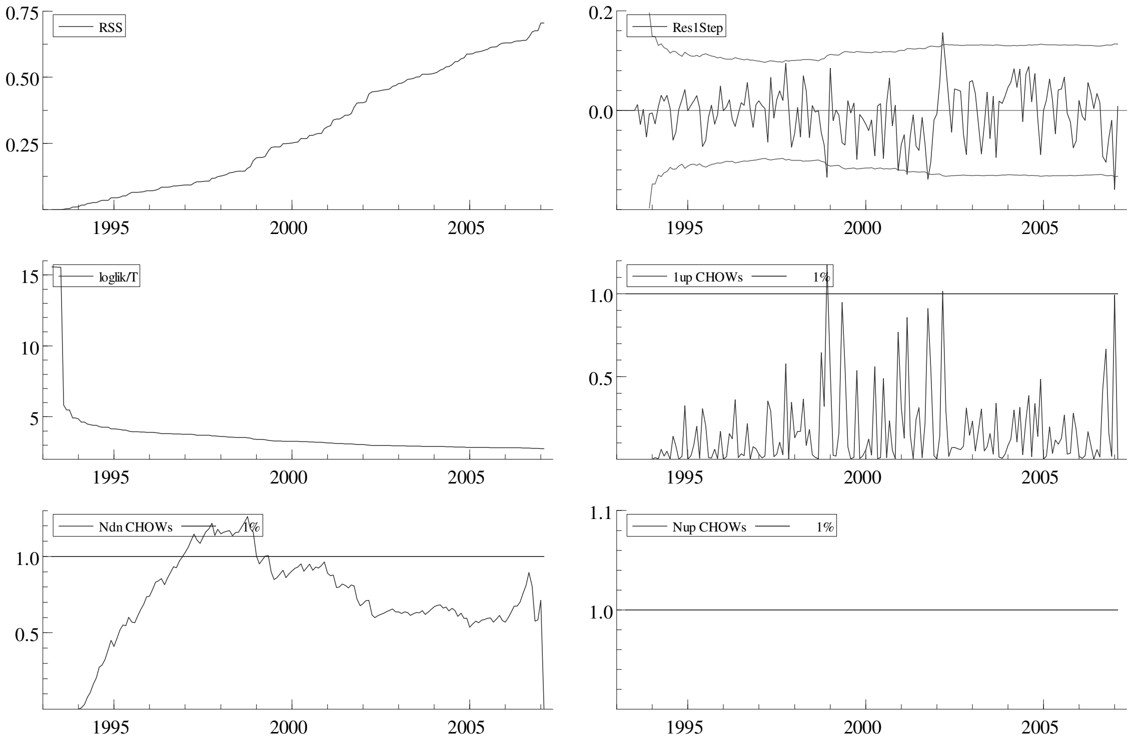

Note: Recursive residuals are from an increasing-window recursive NLS estimation of the STR2 model.

5.7 Summary of findings

This section summarises the findings from the estimation and evaluation of the BAS, QLL, STR0, STR1 and STR2 models above.

5.7.1 Linear regression models on 1992-2003 vs. 1992-2007

Comparing the results from the OLS estimation of the linear models on the 1992-2003 dataset with the results from the 1992-2007 dataset reveals that intercept and autoregressive coefficients are severely affected by the extension of the dataset. In the 1992-2003 dataset, these coefficients suggest that oil price is mean-reverting, with a long-run mean oil price around 20. For the 1992-2007 dataset, the autoregressive coefficient suggests that oil price is non-stationary. This is expected, as the comparable unit root test specification – ADF with constant term only – shows that both OP and op are $I(1)$. As for the long-run multipliers (Table 8) of rin and rin^2, they are significant on the 1992-2003 dataset. The instability of the model limits conclusions about the long-run multipliers from the 1992-2007 dataset.

