Summer 8-12-2011

Oil Prices and Long-Run Risk

ROBERT READY
rready@wharton.upenn.edu

Oil Prices and Long-Run Risk

Abstract
I show that relative levels of aggregate consumption and personal oil consumption provide anexcellent proxy for oil prices, and that high oil prices predict low future aggregate consumptiongrowth. Motivated by these facts, I add an oil consumption good to the long-run risk model of Bansal and Yaron [2004] to study the asset pricing implications of observed changes in the dynamicinteraction of consumption and oil prices. Empirically I observe that, compared to the rst half of my1987 - 2010 sample, oil consumption growth in the last 10 years is unresponsive to levels of oil prices,creating an decrease in the mean-reversion of oil prices, and an increase in the persistence of oil priceshocks. The model implies that the change in the dynamics of oil consumption generates increasedsystematic risk from oil price shocks due to their increased persistence. However, persistent oil pricesalso act as a counterweight for shocks to expected consumption growth, with high expected growthcreating high expectations of future oil prices which in turn slow down growth. The combined eectis to reduce overall consumption risk and lower the equity premium. The model also predicts thatthese changes aect the riskiness of of oil futures contracts, and combine to create a hump shapedterm structure of oil futures, consistent with recent data.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Finance

First Advisor
Amir Yaron

Second Advisor
Andrew Abel

Third Advisor
Nikolai Roussanov

Keywords
Finance, Asset Pricing, Oil

Subject Categories
Finance and Financial Management

OIL PRICES AND LONG-RUN RISK

Robert Clayton Ready

A DISSERTATION

in

Finance

For the Graduate Group in Managerial Science and Applied Economics

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2011

Supervisor of Dissertation

Signature ____

Dr. Amir Yaron, Professor of Finance

Graduate Group Chairperson

Signature ____

Dr. Eric Bradlow, Professor of Marketing, Statistics, and Education

Dissertation Committee

Dr. Amir Yaron, Professor of Finance

Dr. Andrew Abel, Professor of Finance

Dr. Nikolai Roussanov, Assistant Professor of Finance

ABSTRACT

OIL PRICES AND LONG-RUN RISK

Robert Clayton Ready

Amir Yaron

I show that relative levels of aggregate consumption and personal oil consumption provide an excellent proxy for oil prices, and that high oil prices predict low future aggregate consumption growth. Motivated by these facts, I add an oil consumption good to the long-run risk model of Bansal and Yaron [2004] to study the asset pricing implications of observed changes in the dynamic interaction of consumption and oil prices. Empirically I observe that, compared to the first half of my 1987 - 2010 sample, oil consumption growth in the last 10 years is unresponsive to levels of oil prices, creating an decrease in the mean-reversion of oil prices, and an increase in the persistence of oil price shocks. The model implies that the change in the dynamics of oil consumption generates increased systematic risk from oil price shocks due to their increased persistence. However, persistent oil prices also act as a counterweight for shocks to expected consumption growth, with high expected growth creating high expectations of future oil prices which in turn slow down growth. The combined effect is to reduce overall consumption risk and lower the equity premium. The model also predicts that these changes affect the riskiness of of oil futures contracts, and combine to create a hump shaped term structure of oil futures, consistent with recent data.

Contents

List of Tables

List of Figures

1 Introduction

The significance of oil as an input into the macroeconomy, and its ability to predict future growth in economic variables, suggests that the oil price is an important variable to consider in the context of consumption based asset pricing models.[1] Though these models have had substantial success in linking exposure to macroeconomic risk to the observed behavior of equity prices, there has been little work examining oil price risk in this context. I develop a model to study how changes in the dynamics of oil consumption and aggregate consumption over the last decade affect the risk premia associated with oil prices.

