# Demand Effects and Speculation in Oil Markets: Theory and Evidence

Eyal Dvir*
Boston College

Kenneth Rogoff
Harvard University

First Draft: September 10, 2012
This Draft: November 11, 2013

**Abstract**

We present evidence showing the existence of stable cointegrating vectors connecting four important variables in the U.S. and global oil markets: oil production, stocks of crude oil, the real price of oil, and broad measures of income. Our data are monthly, and go back to the 1930s, split into sub-samples which correspond to periods before and after the 1973 crisis. We further show that the cointegrating vectors found in the data accord well with an extended commodity storage model which allows for demand growth dynamics and for supply regimes. Specifically, inventories and price move in opposite directions when supply is flexible, but the relationship reverses so that they comove when supply is inflexible.

## 1 Introduction

The role of speculation in driving the price of crude oil has been the object of renewed interest recently. The decades-old debate, between those who argue that market developments can be directly attributed to changes in fundamentals and those who believe that speculators are creating price volatility,

*Corresponding author. *Email*: eyal.dvir@bc.edu

is showing no signs of abating[1]. In this paper we put forward the argument that a simple model with four variables - inventories, production, income, and price - can be useful in capturing important long-run features of the market for crude oil, and in particular elucidating the seemingly unstable relationship between inventories and price. We estimate the long-run equilibrium relationships predicted by the model separately for the periods 1931 - 1972 and 1975 - 2011and show that our model's long-run predictions are borne out by the data. We show evidence in support of our model's prediction that a stable equilibrium relationship exists among these four variables, that these equilibrium relationships are stable before and after the 1973 crisis, and that they broadly comport in sign to the model's predictions. Specifically, we find, as predicted by the model, that post - 1973 there is a stable **positive** long-run relationship between inventories of crude oil and the real price of crude, whereas before 1973 that relationship was **negative.**

In a previous paper, Dvir and Rogoff (2009), we argue that the real price of oil has gone through three distinct periods. First, from 1861 to about 1878 (a period not covered in the current paper), the price of oil was generally high (in real terms), and was moreover highly persistent and volatile. Then came a much less volatile period, between 1878 and 1973, in which prices were also generally lower and not at all persistent. This long period can be further divided into two sub-periods: before and after 1933, where price volatility is significantly lower after 1933 compared with the years 1878-1933. Finally, from 1973 onwards, there is a recurrence of high persistence and volatility accompanied again by higher prices. In that paper, we offered a narrative, based on our reading of the historical events, for the recurrence of high price persistence in the two end-periods mentioned, 1861-1878 and 1973-2010. We argued that in these periods two forces coincided: first, demand (governed by income) was high and very persistent, i.e. it was governed by growth shocks. Second, access to supply was restricted by agents who had the capability and incentive to do so. In particular, whereas before the crisis oil supply was easily accessible, and indeed was actively managed by regulators, after the crisis there was no easy way of increasing oil production. Post 1973, all excess

[1]See Singleton (2013) for a recent review. Recent theoretical contributions include Basak and Pavlova (2012), where institutional investors can cause commodity prices to rise and become more volatile, and Sockin and Xiong (2012), where feedback effects from futures prices affect commodity spot prices through an information channel. However, in a survey of recent empirical work Fattouh, Kilian, and Mahadeva (2012) find scant evidence for the effect of speculation on prices. See also Hamilton (2009).

capacity existed in the Middle East, where producers were more interested in maintaining high prices than in accommodating demand increases.

In that paper we also presented our model, which is an extension of the canonical commodity storage model à la Deaton and Laroque (1992, 1996). Our model introduces income growth dynamics to that framework, and in particular, can accommodate both $I(0)$ and $I(1)$ income processes. In that paper we focused on predictions of the model with regards to demand and supply shocks. In the current paper we emphasize the related but distinct **long-run** relationships which are also predicted by the model. This focus is important when attempting to account for actual market behavior. Perhaps due to the existence of unit roots in the various series - inventories, production, and price - identifying long-run relationships among them has not been a priority in the literature on the market for crude oil. However predictable long-run relationships between these variables are a hallmark of the commodity storage model. Therefore our contribution here is two-fold: first, we show that these long-run relationships do exist in the data, as predicted by the model. Second, we show that the relationships we identify in the data are (mostly) sign - consistent with our version of the commodity storage model.

Introducing income growth dynamics to the classical model adds considerably to its predictive capacity, partly by changing some of the classical model's predictions. The model can now predict how inventories and price behave when income rises or falls, conditioning on production behavior. The relationship between inventories and price is therefore no longer simple, becoming a function of production behavior. In particular, inventories may rise or fall with income. Specifically, in periods when production is flexible, i.e. when a rise in income, which raises demand, is predicted to result in a commensurate rise in production, inventories should fall, since the effects of high demand should dissipate quickly. This should help mitigate any rise in price associated with the surge in demand so that inventories and price should exhibit a negative relationship. Conversely, in periods when production is inflexible, i.e. when a rise in income is not predicted to raise production significantly, inventories should rise. This would actually enhance any rise in prices associated with high demand, so that inventories and price should exhibit a positive long run relationship. Note importantly that causality runs in both directions in a commodity storage model: in the flexible supply case a rise in prices will cause inventories to drop, thereby releasing more oil to the market and exerting a downward effect on prices. In the restricted supply case a rise in price leads to a surge in inventories, pulling oil off the market and leading to further price rises. The relationships we identify should therefore be understood as long-run equilibrium conditions, and not given a causality interpretation. It is therefore important to specify the long-run relationship among all four variables - inventories, production, income, and price - at the same time.

We use U.S. monthly data on crude oil stocks, production, and prices, as well as on the size of the economy, going back to 1/1931. Demand is assumed to rise with income, or the size of the economy. Our chosen variable is monthly U.S. industrial production, which is available from the Federal Reserve. Petroleum is used, directly or indirectly, by every sector of the economy; therefore demand for oil is commensurate with overall economic activity. The measure we use is the most complete measure of economic activity in the U.S. at a monthly frequency which is available for the entire period. Our measures of U.S. crude oil stocks and production are from the Energy Information Administration. Oil price in our data set is a composite series of monthly prices quoted in Texas and Oklahoma. For the post - 1973 oil market, U.S. data alone are insufficient, the U.S. market having become dependent on global conditions. We therefore use the broadest measures of the global oil market available to us. We use the OECD +6 monthly GDP data (available from 1970/1) as our income variable. This is the broadest measure of global income available at a monthly frequency, including all developed economies as well as the BRIC countries, Indonesia and South Africa. It is constructed by the OECD in constant 2005 PPP prices. World oil production data is available from the U.S. dept. of Energy starting in 1973/1. OECD petroleum stocks data is available at a monthly frequency from 1988/1, with quarterly and annual data going back to 1973, which we then use to impute stocks at monthly frequency, following Kilian and Murphy (2013). We use the Dept. of Energy's refiners' acquisition cost price for imported oil, also deflated by the U.S. CPI as our measure of the real price of oil. This series is thought to be less subject to regulatory pressure relative to the West Texas Intermediate during the 1970s and early 1980s.