5.7.2 Baseline model vs. quadratic log-log model

The results from the 1992-2003 dataset suggest that the BAS model is misspecified in that the residuals are heteroskedastic, and that this issue is remedied by the QLL model. In addition, the QLL specification reduces the AIC statistic from 3.8 to -2.5, suggesting that the QLL model represents the data better.

Furthermore, the results suggest that regression error varies with the price level in the BAS model, but not in the QLL model. This finding supports the argument that the elasticity of OP with respect to RIN is not as specified in the original relative inventory level model, but more likely independent of the oil price level, as specified by the QLL model.

In addition, the theoretical model behind the QLL postulates that the elasticity of OP with respect to RIN varies with RIN, formalised through the assertion of a positive η in the elasticity expression. This construction is supposed to capture that an inventory change is more important at low inventory levels. In the empirical model, η is estimated by the long-run rin^2-multiplier. The estimations suggest that the long-run rin^2-multiplier is positive, but the estimate is only significant at the 20% level.

5.7.3 Nonlinear (smooth transition) regression models on 1992-2007

The various estimations of the smooth transition models show that the intercept and autoregressive parameters are time-varying, while relative inventory level coefficients are fairly constant. If the time-variation is represented by a logistic function, a smooth transition towards a higher mean oil price starts around 2000, is steepest in 2004, and flattens out in

2007. The rate of mean reversion is high before the transition, and lower at the end of it. At the end of the transition, the autoregressive coefficient is reduced from 0.88 to 0.64. As for the long-run oil price, the smooth transition models suggest that it is around 19 USD per barrel until 2002, and then starts rising smoothly up to 68 USD per barrel.

5.8 Discussion of findings

In this section, the findings are related to the research questions and hypotheses.

5.8.1 Is there a break in the price-inventory relationship?

The linear regression model estimations apparently confirm the break in the price-inventory relationship referred to by e.g. Ye, Zyren et al. (2006b).

However, these linear models presume that oil price is fluctuating around a constant long-run oil price level. The estimation results show that the relationship between short-term oil price and relative inventory level is fairly stable, provided that the model accommodates for a slow shift in the parameters related to the long-run oil price.

Thus, the conclusion is that there is no break in the price-inventory relationship. The break in the relative inventory level model is concentrated to the intercept and autoregressive parameters. These parameters have changed in the past as well, which is one of the reasons why previous relative inventory level models have resorted to dummy variables in order to achieve sufficient fit.

5.8.2 Is a quadratic log-log specification better than the baseline linear?

The findings, summarised in paragraph 5.7.2 above, suggest that the quadratic log-log specification resolves the misspecification in the linear relative inventory model.

The respecification from the baseline to the quadratic log-log specification consists of two changes. The major change allows for an initially constant elasticity of oil price to relative inventory level, and the minor change is the introduction of a nonlinear correction to the constant elasticity. The results indicate that the nonlinear correction is relevant, but that the greater part of the improvement may be achieved by a respecification to constant elasticity.

The QLL model has two primary advantages: Firstly, it is more aligned with data and represents some of the theoretical nonlinearities that are misrepresented in the BAS model. Secondly, it is well-specified enough to be used for hypothesis testing and computation of confidence intervals. However, data transformation and interpretation is somewhat hampered by the use of logarithmic scale and quadratic terms.

The main advantage of the BAS model is its simplicity and transparency: Data transformation is not very much easier, but intermediate results and outputs may be easier to check and interpret, due to the use of variables in levels. Furthermore, the misspecification is probably not severe enough to distort short-term point forecasts very much. However, hypothesis testing and computation of confidence intervals require corrections that are not straightforward.

Even if the purpose is coarse short-term point forecasting, the baseline relative inventory level model has little to recommend it. A considerable improvement may be achieved by simply calculating relative inventories by division instead of subtraction, and thereafter using variables in logs, as proposed by the quadratic log-log model. However, the value of the additional quadratic term is more questionable.

5.8.3 Is there a break in the oil price process?

The estimation results suggest that the oil price process has changed gradually since 2002: The long-run oil price has increased from 20 to 68, and the oil price process has become less mean-reverting.