The model is an endowment model of consumption. Motivating this choice is a new fact about the relation between oil prices and personal consumption, namely that real oil prices can be closely approximated by a function of the relative levels of household oil consumption and aggregate consumption (excluding oil consumption), where oil prices are high when oil consumption is low relative to aggregate consumption[2]. I also find that high oil prices predict low aggregate consumption growth. This predictive relation has particular importance for the Long-Run Risks (LRR) model of Bansal and Yaron [2004], which relies on a predictable component of consumption growth to explain observed behavior of asset prices. In order to study these effects I add an oil consumption good to the LRR framework.

I use the model to study how observed changes in the dynamics of oil consumption over second half of my 1987 - 2010 sample translate to changes in the risk premia associated with oil prices. Over the second half of the sample, oil consumption growth becomes unresponsive to levels of oil prices. This change means oil prices exhibit significantly less mean-reversion in recent years, so

[1] Hamilton [2005] documents that oil shocks have a significant negative relation with future GDP growth for 1973 - 2005

[2] Oil consumption is defined personal consumption of energy goods taken from the NIPA survey. Following Yogo (2006) and Yang (2010), aggregate consumption is an aggregation of expenditure on nondurables and services (excluding energy goods) and the flow of services from the stock of durable goods. Further details are in the data section.

that shocks to oil consumption result in a much more permanent change in the level of oil prices. Since the persistence of shocks is the main driver of riskiness in the LRR framework, these shocks will command a larger risk premia in the recent period. In addition, this unresponsiveness means that shocks to aggregate consumption growth will have a larger effect on oil prices. Therefore, high expected future growth will result in high expectations of future oil prices, so that oil prices will act as a counterweight to shocks to expected growth thus reducing risk associated with these shocks. These two intuitive effects imply significant changes for both the overall level of risk in the economy and the riskiness of exposure to oil prices. In fact, they generate changes in expected returns to futures contracts which can explain the significant changes in the term structure of oil futures over the last 10 years, most notably the development of a "hump" shaped term structure of oil futures.

This mapping of the change in consumption dynamics to a change in the riskiness of oil futures is important not only because it gives insight into how changing conditions in the oil market translate to changes in risk, but also because it provides evidence for two very important aspects of the LRR model. (i) The relation between the persistence of shocks to growth and their associated level of risk, and (ii) the relation between the timing of cash flows and their associated risk premium. These two effects are difficult to observe in the standard consumption and equity data, since expectations of future consumption growth and the persistence of shocks to this growth are generally difficult to identify, and since the risk associated with specific dividend payments at different horizons is hard to identify from equity prices.[3] Since oil prices predict future consumption growth, the existence of oil futures contracts make oil prices an ideal laboratory to study these effects, both because futures contracts allow for measurement of expectations of persistence, and because the cross-section of different maturities allows for observation of risk premia at different time horizons.

Most models of oil prices consider oil as an input to production, and therefore require modeling

[3]Binsbergen et al. [2010] construct synthetic dividend strips from option values to study the risk associated with individual cashflows. They find that risk premia do not increase as the time to realization of the cashflow increases, and interpret this as evidence against the long-run risk formulation.

the decisions of oil producers, as well as producers of final consumption goods. This would greatly complicate the modeling of oil prices in this framework, but this issue can be avoided by utilizing the fact that oil itself enters into the consumption basket through the personal consumption of gasoline. The intratemporal utility function I propose is a generalized constant elasticity of substitution function (GCES). This function allows for non-homotheticity, which I find to be important to match the observed data. I find that empirically, oil consumption is both highly complementary to aggregate consumption, and that oil is a necessary, rather than a luxury, good. I also find that oil consumption expenditure is very small relative to aggregate consumption expenditure, so that the importance of oil is not in its direct impact on consumption, but rather in the ability of the oil price to predict future consumption growth.