A common feature of these series is that they exhibit unit roots, at least for some sub-periods. The most pressing empirical question then becomes: Are these series cointegrated? For our commodity storage model to be a reasonable account of the market for crude oil, these series must, in the long run, co-move in a predictable way. This is because our model has a stable rational expectations equilibrium, with well defined equilibrium relationships among the constituent variables. If there is little evidence of the existence of a cointegrating vector, then the commodity storage model cannot be the simplest way of describing the market. Our first task then is to establish that stationary cointegrating vectors do indeed exist. Since our model leads us to expect that the market should behave differently before and after the crisis of 1973, we believe it is necessary to split the sample around that time. In Section 3 we show that using standard Johanssen tests we can establish quite clearly that, yes, these series are cointegrated.

Our second task is to estimate these cointegrating vectors to see whether the variables interact in the way predicted by the model. We run vector error correction regressions and arrive at statistically and economically significant long-run relationships, in two separate sub-samples: 1931/1 - 1973/12, and 1975/1 - 2011/12. In the earlier sub-sample supply was quite flexible, since U.S. oil production was well below its capacity. We would therefore expect inventories to increase with production, and to decrease with demand and with price. That is indeed what we find. In the later sub-sample, supply was inflexible, since U.S. production was at its limit (wells were operating at capacity) and Western firms had limited access to additional oil, excess capacity having been effectively nationalized by Middle East governments. We would expect then to see inventories **increase** with demand and with price. Our estimates of the coefficients of the cointegrating vector accord well, but not perfectly, with the model's expectations[2]. Our reliance on aggregate data at the OECD and global level may be the reason. Experimenting with different break points and different specifications in terms of lag length, alternative variables, etc. leads to different magnitudes for the coefficients, but importantly their sign remains stable, and they remain statistically significant.

Recent contributions have also looked at the relationship between oil supply, demand, price, and inventories. Kilian and Murphy (2013) use a structural VAR framework to identify separate types of shocks to the global oil market. In particular they are able to use inventory data to identify shocks to speculative demand and estimate their relative importance to oil price behavior. Kilian and Lee (2013) extend this work to include new proxies for global oil inventories, and use these to re-examine the relative importance of speculative demand shocks. Our paper is complementary in many ways to these contributions, in that we focus on the existence and characterization of

[2]Our estimates of the long-run relationship show inventories increasing with price, as expected, but decreasing with income.

long-run relationships among the variables, but do not attempt to identify separate shocks to the system. Knittel and Pindyck (2013) also develop a model of oil storage and use it to examine recent claims of speculation in the oil market. Their analysis covers the U.S. market between 1998 - 2012, whereas we focus on the long term, with monthly data going back to 1931/1.

The paper proceeds as follows: Section 2 presents our model and characterizes the rational expectations equilibrium. Section 3 describes the model's predictions for each period in detail and presents our empirical findings. Section 4 concludes.

# 2 An Extended Commodity Storage Model

Our model is an extension of the classic commodity storage framework. Chambers and Bailey (1996) and Deaton and Laroque (1996) extend the model to allow for auto-regressive shocks. We extend it further to explicitly incorporate demand, and to allow for growth shocks[3].

## 2.1 Availability and Storage

Time is discrete, indexed by $t$. The market for oil consists of consumers, producers, and risk neutral arbitrageurs. The latter have at their disposal a costly storage technology which may be used to transfer any positive amount of oil from period $t-1$ to period $t$. Storage technology is limited by a non-negativity constraint, i.e. the amount stored at any period cannot drop below zero. This implies that inter-temporal arbitrage, although potentially profitable, cannot always be achieved. In these cases the market is "stocked out". Let $A_t$ denote *oil availability*, the amount of oil that can potentially be consumed at time $t$. This amount has already been extracted from the ground, either in period $t$ or at some point in the past, and has not been consumed before period $t$. It is given by

$$A_t = X_{t-1} + Z_t, \tag{1}$$

[3]This is essentially the same model we presented in Dvir and Rogoff (2009). We include it here for completeness and also because here we emphasize its predictions of stable relationships between the constituent series. Empirical work on commodity futures curves has shown the need for a permanent shock component: see Routledge, Seppi, and Spatt (2000), and Schwartz and Smith (2000).

where $X_{t-1}$ denotes the stock of oil transferred from period $t-1$ to $t$, and $Z_t$ denotes the amount of oil that is produced at time $t$. For simplicity, we assume that no oil is lost due to storage[4]. Decisions concerning both variables - how much to store, how much to produce - are assumed to have been made before period $t$ began. In period $t$ agents decide how to divide $A_t$ between current consumption $Q_t$ and future consumption, so that demand - the sum of current consumption and the amount stored for the future - must always equal current availability:

$$A_t = Q_t + X_t. \tag{2}$$

## 2.2 Demand for Oil

Let $P_t$ denote the price of crude oil, and let $Y_t$ be a demand parameter, which should be thought of as capturing the economy's derived demand for energy stemming from industrial, residential, and transportation uses. For simplicity, we will refer to $Y_t$ as income. We can then write an inverse demand function for oil as follows:

$$P_t = P(Q_t, Y_t), \tag{3}$$

where inverse demand is decreasing in its first argument, and increasing in its second. This constitutes a mild departure from the canonical model, where demand for the commodity is a function of its price alone. This departure is a natural one to make, however, in the context of oil, as oil consumption and income are very highly correlated. We posit an inverse demand function in which only the *ratio* of consumption to income matters, i.e. inverse demand is homogeneous of degree zero:

$$P_t = P(Q_t, Y_t) = P(\frac{Q_t}{Y_t}, 1) = p(q_t), \tag{4}$$

where lowercase letters denote variables normalized by $Y_t$. We will refer to normalized variables as "effective" amounts, in the sense that a growing economy leads to higher energy needs, spreading any given amount of oil more thinly.

We will use a CES inverse demand function:

$$P_t = q_t^{-\gamma} = (a_t - x_t)^{-\gamma}, \tag{5}$$

[4]Alternatively, we could have specified storage costs by a given loss percentage, as in Deaton and Laroque (1996).

where $\gamma > 1$ is the inverse elasticity of demand, and $a_t$, $x_t$ denote effective availability and storage in period $t$, respectively. It is natural to assume that the effective demand for oil is inelastic with respect to price. As equation (5) makes clear, for a given supply of oil, price is a function of the competing demands of current and future consumption. If the desire to consume more in the future grows (driven by expectations of future conditions), more oil is stored rather than consumed today, resulting in a price rise today even though supply has not changed.

Let $\overline{Y}_t$ denote trend income, i.e. the level of income that would prevail at time $t$ in a world without income shocks. $\overline{Y}_t$, which we think of as a measure of current production technology, is assumed to increase over time at a constant rate $\overline{\mu} > 0$. We now consider two alternative stochastic processes for $Y_t$: one where income moves around a deterministic trend, and another where the trend itself is stochastic. The former is a simple AR(1) process, analogous to the stochastic process that Deaton and Laroque (1996) consider for supply. Under this assumption we have:

$$\frac{Y_{t+1}}{\overline{Y}_{t+1}} = \left(\frac{Y_t}{\overline{Y}_t}\right)^{\rho} e^{\varepsilon_{t+1}}, \tag{6}$$

where $\rho \in (0,1)$ and $\varepsilon_{t+1} \sim N(0, \sigma_{\varepsilon}^2)$ is an iid shock. We think of this case as more closely relevant to income shocks in developed economies, where the economy exhibits business cycles around a stable trend. In the latter case, we assume instead:

$$Y_{t+1} = e^{\mu_{t+1}} Y_t, \tag{7}$$

such that

$$\mu_{t+1} = (1-\phi)\overline{\mu} + \phi \mu_t + \upsilon_{t+1}, \tag{8}$$

where $\phi \in (0,1)$ and $\upsilon_{t+1} \sim N(0, \sigma_{\upsilon}^2)$ is an iid shock. Dividing both sides of (7) by $\overline{Y}_{t+1}$ we get:

$$\frac{Y_{t+1}}{\overline{Y}_{t+1}} = e^{\mu_{t+1} - \overline{\mu}} \frac{Y_t}{\overline{Y}_t}. \tag{9}$$

We think of this case as more relevant to income shocks in some developing countries, in particular quickly industrializing economies where very high growth rates can be quite persistent.