This increase in the long-run oil price, evidenced by the results from estimation of the STR1 and STR2 models above, may be consistent with a step up in Pindycks (1999) quadratic-trend model of long-run oil price evolution. Furthermore, if the smooth transition model STR2 is interpreted as a two-regime model, the estimation results suggest that the oil price may be moving from a regime with low mean, low variance and high persistence to a new regime with high mean, high variance and low persistence. This is consistent with the findings of Videgaray-Caso (1998).

On the other hand, the findings require that the oil price is mean-reverting. If oil price is in fact non-stationary, there is no implicit long-run oil price in any of the models estimated here. Thus, it would be pointless to derive statements about long-run oil prices and oil process breaks. In fact, the modelling may very well be void and regressions spurious, since all the models are specified in levels instead of differences.

Furthermore, the findings require that the models estimated do not neglect any explanatory variables. Given that the estimated models are exceptionally simple, with oil price being a function of past oil price and a single variable derived from the inventory level, this is not very likely. There are most likely variables that influence the oil price significantly without any pre-warning in the relative inventory level, e.g. the futures market activity identified by Merino and Ortiz (2005), or surplus production capacity, as suggested by Ye, Zyren et al. (2006c).

However, assuming that the oil price process is mean-reverting and that no explanatory variables are neglected, the conclusion here is that there is a break in the oil price process. The estimation results are consistent with a shift from a low-mean/low-variance/high-persistence regime to a high-mean/high-variance/low-persistence regime.

5.8.4 Is the price-inventory relationship “inverted” or “out of line”?

Figure 1 in the introduction seems to confirm the apparent inversion in the relationship between oil price and inventories referred to by analysts and researchers, e.g. Ye, Zyren et al. (2006b) .

However, the inversion is a result of an attempt to fit a straight line to (price, inventory) data points that are gradually located at a higher price level, possibly due to a slow shift in the long-run oil price. By splitting up the 2004-2007 period, as in Figure 22 below, it is evident that although the static price-inventory regression line is very stable in the 1992-2003 period, it may have been on the move the last years.

The estimations above suggest that the price-inventory relationship is intact, but that the long-run oil price has increased. This may be consistent with the 1992-2003, 2006 and 2007 lines in Figure 22, and is also consistent with analysts observations of a price-inventory relationship that is “out of line” (Tchilinguirian (2006:10)).

Thus, the estimations here offer no evidence that the price-inventory relationship is inverted. On the contrary, the results suggest a fairly consistent price-inventory relationship. Furthermore, assuming that the long-run oil price is changing, “out of line” is an expected outcome.

Figure 22 The price-inventory relationship is probably not inverted, and being “out of line” is an expected consequence if the oil price process shifts.