I find that empirically, the implied price performs very well, explaining 85% of the total variation in oil prices over the observed time period. To my knowledge this is a novel formulation of oil prices, however it is in the same spirit of tests of Bentzen and Engsted [1993], Ramanathan [1999] and others, who estimate the response of gasoline consumption to changes in personal income and the price. These studies rely on a measures of economy wide gasoline or oil consumption taken from the Energy Information Administration (EIA). I perform similar analysis using GDP and personal income in place of aggregate consumption, and the EIA measure in the place of personal consumption of gasoline, and find that using aggregate consumption provides a small increase in explanatory power of oil prices over GDP and income. I also find that the NIPA measure personal consumption of gasoline provides a very large increase (almost 20% in terms of R^2) over the usual measure of economy wide oil consumption. These findings motivate my choice of variables, and more importantly illustrate the close links between oil prices and personal consumption, providing a more general motivation for a consumption based explanation of oil prices and risk premia.

Much of the literature on commodities prices has its roots in the theory of storage (Kaldor [1939], Working [1949], Telser [1958]) and until very recently, most work in this area fell into one of two categories. The first specifies an exogenous process for the stock price to examine the pricing implications for derivative contracts (Brennan and Schwartz [1985], Gibson and Schwartz [1990], Schwartz [1997]), while the second uses the theory of storage to derive implications of the price of oil (Williams and Wright [1991], Deaton and Laroque [1992], Deaton and Laroque [1996], Routledge et al. [2000]). More recent research (Carlson et al. [2007], Kogan et al. [2009]) has focused on oil production to generate futures price dynamics. These recent studies focus primarily on dynamics of the futures prices under the physical measure, and while they allow for a specification of the risk premium, they do not provide a theoretical explanation of the price of commodities risk. Casassus et al. [2005] develop a general equilibrium model with oil as an input into the production of a single consumption good, and study the implications of oil price risk in this context. Their model generates a curve which is sometimes hump-shaped, but the shape is generated by the expected change in future oil spot prices rather than differing risk premia across the curve. In addition, they also find that oil price risk can change based on the condition of oil production. However, the mechanism relies on the distance of the oil price from the level necessary to induce further investment in oil wells, and is therefore distinct from the effects described here, which reflect a more fundamental shift in the dynamics of oil consumption.

Studies applying traditional asset pricing models to explain risk premia in commodity prices have met with limited success (see Dusak [1973], Breeden [1980] and Jagannathan [1985]). Another common theory to explain the observed positive risk premia, or "Normal Backwardation", as introduced by Keynes [1930] postulates that producers who are seeking to hedge risks of future price movements are willing to pay a premium to speculators. Gorton et al. [2007a] show that Sharpe ratios of commodities prices over the last 40 years are significantly higher for commodities futures than for equities, and that levels of inventory predict futures returns, which they interpret as support for this theory. While the results here may help shed some light on the source of risk premia in commodities, it is important to keep in mind that the results in this paper depend greatly on the relations between oil prices and consumption which are unique among commodities.

The rest of the paper is organized as follows. Section 2 describes the model and shows how changes in parameters governing the responsiveness of oil consumption create changes in risk. Section 3 describes the observed behavior of consumption and oil prices, and documents the changes in these dynamics as well as the changes in the term structure of oil futures prices over the the sample period. Section 4 discusses extensions to the model. Section 5 calibrates the model to match salient moments of asset prices and consumption. Section 6 Concludes.

2 The Model

The model adds an oil consumption good to the long run-risk framework. Recent work by Yang [2010] emphasizes that durable consumption growth exhibits much higher persistence than nondurable consumption growth, and that this higher persistence can be used in a model of long-run risk to explain the equity premium and risk-free rate puzzles. I find that this higher persistence is important in explaining the observed term structure of oil futures. I also find that including durable goods strengthens the relation between levels of consumption and the spot price of oil. For both of these reasons including durable consumption is important to generate the implications of the model.

Considering durable consumption and nondurable consumption separately generates an extra term in the stochastic discount factor when using Epstein-Zin Preferences, reflecting the fact that consuming a durable good exposes the representative agent to price risk generated by the changing composition of consumption[4]. I assume that $C_t = N_t^{1-\alpha} D_t^{\alpha}$, where N_t is the expenditure on non-durables and services excluding oil, and D_t is the services flow from the stock of consumer durable goods, which is assumed to be linear in the stock. I consider this aggregation as the consumption good. Oil prices will be in terms of the price of this good.