## 2.3 Supply of Oil

In the canonical commodity storage model, supply $Z_t$ varies according to some stochastic process $\psi_t$ around a predetermined mean $\widetilde{Z}_t$, and it is this variability in supply that creates an incentive for inter-temporal smoothing by the large pool of risk neutral arbitrageurs. As the literature has long recognized, demand and supply shocks in the canonical model are isomorphic: one can think of a negative realization of $\psi_t$ as representing an especially cold winter (demand) or a breakdown in a major pipe (supply). For this reason, since we model demand shocks explicitly, it would be redundant to model supply shocks separately.

We do model supply *choices*, however. In particular, we assume that either of the following two regimes holds: a regime where oil supply does not react at all to demand shocks due to capacity constraints (such as railroad infrastructure or number of operational wells), and a regime in which oil supply fully accommodates any shock to demand (for example, when potential production is much higher than current production). We think of the former regime as describing supply behavior when access to excess supply sources is restricted, so that suppliers are constrained to produce at their installed capacity[5]. Under the latter regime, suppliers seek to stabilize prices by varying quantities as needed. We think of this regime as representing either perfectly competitive supply, where producers will offer any amount at a given price, or else the effect of purposeful government intervention, seeking to control market prices by adjusting supply.

Formally, in the former regime we assume that supply grows at the trend income rate $\overline{\mu}$, so that

$$Z_{t+1} = \widetilde{Z}\overline{Y}_t, \tag{10}$$

where $\widetilde{Z}$ is a supply parameter. Next period's oil supply depends then on current technology, since overall technological progress, which drives global GDP growth, applies to the oil extraction and exploration sectors as well, and therefore determines overall capacity.

This assumption deserves some comment. The total amount of oil existing in the earth's crust is finite. However technological progress is key to exploiting an increasing fraction of it over time. The global ratio of oil

[5]Naturally, capacity constraints can be relaxed in the medium run. However, as long as capacity does not fully accommodate all demand shocks, dynamic behavior will be qualitatively similar to the case where it does not react at all. A similar point has been made by Williams and Wright (1991).

production to known oil reserves is slightly less than 2.5% , and has been quite steady at that level since 1985 (BP Statistical Review), even though global production has increased by about 39% from 1985 to 2010. The world economy is no closer to running out of oil now than it was in 1985 due to the rate at which new reserves are discovered and known reserves become exploitable due to better technology. This is the context which drives our modeling choice, since it suggests that a stationary equilibrium relationship among the important variables might exist.

Note that in this regime oil supply depends on the technology driving income growth, but not on income growth itself. Therefore shocks to demand will drive a wedge between supply and demand, causing a shift in equilibrium price. In contrast, under the alternative supply regime oil suppliers will accommodate all income shocks, i.e. oil supply will be perfectly elastic. Next period's supply then will also depend on current income level (and growth rate if appropriate). Supply is then given by:

$$Z_{t+1} = \widetilde{Z}\overline{Y}_t \left( \frac{Y_t}{\overline{Y}_t} \right)^{\rho}, \tag{11}$$

for the AR(1) case or by:

$$Z_{t+1} = \widetilde{Z} e^{(1-\phi)\overline{\mu}+\phi\mu_t} Y_t, \tag{12}$$

for the stochastic trend case.

## 2.4 Storage of Oil

The defining characteristic of the canonical model is the availability of storage technology, i.e. the ability to perform inter-temporal arbitrage. Here we follow the literature closely. We assume free entry into the storage sector as well as risk neutrality, implying that the actions of arbitrageurs will raise or lower the current price until it is at a level which renders the strategy unprofitable in expectation, unless that would require holding negative stocks, at which case inter-temporal arbitrage will be incomplete. In all other cases, i.e. when equilibrium at time $t$ is fully optimal, the price of oil must obey the following arbitrage condition:

$$P_t = \beta E_t[P_{t+1}] - C, \tag{13}$$

where $\beta = 1/(1+r)$ is the discount factor, and $r > 0$ is the exogenously given interest rate. The parameter $C > 0$ denotes the per barrel cost of storage. Equilibrium price $P_t$ must be such that there is no incentive to increase or decrease $X_t$, the amount stored[6].

Note that storage involves an inter-temporal choice, whereas the production decision does not. This is worth mentioning since models of the oil market which emphasize non-renewability imply that producers must decide whether to extract a barrel of oil today or tomorrow. That is not the case here: in our model, as in the canonical storage model, production decisions are made based on current and expected market conditions. Hence the real interest rate enters the storage equation, but does not enter the production equations.

## 2.5 The Rational Expectations Equilibrium

The canonical commodity storage model is a rational expectations model with one state variable - availability of oil $A_t$ - and one choice variables - storage of oil $X_t$. A solution of the model - the rational expectations equilibrium - consists of a storage rule, which specifies the level of storage for every possible value of the state variable. Determination of price and consumption follows immediately from this rule. In our extended version of the model the rule retains its salient characteristics, well known from the literature (see below). However in the extended version, as in the AR(1) case considered by Chambers and Bailey (1996), storage is also the function of one (or two) exogenous variables, depending on assumptions regarding the income process. Relative income $Y_t/\overline{Y}_t$ - how far above or below its mean is the current level of income - serves as the second state variable of the model when we assume that income follows a stable trend. For the case where income is subject to growth shocks, we need a third state variable: the current growth rate of income, denoted by $\mu_t$.

In order to solve the model we express all quantity variables in their normalized forms. The model can be then be summarized by two (or three)

[6]The inter-temporal price condition (13) does not hold in the case of a stock-out, i.e. the case where $X_t = 0$ because the storage non-negativity constraint is binding; every barrel of extracted oil is being used for consumption. As a result, current price is above its unconstrained level:

$$P_t > \beta E_t[P_{t+1}] - C. \tag{14}$$

transition functions which govern the state variables, and one response equation which determines storage, the decision variable. We therefore arrive at a $2 \times 2$ framework: two alternatives for the demand process and two for the supply regime. Agents in the model observe all the state variables every period, and decide on storage accordingly, taking into consideration expectations regarding the next period's price, and implicitly producers' behavior.

The transition functions for the stable trend case are:

$$a_{t+1} = \frac{x_t + z_{t+1}}{(Y_t/\overline{Y}_t)^{\rho-1} e^{\overline{\mu}+\varepsilon_{t+1}}}, \tag{15}$$

$$\frac{Y_{t+1}}{\overline{Y}_{t+1}} = \left(\frac{Y_t}{\overline{Y}_t}\right)^{\rho} e^{\varepsilon_{t+1}}, \tag{16}$$

where equation (15) is derived by normalizing equation (1) by $Y_{t+1}$ and using (6). Effective supply $z_{t+1}$ is arrived at by dividing either equation (10) or (11) through by $Y_t$, depending on the supply regime in effect.