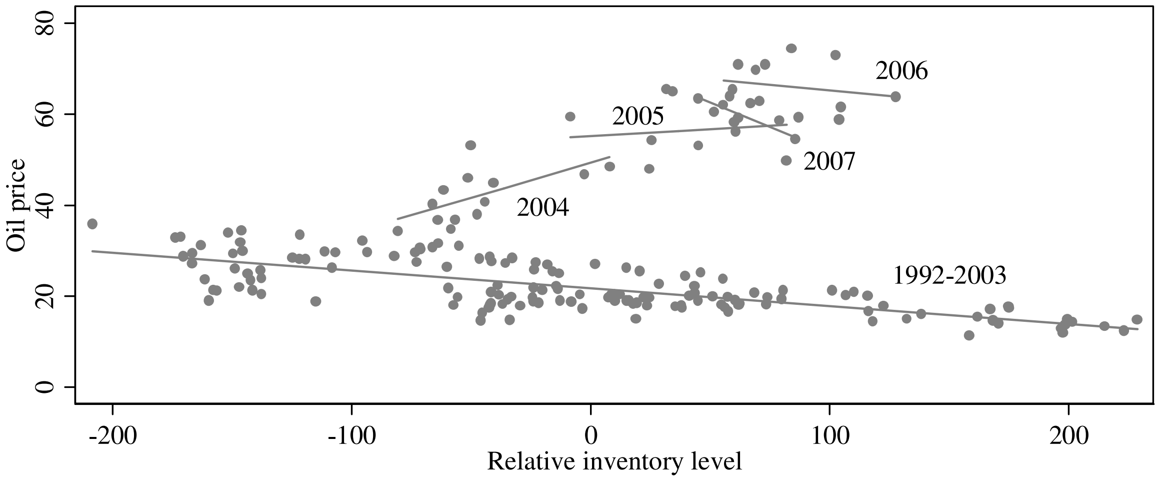

Note: The scatter plot shows the nominal WTI USD oil price (per stock tank barrel) versus the linear relative inventory level, measured as the difference (in million barrels) between current total OECD inventories and a presumed normal inventory level. The regression lines are from linear regression of the labelled dataset segments.

6 Concluding remarks

Short-term oil price used to be predictable using the so-called relative inventory level as an indicator of imbalance in the crude oil market (Ye, Zyren et al., 2002, 2005a). In this thesis, a previously successful relative inventory level model is exploited to investigate oil price and inventory dynamics and the failure of the relative inventory level model on updated data.

The thesis aims to answer two questions: If there is a break in the relative inventory model that can be explained by the crude modelling of the price-inventory relationship, or if the break, if any, is in the oil price process, unrelated to the price-inventory relationship.

The main finding is that the oil price process has shifted slowly between 2002 and 2007, from one regime with a low mean, low variance and high persistence, towards a new regime with a higher mean, higher variance and lower persistence. The relative inventory models fail because they are unable to capture this shift. A secondary finding is that the price-inventory relationship does not appear to be broken. However, the relationship is found to be somewhat misrepresented in the original model. The proposed quadratic log-log model, with constant elasticity of oil price to inventory and a nonlinear inventory-dependent correction to the elasticity, is found to remedy misspecifications in the original model.

The findings have some interesting consequences: Firstly, some of the under-prediction in relative inventory level models may be alleviated. These models assume a constant long-run oil price, thus under-prediction is an expected consequence if the long-run price is rising. By introducing a slow long-run price shift into the model, its short-term forecasting abilities may be restored. It should be emphasised, however, that the introduction of a shift in the model does not bring along any information about the causes of the shift. Secondly, the proposed specification allows for a more proper treatment of uncertainty, e.g. for computation of confidence bands. The third consequence is possibly the most exciting: The oil price shift bears the mean/variance/persistence marks of the kind of shifts that Videgaray-Caso (1998) found to come about every 11th year or so, before they reverted back after 4 or 5 years.

Some of the weaknesses in this study concern the use of data. The oil price data are nominal, and thus contain inflation components. Inventories are measured in levels, not days of demand, and thus contain some general economic growth. Other shortcomings are methodological: Findings are based on in-sample evaluation of fit, not out-of-sample forecasting ability. Inventory coefficients are not time-varying, weakening the conclusions about the stability of the price-inventory relationship. Some shortcomings may merit further study: Evaluating the forecasting properties of the proposed models is necessary, and relating the regime shift to fundamental causes would enable an understanding of the shift, not a mere characterisation.

References

Akaike, H. (1974). "New Look at Statistical-Model Identification." Ieee Transactions on Automatic Control **Ac19**(6): 716.

Amano, A. (1987). "A small forecasting model of the world oil market." Journal of Policy Modeling **9**(4): 615.

Bernard, J.-T., L. Khalaf and M. Kichian (2004). Structural change and forecasting long-run energy prices. Bank of Canada Working Paper Series. Ottawa, Bank of Canada.

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

D'Agostino, R. B., A. Belanger and R. B. J. D'Agostino (1990). "A Suggestion for Using Powerful and Informative Tests of Normality." American Statistician **44**(4): 316.

Davidson, R. and J. G. MacKinnon (1993). Estimation and inference in econometrics. New York, Oxford University Press.