Pakos [2004], considers a model with utility arising from an aggregation of nondurable and

[4]For a full discussion of the issues involved using durable consumption in a model with Epstein - Zin preferences see Yogo [2005] and Yang [2010]

durable goods using a Generalized Constant Elasticity of Substitution (GCES) felicity function. Here I follow Yang [2010] and consider a Cobb-Douglas aggregate of durable and nondurable goods, I then use the GCES functional form to represent utility across the aggregate consumption good, C_t, and an oil consumption good, O_t. The representative consumer has utility $V_t(C_t, O_t)$ in each period, where

$$V_t(C_t, O_t) = \left[(1-a)C_t^{1-\frac{1}{\rho}} + aO_t^{1-\frac{\eta}{\rho}}\right]^{\frac{\rho}{\rho-1}} \tag{1}$$

This function nests several of the commonly used utility functions. For $\eta = 1$, V_t is the standard Constant Elasticity of Substitution (CES) function. For $\rho = 0$ the function is the Leontieff function and for $\rho = 1$ the function is Cobb - Douglass. I find empirically that $\rho < 1$ suggesting that oil consumption is a complement to aggregate consumption rather than a substitute, and that η is substantially greater than one, suggesting that oil consumption goods are necessary, rather than luxury, goods. Given this function, optimal behavior by the consumer implies that the price of oil in terms of units of the aggregate consumption good is the ratio of the marginal utilities of oil and aggregate consumption.

$$P_t = \frac{a(1-\frac{\eta}{\rho})C_t^{\frac{1}{\rho}}}{(1-a)(1-\frac{1}{\rho})O_t^{\frac{\eta}{\rho}}} \tag{2}$$

Taking logarithms, where p_t, c_t, o_t representing logs of price, aggregate consumption and oil consumption, yields

$$p_t = \text{constant} + \frac{1}{\rho}(c_t - \eta o_t) \tag{3}$$

I then embed this intratemporal utility function within Epstein and Zin preferences so that total utility is

$$U_t = \left[(1-\delta)V_t^{\frac{1-\gamma}{\Theta}} + \delta\left(E_t[U_{t+1}^{1-\gamma}]\right)^{\frac{1}{\Theta}}\right]^{\frac{\Theta}{1-\gamma}} \tag{4}$$

Where γ is the coefficient of risk aversion and ψ is the intertemporal elasticity of substitution (IES). Having specified the utility of the representative agent, what is left is to specify dynamics of oil consumption and aggregate consumption. The consumption dynamics I consider have the following form.

$$
\begin{aligned}
\Delta c_{t+1} =& \mu^c + \pi^c \left[c_t - \eta o_t - \bar{p}\right] + x_t + \sigma_{c,t} e^c_{t+1} \\
\Delta o_{t+1} =& \mu^o + \pi^o \left[c_t - \eta o_t - \bar{p}\right] + \Phi_x x_t + \sigma_{o,t} e^o_{t+1} \\
x_{t+1} =& \rho_x x_t + \varphi_x \sigma_{c,t} e^x_{t+1} \\
\sigma_{o,t+1} =& \nu_o(\sigma^2_{o,t} - \bar{\sigma_o}^2) + \bar{\sigma_o}^2 + \sigma^o_w w^o_{t+1} \\
\sigma_{c,t+1} =& \nu_c(\sigma^2_{c,t} - \bar{\sigma_c}^2) + \bar{\sigma_c}^2 + \sigma^c_w w^c_{t+1} \\
y_t =& \mu^y + \chi \left(x_t + \pi^c \left[c_t - \eta o_t - \bar{p}\right]\right) + \varphi_y \sigma_{c,t} e^y_{t+1}
\end{aligned} \tag{5}
$$