For the stochastic trend case, there are three transition functions:

$$a_{t+1} = (x_t + z_{t+1})/e^{\mu_{t+1}}, \tag{17}$$

$$\frac{Y_{t+1}}{\overline{Y}_{t+1}} = e^{\mu_{t+1}-\overline{\mu}} \frac{Y_t}{\overline{Y}_t}, \tag{18}$$

$$\mu_{t+1} = (1-\varphi)\overline{\mu} + \varphi\mu_t + \upsilon_t, \tag{19}$$

where the transition function (17) is derived again by normalizing equation (1) by $Y_{t+1}$ , now using (7) instead. Here as well, the supply regime in effect determines how we arrive at $z_{t+1}$: dividing either equation (10) or (12), as appropriate, by $Y_t$.

The response equation for both cases is:

$$(a_t - x_t)^{-\gamma} = \beta E_t[P_{t+1}] - C. \tag{20}$$

Note importantly that equation (20), which determines optimal storage, holds only when the state variables are such that the optimal storage is non-negative. If the state variables dictate negative storage, this response condition breaks down and we have simply $P_t = a_t^{-\gamma}$.

Commodity storage models generally cannot be solved analytically even in their most simple form (Newbury and Stiglitz, 1981, Williams and Wright, 1991). We therefore follow the literature since Gustafson's (1958) original contribution and proceed to solve the model numerically[7]. It turns out from our numerical solutions that the storage rules which result from any of our four sets of assumptions regarding supply and demand are very similar. All four of these rules are essentially identical in form to the rule that results from the canonical model. The difference is that in our extended model these rules hold for the *normalized* variables instead of the original quantities. In other words, effective storage has a relationship with effective availability in the extended model, under both sets of assumptions regarding demand, and both supply regimes, that is qualitatively similar to the relationship between actual storage and actual availability in the canonical model. As far as we know this is a new result as well.

Figure 1 shows a typical storage rule as well as the corresponding equilibrium price, both as functions of effective oil availability $a_t$ (on the horizontal axis)[8]. Both curves are qualitatively similar regardless of our assumption on income's stochastic process or the supply regime. Together these curves signify the location of equilibrium at every possible level of effective availability. As in the canonical model, storage is a positive function of availability beyond a certain point (below this point the non-negativity constraint is binding), whereas price is a negative function of availability, the curve becoming less steep once storage is positive.

Figure 2 exhibits the novel results of our model. In its two panels we show the effect of a rise in relative income $Y_t/\overline{Y}_t$ (horizontal axis) on effective storage $x_t$. In the upper panel we show the rational expectations equilibrium where supply is flexible and demand grows around a deterministic trend. In the lower panel we show the RE equilibrium where supply is restricted and demand exhibits a stochastic trend. Our model predicts that in the former case (flexible supply, deterministic trend), a rise in relative income will be accompanied by a reduction in inventories. The reason is as follows: as income rises above its long-run trend, production will increase to accommodate the higher demand, and also income will be expected to revert back to its trend.

[7]See Dvir and Rogoff (2009), appendix B for details of the solution method.

[8]Certain assumptions need to be made regarding the model's parameters in order to solve the model numerically. Demand elasticity $-1/\gamma$ is set at -0.5. The cost of storage $C$ is 0.02 per barrel. The discount factor $\beta$ is set at 0.97. The trend income growth rate $\overline{\mu}$ is set at 0.02, the income persistence parameter $\rho$ is set at 0.6, and the growth persistence parameter $\phi$ is set at 0.45. Effective supply capacity $\widetilde{Z}$ is set at $e^{\overline{\mu}}$. Lastly, the income shock's standard deviation $\sigma$ is set at 0.1, and the growth shock's standard deviation $\upsilon$ is set at 0.02.

Both forces imply that any rise in price will be short-lived, and therefore rational agents will sell some of their inventories in order to profit from the relatively higher price. On the other hand, when supply is restricted and demand exhibits a unit root (lower panel), a rise in income is not predicted to induce a rise in production or any mean reversion. For this reason rising prices due to rising demand can be seen as a process which is likely to continue, and rational agents will accumulate inventories as a result. Note that in both panels we also show that higher availability (i.e. higher production for any given relative income) will in both cases be associated with higher inventories, as already seen in Figure 1.

# 3 Stocks, Production, Demand, and Price: Empirical Links Over Time

We have monthly production and stocks data from the U.S. Dept. of Energy going back to 1920/1, covering the entire U.S. Our oil price series, reflecting prices in Oklahoma (what became in the 1980's the West Texas Intermediate price), is constructed from Commodity Research Bureau (1940, 1950, 1960), for 1931/1 - 1958/12, and from the IMF's International Financial Statistics database for 1959/1 - 2011/12. We deflate this series by the U.S. CPI (Bureau of Labor Statistics), to arrive at the real price of oil. An alternative oil price for the period from 1974/1 is the EIA's refiners' acquisition cost price for imported oil, also deflated by the U.S. CPI. This series is thought to be less subject to regulatory pressure relative to the WTI during the 1970s and early 1980s. For industrial production, we use the Federal Reserve Board's Industrial Production series, which starts in 1919/1. We utilize the most inclusive index available. For the global oil market, we use the OECD +6 monthly GDP data (available from 1970/1) as our income variable. This is the broadest measure of global income available at monthly frequency, including all developed economies as well as the BRIC countries, Indonesia and South Africa. World oil production data is available from the U.S. dept. of Energy starting in 1973/1. OECD petroleum stocks data is available at a monthly frequency from 1988/1, with quarterly and annual data going back to 1973. Monthly stocks were imputed back to 1973/1 by scaling U.S. crude oil inventory data by the ratio of OECD petroleum inventories over U.S. petroleum inventories, following Kilian and Murphy (2013). This is the broadest measure available to us[9]. It should be noted that OECD and world data are, by necessity, less reliable than U.S. data. OECD stocks include government as well as private storage, a relevant fact if governments do not respond to the market in the same way that private agents would. World oil production data may be biased by intentional misreporting, especially in OPEC countries. Finally, the challenges of constructing a broad measure of GDP across many countries and over long periods are well known[10].

We split the sample in two: the first sub - period (Period I), when supply was flexible and the U.S. oil market was largely self - sufficient, starts at 1931/1 (when our price series begins) up to and including 1972/12. This is not an arbitrary choice, as explained below. It turns out not to make much difference to the results, however. The second period (Period II), when supply was restricted, and the oil market became more globalized, is from 1975/1 (so as to allow for a maximum lag order of 24 months) up to and including 2011/12. Our variables are: crude oil production in million barrels (in the U.S. for Period I, in the world for Period II) , crude oil inventories in million barrels (in the U.S. for Period I, in the OECD for Period II), an index of production (industrial production in the U.S. for Period I, GDP in the OECD+6 for Period II), and the real price of oil. Figures 3 and 4 present the four series for the sub-periods 1920/1 - 1972/12 and 1973/1 - 2011/12. We will construct two series of availability, one of the state variables of the model, defined as the sum of available inventories in the beginning of the month and of the relevant oil production stream in the current month.

These time series seem to be non-stationary. In what follows, we first derive the model's predictions taking this feature of the data into account. We show that our model, while able to accommodate stochastic trends, nevertheless posits a stationary relationship between the variables of interest. Crucially, the model posits a different relationship across periods: storage should decrease with income and price in Period I, but increase with income and price in Period II. We then present our empirical results, where we show that the relevant series are cointegrated, as predicted by our model, and further that the signs of the cointegrating equation mostly correspond to the

[9]Kilian and Lee (2013) explore other measures of global oil inventories.