Doornik, J. A. and H. Hansen (1994). A Practical Test for Univariate and Multivariate Normality. N. C. Discussion Paper, University of Oxford.

Doornik, J. A., D. F. Hendry and L. Timberlake Consultants (2006). Empirical econometric modelling - PcGive 11 : Volume I. London, Timberlake Consultants Ltd.

Engle, R. F. (1982). "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United-Kingdom Inflation." Econometrica **50**(4): 987.

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

Kaldor, N. (1939). "Speculation and Economic Stability." Review of Economic Studies **7**(1): 1.

Kaufmann, R. K. (1995). "A Model of the World Oil Market for Project Link - Integrating Economics, Geology and Politics." Economic Modelling **12**(2): 165.

Kaufmann, R. K., S. Dees, P. Karadeloglou and M. Sánchez (2004). "Does OPEC Matter? An Econometric Analysis of Oil Prices." Energy Journal **25**(4): 67.

Kwiatkowski, D., P. C. B. Phillips, P. Schmidt and Y. C. Shin (1992). "Testing the Null Hypothesis of Stationarity against the Alternative of a Unit-Root - How Sure Are We That Economic Time-Series Have a Unit-Root." Journal of Econometrics **54**(1-3): 159.

Lanne, M. and H. Lütkepohl (2002). "Unit root tests for time series with level shifts: a comparison of different proposals." Economics Letters **75**(1): 109.

Leybourne, S., P. Newbold and D. Vougas (1998). "Unit roots and smooth transitions." Journal of Time Series Analysis **19**(1): 83.

Lin, C.-F. J. and T. Teräsvirta (1994). "Testing the constancy of regression parameters against continuous structural change." Journal of Econometrics **62**(2): 211.

Merino, A. and A. Ortiz (2005). "Explaining the so-called "price premium" in oil markets." OPEC Review **29**(2): 133.

Pindyck, R. S. (1993). "The Present Value Model of Rational Commodity Pricing."

Pindyck, R. S. (1999). "The long-run evolution of energy prices." Energy Journal **20**(2): 1.

Radchenko, S. (2005). The Long-Run Forecasting of Energy Prices Using the Model of Shifting Trend. University of North Carolina at Charlotte Working Paper.

Tchilinguirian, H. (2006). "Market Prices – What do Oil Prices Say about Fundamentals." The 7th China Oil Traders' Conference (COTC-VII), Mandarin Garden Hotel, Nanjing, Jiangsu, China, 18-20 April 2006 Retrieved 11 November 2007, from http://www.iea.org/textbase/speech/2006/ht_ChinaOil.pdf.

Telser, L. G. (1958). "Futures Trading and the Storage of Cotton and Wheat." Journal of Political Economy. **66**(2): 233.

Teräsvirta, T. (1998). Modelling Economic Relationships with Smooth Transition Regressions. Handbook of Applied Economic Statistics. D. Giles and A. Ullah**:** 507.

Videgaray-Caso, L. (1998). The fiscal response to oil shocks**:** 121.

White, H. (1980). "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity." Econometrica **48**(4): 817.

Working, H. (1949). "The Theory of Price of Storage." The American Economic Review **39**(6): 1254.

Ye, M., J. Zyren and J. Shore (2002). "Forecasting Crude Oil Spot Price Using OECD Petroleum Inventory Levels." International Advances in Economic Research **8**(4): 324.

Ye, M., J. Zyren and J. Shore (2005). "A monthly crude oil spot price forecasting model using relative inventories." International Journal of Forecasting **21**(3): 491.

Ye, M., J. Zyren and J. Shore (2006a). "Forecasting short-run crude oil price using high- and low-inventory variables." Energy Policy **34**(17): 2736.

Ye, M., J. Zyren and J. Shore (2006b). "Recent Disconnects in the Crude Oil Price and Inventory Relationship." International Advances in Economic Research **12**(4): 556.