[10]A previous version of the paper included an analysis of the U.S. oil market post - 1973, where data reliability is less of an issue. Both anonymous referees pointed out, correctly, that the global nature of the oil market after 1973 makes an analysis of the U.S. by itself misleading. In this version we provide results using the broadest measures we could find. The U.S. results are available upon request.

model's predictions.

## 3.1 Empirical Predictions of the Model: Period I

Period I begins with the discovery of large amounts of oil in East Texas. This was a transformative moment in the history of the oil industry since, in the space of a few years, it became clear for the first time that oil was indeed abundant. Previously, oil fields in other parts of the U.S. (e.g. Pennsylvania, Indiana, Ohio) became depleted quite quickly. In Texas, oil was so plentiful that government intervention was required to keep price volatility in check. The U.S. government came to effectively control supplies: since East Texas production was, for decades, far below its potential, and given the authority to raise and lower production quotas as circumstances required, the U.S. government (both Federal and state, in particular the Texas Railroad Commission) had the power to increase or decrease oil supply almost at will. Over the decades since, while it still had that power, the U.S. government would use it to stabilize the market on numerous occasions. It increased production enormously during World War II, as well as during supply crises involving the Middle East, in 1953 (Iran), 1956 (Suez), and 1967 (Six-Day War). When the surge of oil was no longer needed, it had the power to reduce production once more. U.S. regulation thus acted as an automatic stabilizer (Yergin, 1991, page 259). Period I ends just before 1973, as production in Texas reached capacity. We call the supply regime in Period I "flexible" since in that period increased demand could be easily accommodated by increasing production, and indeed the regulatory structure in the U.S. was designed to do just that.

Since in Period I we think of oil supply as flexible, we model it as being responsive to demand shocks, as in equation 12. Note that our model's state variable is not supply but oil availability, which we measure as the sum of U.S. oil stocks at the beginning of the month and of U.S. oil production during the month (i.e. supply). From the equilibrium solution in Section 2 we know that, all else equal, storage (measured as U.S. oil stocks at the end of the month) should co-move with availability (see Figure 1), but should have a negative relationship with income (proxying for oil demand, and measured by U.S. industrial production; see Figure 2). Similarly, we know that the equilibrium relationship between availability and price is negative (Figure 1, lower panel). The relationship between income and price is positive by construction. In the model storage and price are co-determined in equilibrium by the no - arbitrage condition 13, implying that in Period I they should move in opposite directions due to the flexibility of supply: a price rise would lead to increased supply and a lower price in the future, thus causing a decrease in storage, all else equal. To sum up the model's predictions for Period I: availability and storage should move together, whereas storage should be negatively associated with both income and price.

The monthly data series pertaining to Period I (see Figure 3) seem to contain unit roots. It turns out that preliminary tests cannot reject the null that the series contains a unit root for each of the four series[11]. Since we cannot exclude the possibility that any of the variables we observe exhibits a unit root, we will treat all of them as integrated processes. However, our model predicts that there exists a stationary equilibrium: for any given levels of the state variables (availability and income), there are corresponding equilibrium levels of storage and price. Specifically, let

$$s_t = (storage_t, availability_t, income_{t,}price_t),$$

, and let $\gamma = (\gamma_1, \gamma_2, \gamma_3, \gamma_4)'$ denote a vector of coefficients, then the model predicts the existence of a long run equilibrium of the form

$$\gamma' s_t = 0.$$

In any particular period this relation will not be satisfied exactly: the prediction of a stationary equilibrium translates to a requirement that $\gamma' s_t$ is shown to be stable over time, i.e. it does not contain a unit root. In other words, given that the observed series may be integrated processes, the model predicts the existence of a cointegrating vector $\gamma$. In particular, the model predicts the signs of the components of $\gamma$: picking the sign of $\gamma_1$ as positive, we get that $\gamma_2 < 0$ (since storage and availability co-move in equilibrium), $\gamma_3 > 0$ and $\gamma_4 > 0$ (since storage has a negative equilibrium relationship with income and price). In the results section which follows we proceed in stages: first we determine the appropriate lag length $k$ for $s_t$, then we test whether a cointegrating vector exists using the standard Johanssen test, and finally we estimate a vector error correction model. But before that we turn to discuss the model's empirical predictions for Period II.

[11]The series were tested using the GLS version of Dickey-Fuller, separately for each sub-period. Full results available upon request. We could not reject the null even at the 10% confidence level, at any lag length up to the Schwartz maximum.

## 3.2 Empirical Predictions of the Model: Period II

We start Period II in 1973/1. This is not an arbitrary cutoff; rather, it is based on careful study of the structural breaks in persistence and volatility which exist in the long run real price data for crude oil. For details, the reader is referred to our earlier paper, Dvir and Rogoff (2009). In that paper, we find that in 1972 or 1973 there occurred a structural break in the price series, where the real price series became significantly more persistent, as well as significantly more volatile. We associate this break with the abrupt change in the supply regime faces by the West, from easy access to oil, to severely limited access. By 1973, the U.S. was producing oil at capacity; the ability of U.S. government agencies to increase production in times of need was gone (Yergin, 1991, pp. 567-8). Excess capacity existed now only in the Middle East, giving the rulers of these countries the ability to extract large rents from consumers by limiting access to oil supplies. As is well known, they took advantage of this new environment, most dramatically in late 1973, when the price of oil rose following the Arab - Israeli war. These years also saw the transfer of some ownership rights of the oil resources located on Arab land from the international oil companies to the Arab governments. These developments changed fundamentally the nature of the market: the oil producing countries were now owners (whole or part) of their reserves, the only easily-exploitable oil reserves left in the world [12]. For our purpose, Period II is characterized by "inflexible" oil supply, in the sense that there is no mechanism which forces oil producers to accommodate demand increases, and they may choose to let prices rise far above the marginal cost of production. Our assumption is that supply behaves according to equation 10, i.e. supply evolves according to technological progress, and (for simplicity) does not respond to demand conditions at all. As a result, persistent demand increases will lead to persistent price increases. Storage will respond endogenously by increasing as well, exacerbating the price increase. This dynamic is reflected in the model's equilibrium relationships. Importantly, storage is now expected to have a positive relationship with price, whereas in Period I

[12] There is a debate in the literature on whether OPEC can be shown to have acted collusively to withhold supplies from the market. Clearly OPEC's degree of control over prices has been inconsistent over the years, however that in itself does not settle the issue: Smith (2005, 2008) and Almoguera *et al.* (2011) are recent references. In our context, what matters is only OPEC's ability to restrict access to the world's only easily exploitable reserves of oil. This ability is undisputed (Smith, 2008).

this relationship was expected to be negative.

The series pertaining to Period II (see Figure 4) also may contain unit roots; Dickey-Fuller tests fail to reject a unit root in any of them. However, as in Period I, our model predicts the existence of a stationary equilibrium and we can proceed to estimate the cointegrating vector $\gamma$. The model's equilibrium relationships between the variables are again given by Figures 1 and 2: as before, storage (measured by end - of - month OECD stocks) and availability (measured by the sum of world production and beginning - of - month OECD stocks[13]) have a positive relationship, and price (measured by the refiner's acquisition cost, CPI deflated) and availability have a negative relationship (Figure 1, upper and lower panel respectively). However now storage and income (measured by OECD+6 GDP) have a positive relationship (Figure 2 lower panel), and, since income and price co-move by construction, storage and price will have a positive equilibrium relationship as well. In Period II, again picking the sign of $\gamma_1$ as positive, we expect that $\gamma_2 < 0$ (since storage and availability here also co-move in equilibrium), but now we also expect $\gamma_3 < 0$ and $\gamma_4 < 0$ (since storage now has a positive equilibrium relationship with income and price).

In the next subsection we present our empirical results for both periods.

## 3.3 Results

Our first task is to test whether a stationary cointegrating vector exists, as predicted by the model for both sub-samples separately. Table 1 presents the results of standard Johanssen tests conducted on the two sub-samples, with all variables included, as well as a constant and seasonal dummies[14]. The number of lags included is determined by the HQ information criterion, since it is a consistent statistic of the true number of lags[15]. We see that for both sub-samples the null hypothesis of no cointegrating vector (rank zero) is strongly rejected. In 1931/3 - 1973/12, there is no evidence of more than one

[13] This is an imperfect measure of availability as it does not include any stocks existing in non - OECD countries.

[14] Including deterministic trends in the cointegration relationships does not make a qualitative difference, either when testing for cointegration, or later when estimating the vector error correction models.

[15] Results of the Johanssen tests are not sensitive to the choice of lag number. The VECM estimates are sensitive to this choice, in coefficient size, but importantly not in sign or level of significance.

cointegrating vectors. However in 1975/1 - 2011/12 there is some evidence of more than one cointegrating vectors. We find very little support for that in further testing, and do not explore this here. Changing the beginning and ending months, within the limits detailed above, does not qualitatively change the test results. Our conclusion from these tests is that the variables we examine seem to have a common stochastic trend, separately for each sub-sample. We can therefore proceed to estimate the cointegrating vector.

Table 2 presents the results from an estimation of vector error correction models, under the assumption that in each sub-period there is exactly one cointegrating vector. The estimated model is

$$\Delta s_t = \alpha_0 + \alpha\gamma' s_{t-1} + \Lambda_1 \Delta s_{t-1} + ... + \Lambda_k \Delta s_{t-k} + u_t,$$

where $\alpha_0$ is a constant term accounting for the possible existence of time trends in some of the component series, $\alpha$ is a $4 \times 1$ vector of coefficients, $\Lambda_1, ..., \Lambda_k$ are $4 \times 4$ matrices of coefficients, and $u_t$ is a standard error term (see, for example, Lutkepohl [2006] for details on the estimation of vector error correction models). The number of lags and periods is the same as in Table 1. Note that the coefficient for log of inventories is normalized to one. The table shows the coefficients of the lagged variables in the estimated cointegration equation only, i.e. the Table only shows the estimated value of $\gamma$[16]. Note that for both sub-periods the cointegrating equations are extremely significant. All coefficients are significant at the 1% level. Note that a negative sign implies that the variable has a positive long-run relationship with storage.

In both sub-samples, corresponding to Periods I and II, we see that storage and availability have a positive equilibrium relationship, i.e. $\gamma_2 < 0$ as predicted by the model and described in Figure 1, upper panel. This is a standard prediction of commodity storage models, namely, that as the amount potentially available for consumption increases, storage should increase to take advantage of the relative abundance. Both variables are measured in the same units, so that the coefficients can be interpreted as the linear effect. Note that in Period II the effect is quite large - for any increase in availability, storage increases by more than two and a half times as much.

In the early sub-sample, corresponding to Period I, we estimate a negative equilibrium relationship between storage and price, whereas in the later sub-sample, corresponding to Period II, we estimate the opposite: a positive

[16]More results are available upon request. In particular, both VECMs are stable, and the we can reject the null of nonstationarity for both estimated cointegrating vectors at the 5% level using Dickey-Fuller GLS.

equilibrium relationship. This result is exactly as expected by our model: in periods of flexible supply, price increases should reduce storage, since those are good times to sell stocks. In contrast, in periods of inflexible supply, price increases can signal further increases in the future, and are therefore good times to buy stock, i.e. increase storage. This simple interpretation of our model's results seems to accord well with the equilibrium relationships found in the data, as shown in Table 2. The coefficients are large (storage is measured in millions of barrels)

We estimate a negative equilibrium relationship between storage and income in both sub-samples. This accords well with the model's prediction for Period I, since income increase price, which then leads to reduced demand for storage. However in Period II our model predicts that the same effect would lead to increased storage, a prediction which is not borne out by the data. There could be a number of reasons for this result, among them: measurement error in the OECD income or stocks data or the world oil production data; or a possible misspecification of the estimated model, perhaps due to an omitted variable[17]. We leave this issue to further research.

It is important to stress that these estimates represent the long-run relationship among the variables, i.e. there is no claim here of causality from any one variable to the other, rather the finding is of a long-run stationary link. This strongly supports the relevance of a model which posits such a link among the variables. The fact that the signs seem to (mostly) accord well with our model is encouraging. Experimenting with different starting and ending points, as well as varying the lag order, do not change the signs of the coefficients, nor the cointegration rank, nor the significance of the cointegrating equation or the estimated coefficients. Using a different price series (WTI instead of RAC), or using OPEC oil production instead of world oil production, or U.S. stocks instead of OECD stocks, lead to qualitatively similar results.

# 4 Conclusion

This paper presents evidence that important variables in the market for crude oil are connected by stable relationships, and have been at least since the

[17]The model actually has a third state variable, the rate of growth of income. However including that in the regression, with income or by itself, does not qualitatively change the results. These results are available upon request.

1930s. This evidence, of a single cointegrating vector connecting oil production, oil inventories, income, and the real price of crude oil, turns out to accord well with an extended storage model which allows for income growth dynamics and for changes in supply regimes. In particular, before 1973, when supply was unrestricted, inventories tended to decrease with the real price of oil. After 1973, when supply became restricted, the relationship changed, and the relationship between inventories and the real price became strongly positive. Whether the long-run relationships identified here will remain stable will be a function of the flexibility of oil supply in the future, assuming of course that our interpretation of the data as validating our model is correct. In particular, in view of recent estimates by the U.S. Dept. of Energy (EIA [2013]) regarding the availability of shale oil in the U.S., and assuming that the industry will remain competitive, within a few years we may be again in a period where increased demand can be easily met by more production from U.S. sources. This development may well reverse the long-run relationship between inventories and price. This is an exciting topic for future research.

# References


[1] Almoguera, Pedro A., Douglas, Christopher C., Herrera, Ana Maria, 2011. "Testing for the Cartel in OPEC: Non-cooperative Collusion or Just Non-cooperative?" *Oxford Review of Economic Policy,* 27(1): 144-168.

[2] British Petroleum, 2012. *Statistical Review of World Energy.* London.

[3] Chambers, Marcus J. and Roy E. Bailey. 1996. "A Theory of Commodity Price Fluctuations." *Journal of Political Economy,* 104(5): 924-957.

[4] Commodity Research Bureau. 1940. *Commodities in Industry.* New York: CRB, Inc.

[5] Commodity Research Bureau. 1950, 1960. *Commodity Yearbook 1950, 1960.* New York: CRB, Inc.

[6] Deaton, Angus and Guy Laroque. 1992. "Behaviour of Commodity Prices." *Review of Economic Studies,* 59(1): 1-23.

[7] Deaton, Angus and Guy Laroque. 1996. "Competitive Storage and Commodity Price Dynamics." *Journal of Political Economy,* 104(5): 896-923.