Ye, M., J. Zyren and J. Shore (2006c). "Short-Run Crude Oil Price and Surplus Production Capacity." International Advances in Economic Research **12**(3): 390.

Zamani, M. (2004). "An Econometrics Forecasting Model of Short Term Oil Spot Price." 6th IAAE European Conference, 1-3 September 2004 Retrieved 10 November 2007, from http://www.saee.ch/saee2004/Zamani_Mehrzad.pdf.

A Model diagnostics

Table 5 Residual and model diagnostics

	1992-2003		1992-2007			
	BAS	**QLL**	**BAS**	**QLL**	**STR1**	**STR2**
AR	0.415	0.497	0.289	0.710	0.666	0.504
ARCH	0.045 *	0.096	0.000 **	0.276	0.294	0.329
Normality	0.789	0.867	0.000 **	0.617	0.803	0.859
Hetero	0.000 **	0.435	0.000 **	0.403	0.445	0.564
Hetero-X	0.007 **		0.000 **	0.828		
RESET	0.813	0.823	0.069	0.289		
RRESET	0.835		0.145			
R^2	0.924	0.934	0.977	0.978	0.980	0.980
RSS	290.343	0.531	931.388	0.797	0.717	0.705
AIC	3.756	-2.457	4.589	-2.407	-2.497	-2.507

Note: BAS is the baseline relative inventory level model, and QLL is the competing quadratic log-log respecification. STR1 and STR2 are smooth transition versions of the QLL. * and ** are used to denote significance at the 10% and 5% level, respectively, even if the test results are reported in p-values. AR denotes the Lagrange-multiplier test of (7th order) no autocorrelation, described by Doornik, Hendry et al. (2006:257). ARCH refers to the Engle (1982) test of no autoregressive conditional heteroskedasticity described by Doornik, Hendry et al. (2006:259). Normality denotes the Doornik and Hansen (1994) test of normality described by (Doornik, Hendry et al. (2006:225). Hetero and Hetero-X refers to the White (1980) test of homoskedastic errors, where the first includes levels and squares, and the second also includes cross products of the explanatory variables. RESET refers to the results from the Ramsey's regression error specification test, and RRESET a heteroskedasticity-robust version in Stata 9.2. R^2 denotes the coefficient of determination, RSS the residual sum of squares and AIC the Akaike information criterion.

B Coefficient estimates

Table 6 Coefficient estimates

	1992-2003		**1992-2007**			
	BAS	**QLL**	**BAS**	**QLL**	**STR1**	**STR2**
Intercept and autoregressive coefficients						
a0	2.032 *	0.288 *	0.206	-0.002	0.425 **	0.352 **
a1	0.907 **	0.903 **	1.001 **	1.001 **	0.856 **	0.881 **
Relative inventory level coefficients						
b0	-0.008 **	-1.180 *	-0.011 *	-1.118 *	-1.275 *	-1.290 *
b1	-0.006	-1.063	-0.006	-1.014	-1.402 *	-1.243
b2	0.008	0.966	0.017 **	1.666 *	1.423 *	1.365
b3	-0.001	-0.464	-0.001	-0.341	-0.395	-0.336
b4	0.011 *	1.793 *	0.005	1.309	1.366	1.279
b5	-0.007	-0.866	-0.006	-0.656	-0.755	-0.696
Relative inventory level coefficients, quadratic part						
h0		1.064		-4.543	-4.660	-3.695
h1		5.461		11.546	7.660	9.539
h2		-3.830		-4.689	-3.920	-4.164
h3		-15.968		-14.783	-14.007	-14.294
h4		-6.480		-4.951	-2.875	-4.415
h5		33.043 **		25.464 *	29.384 **	29.232 **
Nonlinear intercept and autoregressive coefficients						
a10					0.213 **	1.186
a11						-0.246
Transfer function coefficients						
gamma					0.109 *	0.147 **
c					147.931 **	153.484 **

Note: BAS is the baseline relative inventory level model, and QLL is the competing quadratic log-log respecification. STR1 and STR2 are smooth transition versions of the QLL. * and ** denote significance at the 10% and 5% level, respectively. Corresponding p-values are found on the next page.