[8] Dvir, Eyal and Kenneth Rogoff, 2009. "Three Epochs of Oil." NBER Working Paper No. 14927.

[9] Energy Information Administration. 2013. Annual Energy Outlook 2013.

[10] Fattouh, Bassam, Lutz Kilian, and Lavan Mahadeva, 2012. "The Role of Speculation in Oil Markets: What Have We Learned So Far?", Working paper.

[11] Hamilton, James D. 2009. "Understanding Crude Oil Prices." *Energy Journal*, 30(2): 179-206.

[12] Kilian, Lutz and Thomas K Lee. 2013. "Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories." CEPR Discussion Paper #9297.

[13] Kilian, Lutz and D.P. Murphy. 2013. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics*, forthcoming.

[14] Knittel, Christopher R. and Robert S. Pindyck. 2013 "The Simple Economics of Commodity Price Speculation", Mimeo, MIT.

[15] Lutkepohl, Helmut. 2006. *New Introduction to Multiple Time Series Analysis.* Berlin: Springer.

[16] NBER. 2008. *NBER Macrohistory Database.* http://www.nber.org/databases/macrohistory/contents/

[17] Newbury, David M.G. and Joseph E. Stiglitz, 1981. *Theory of Commodity Price Stabilization.* Oxford, England: Oxford University Press.

[18] Pavlova, Anna, and Suleyman Basak, 2012. "A Model of Financialization of Commodities", Working paper, London Business School

[19] Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt. 2000. "Equilibrium Forward Curves for Commodities." *Journal of Finance,* 55(3): 1297-1338.

[20] Schwartz, Eduardo, and James E. Smith. 2000. "Short-Term Variations and Long-Term Dynamics in Commodity Prices." *Management Science,* 46(7): 893-911.

[21] Singleton, Kenneth J. 2013. “Investor Flows and the 2008 Boom/Bust in Oil Prices.” textitManagement Science.

[22] Smith, James L. 2008. “Organization of the Petroleum Exporting Countries (OPEC)”. In: Durlauf, Steven N., and Lawrence E. Blume (eds.), *New Palgrave Dictionary of Economics, 2nd Ed.,* New York: Palgrave Macmillan.

[23] Sockin, Michael, and Wei Xiong, 2012. “Feedback Effects of Commodity Futures Prices,” Working paper, Princeton.

[24] Williams, Jeffrey C. and Brian D. Wright. 1991. *Storage and Commodity Markets.* Cambridge, England: Cambridge University Press.

[25] Yergin, Daniel. 1991. *The Prize.* New York: Simon & Schuster.

**Table 1: Johanssen Tests for the Existence of Cointegration Vectors**

| | Column I | | Column II | |
|---|---|---|---|---|
| Period | 1931/3 - 1972/12 | | 1975/1 - 2011/12 | |
| Cointegrating Rank | 0 | 1 | 0 | 1 |
| Trace Statistic | 64.16*** | 24.97 | 64.70*** | 34.00** |
| 5% Critical Value | 47.21 | 29.68 | 47.21 | 29.68 |
| 1% Critical Value | 54.46 | 35.65 | 54.26 | 35.65 |
| Obs. | 502 | | 444 | |
| Differenced Lags | 1 | | 4 | |

Tests include a constant and seasonal dummies. Number of lags chosen by Akaike information criterion. (***) denotes that the trace statistic for the applicable rank is larger than the 1% critical value. (**) denotes that the trace statistic for the applicable rank is larger than the 5% critical value.

**Table 2: Long-Run Relationships of Stocks, Production, Demand, and Price**

| | Column I | Column II |
|---|---|---|
| Period | 1931/3 - 1972/12 | 1975/1 - 2011/12 |
| $\text{Storage}_t$ | 1 | 1 |
| $\text{Availability}_t$ | -0.97*** (0.05) | -2.60*** (0.38) |
| $\text{Income}_t$ | 5.22*** (0.30) | 37.75*** (9.47) |
| $\text{Price}_t$ | 9.07*** (1.09) | -22.57*** (6.53) |
| Obs. | 502 | 444 |
| Differenced Lags | 1 | 4 |
| $\chi^2$(p-value) | 470.13 (< 0.0001) | 91.32 (< 0.0001) |

Data sources and variable definitions: see text. Three asterisks (***) denote significance at the 1% level, two asterisks(**) denote significance at the 5% level. Standard errors are shown in parentheses. All regressions include a constant and seasonal dummies (not shown).

Figure 1: RE Equilibrium: Storage and Price As Functions of Effective Availability

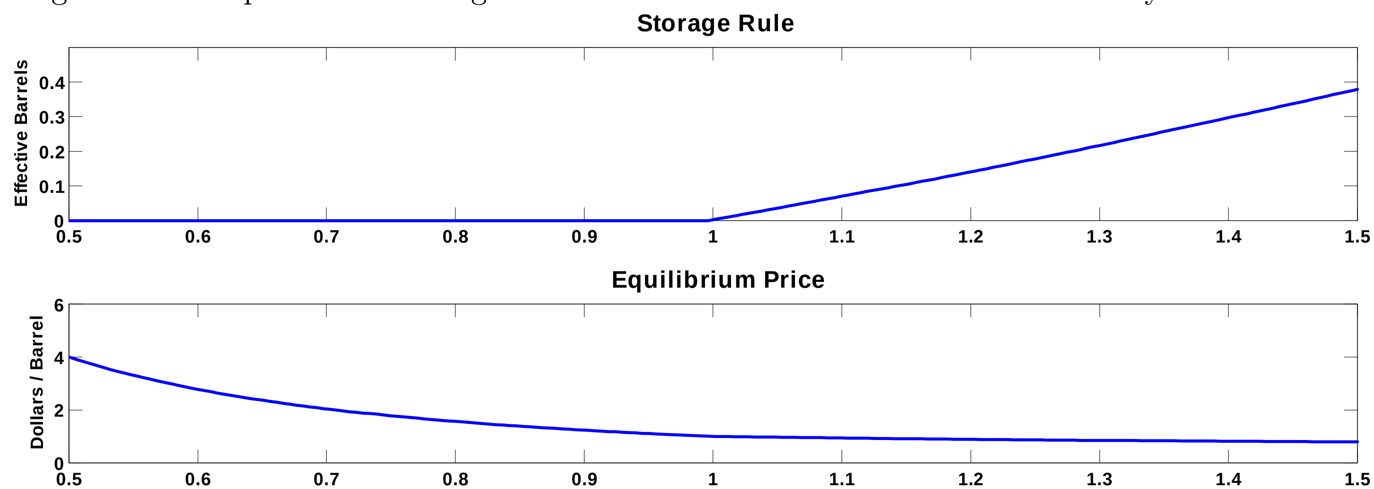

Figure 2: Effect of Change in Relative Income on Storage Across Models

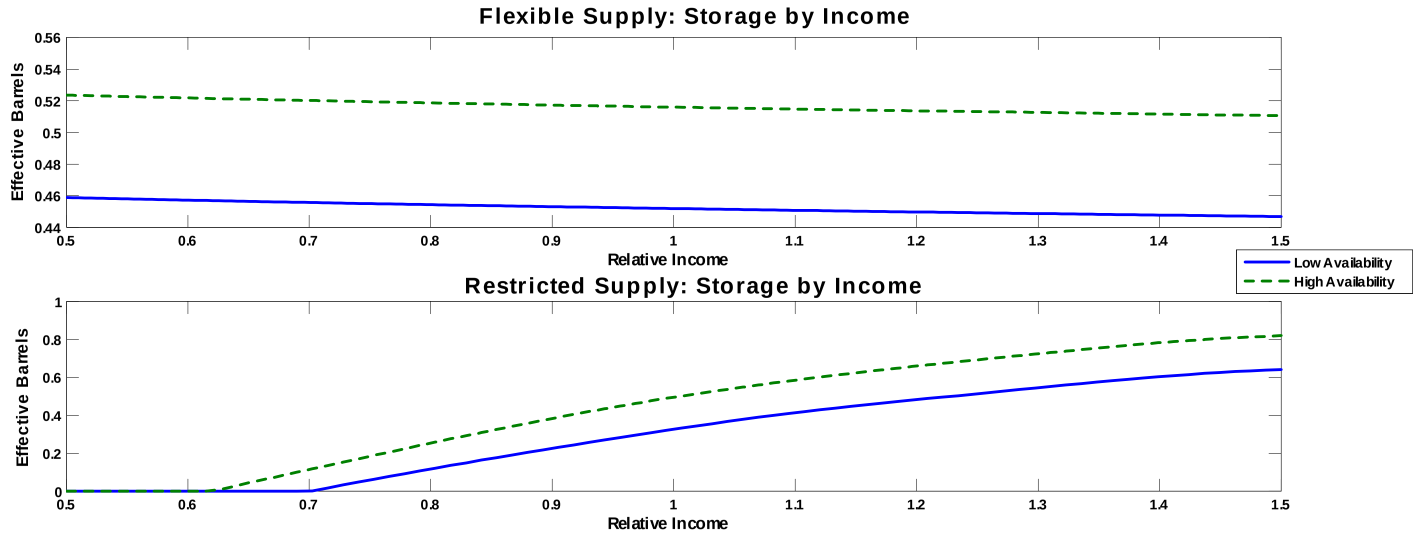

Figure 3: 1920/1 - 1972/12

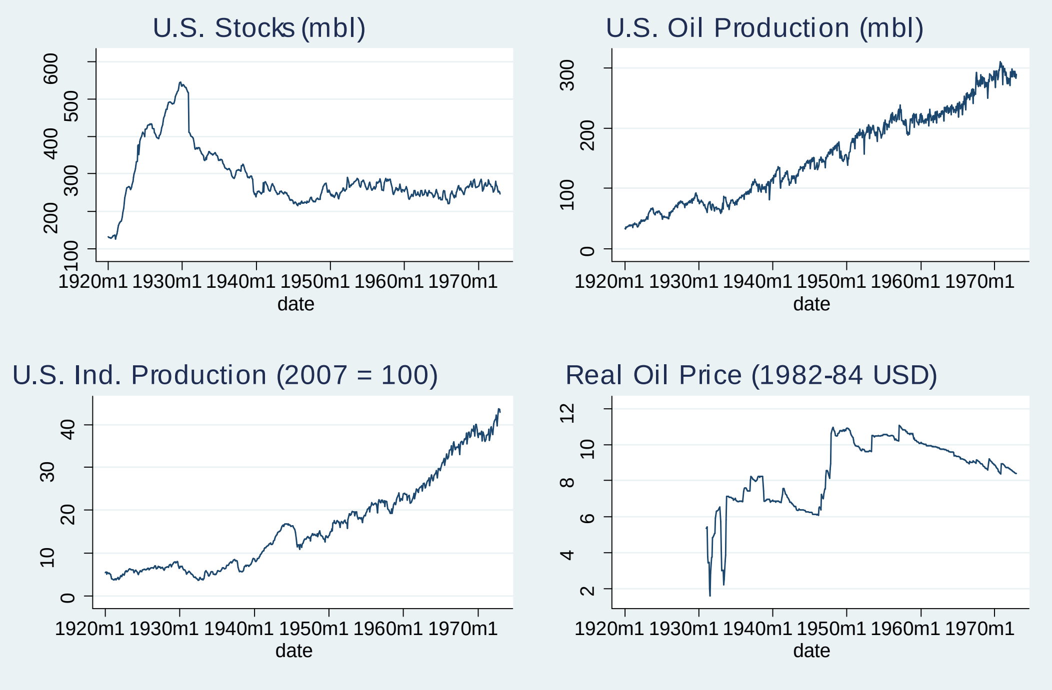

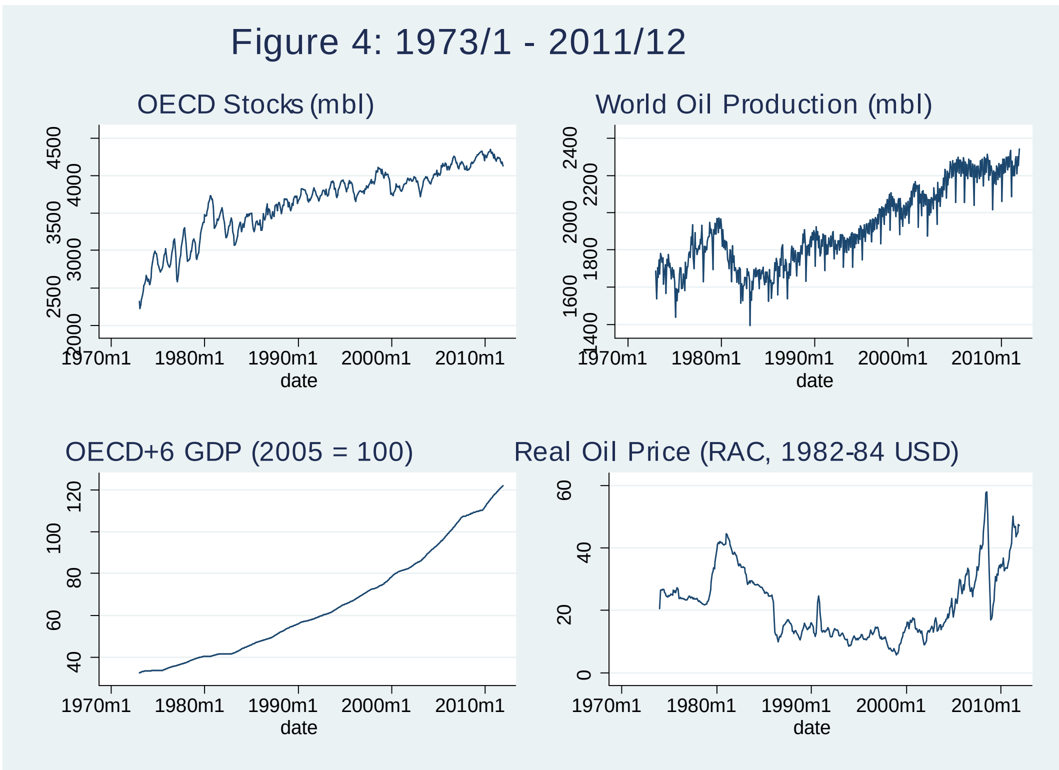

Figure 4: 1973/1 - 2011/12
OECD Stocks (mbl)
World Oil Production (mbl)
OECD+6 GDP (2005 = 100)
Real Oil Price (RAC, 1982-84 USD)
2000
2500
3000
3500
4000
4500
1400
1600
1800
2000
2200
2400
40
60
80
100
120
0
20
40
60
1970m1
1980m1
1990m1
2000m1
2010m1
date