Table 7 P-values for coefficient estimates

	1992-2003		1992-2007			
	BAS	**QLL**	**BAS**	**QLL**	**STR1**	**STR2**
P-values for intercept and autoregressive coefficients						
a0	0.018 *	0.015 *	0.705	0.955	0.001 **	0.004 **
a1	0.000 **	0.000 **	0.000 **	0.000 **	0.000 **	0.000 **
P-values for relative inventory level coefficients						
b0	0.009 **	0.045 *	0.014 *	0.045 *	0.018 *	0.018 *
b1	0.153	0.175	0.348	0.170	0.048 *	0.079
b2	0.117	0.225	0.007 **	0.025 *	0.045 *	0.054
b3	0.739	0.570	0.929	0.645	0.575	0.634
b4	0.014 *	0.030 *	0.486	0.077	0.053	0.071
b5	0.073	0.151	0.238	0.242	0.160	0.196
P-values for the quadratic relative inventory level coefficients						
h0		0.928		0.702	0.680	0.745
h1		0.697		0.416	0.570	0.481
h2		0.791		0.743	0.772	0.759
h3		0.298		0.299	0.300	0.291
h4		0.671		0.727	0.831	0.744
h5		0.009 **		0.031 *	0.010 **	0.010 **
P-values for the nonlinear intercept and autoregressive coefficient						
a10					0.005 **	0.132
a11						0.196
P-values for the transfer function coeffients						
gamma					0.011 *	0.001 **
c					0.000 **	0.000 **

Note: BAS is the baseline relative inventory level model, and QLL is the competing quadratic log-log respecification. STR1 and STR2 are smooth transition versions of the QLL. * and ** denote significance at the 10% and 5% level, respectively. Coefficient estimates are found on the previous page.

C Estimates of long-run oil price and long-run multipliers

Table 8 Estimates of the long-run oil price and multipliers

	1992-2003		1992-2007			
	BAS	**QLL**	**BAS**	**QLL**	**STR1**	**STR2**
RIN or rin-multiplier:	-0.04 **	-8.37 **	2.40	136.53	-7.22 **	-7.75 **
rin2-multiplier:		136.54		-7118.4	80.59	102.60
Long-run oil price:	21.89 **	19.19 **	-286.76	8.79	19.28 **	19.28 **
Long-run oil price at end of shift:					84.57 **	67.92 **

Note: BAS is the baseline relative inventory level model, and QLL is the competing quadratic log-log respecification. STR1 and STR2 are smooth transition versions of the QLL. * and ** denote significance at the 10% and 5% level, respectively. Corresponding test statistics and p-values are found in the table below.

Table 9 Test statistics and p-values for the estimates of long-run price and multipliers

	1992-2003		1992-2007			
	BAS	**QLL**	**BAS**	**QLL**	**STR1**	**STR2**
Significance of estimated relative inventory level multiplier						
Prob>F	0.01 **	0.00 **	0.87	0.88	0.00 **	0.00 **
F	7.63	10.84	0.02	0.02	14.13	11.07
Significance of estimated quadratic relative inventory level multiplier						
Prob>F		0.17		0.87	0.22	0.20
F		1.87		0.03	1.53	1.68
Significance of estimated long-run price						
Prob>F	0.00 **	0.00 **	0.40	0.70	0.00 **	0.00 **
F	225.12	1001.54	0.72	0.15	2801.10	1934.06
Significance of estimated long-run price at end of shift						
Prob>F					0.00 **	0.00 **
F					927.93	5175.23

Note: * and ** denote significance at the 10% and 5% level, respectively. F refers to the result from a Stata 9.2 Wald-type F-test. BAS is the baseline relative inventory level model, and QLL is the competing quadratic log-log respecification. STR1 and STR2 are smooth transition versions of the QLL.