Here o_t represents log of oil consumption, c_t is log of aggregate consumption, x_t is a predictable component in long run aggregate consumption growth, and y_t is the log of the aggregate dividend. This specification combines features of both Bansal and Yaron [2004] in that it includes a separate process for the predictable consumption rate growth, and Hansen et al. [2008] in that it includes an additional source of predictable consumption growth coming from the error correction term $(c_t - \eta o_t)$. Dividends are a levered claim on consumption, as in Bansal and Yaron [2004], χ represents the leverage coefficient. Correlation among the innovations is straightforward to include, but for parsimony here I assume they are independent of each other, and i.i.d. with a $N(0, 1)$ distribution. When calibrating the model I set the correlations to match observed correlations in the data.

The shock to e^o_t represents an innovation to oil consumption, which is also an innovation to the oil price, p_t that is unrelated to a change in c_t. For this reason I will refer to e^o_t as an oil price shock. It is important to note that a positive innovation to e^o_t represents a negative innovation to p_t. I also specify two sources of stochastic volatility, $\sigma_{o,t}$ governing the volatility of oil consumptions shocks, and $\sigma_{c,t}$ governing shocks to the other variables in the economy.

The x_t component represents a predictable component of consumption growth similar to the model of Bansal and Yaron [2004]. This model is sometimes criticized for the low level of predictability in consumption growth. However, as Yang [2010] shows, there is in fact significant predictability in durable consumption growth. This predictability is also present in the Cobb-Douglas aggregation of durable and nondurable consumption used here.

In addition to x_t, there is also predictable growth coming from the error correction term $(c_t - \eta o_t - \bar{p})$. In this sense this model is similar to that of Hansen et al. [2008], which specifies that the difference between consumption and earnings is predictive for future growth. In this model, since oil prices are represented by $\frac{1}{\rho}(c_t - \eta o_t)$, a negative value of π^c captures the idea that high oil prices are predictive for consumption growth. It is important to note here that this specification implies that the oil price is an $I(0)$ variable. Equivalently, it implies a cointegrating relation between c_t and o_t, and two cointegrating relations between c_t, o_t, and p_t. I provide tests for these relations in Appendix A. I find support for this specification from Johansen [1991] tests in estimates of a vector error correction model of oil consumption and aggregate consumption. While these results are potentially interesting, I focus here on the simpler specification of dynamics which allows for an easier interpretation in the familiar context of the long run risk model.

This model here is a slightly simplified version of the model I take to the data. I make two additions to capture two commonly thought of features of oil prices. One is adding drift to the long run price $\bar{p}$. The second is to add an external habit to the specification for oil prices. These changes allow for a better quantitative fit of futures curves but do not in any way effect the qualitative implications of the model. Both extensions are discussed in more detail in Section 4, and the full specification is solved in Appendix B.

2.1 Model Solution

The model solutions, though tedious to derive, produce expressions for asset prices that are easily interpretable as a linear factor model. The log of stochastic discount factor will be a linear function of the state variables, and therefore its innovation will be linear in the innovations to the consumption dynamics specified in system (5), with each innovation being multiplied by an associated price of risk. The expected returns of an asset, such as an oil futures contract, will then be a function of its loadings on the innovations and their associated price of risks. Here I first derive an expression for the stochastic discount factor. Section 2.2 derives expressions for futures prices and their loadings. Section 2.3 provides intuition for how changing two parameters, Φ_x and π^o, in the consumption process changes both the prices of risk and the loadings of futures to generate the observed changes in the term structure.

In order to solve the model, I follow the procedures of Bansal and Yaron [2004] to develop approximate analytical solutions to asset prices. In addition to the Campbell-Shiller approximation of returns, I require an additional approximation in order to handle the GCES function of intratemporal utility. As shown in Appendix B, log of V_t can be approximated as a Cobb-Douglas utility function.

$$\tilde{V}_t = C_t^{1-\tau} O_t^{\tau} \tag{6}$$

The value of τ is equal to the average proportion of consumption expenditure on oil goods, which is approx 3% in the data. Due to the small value of expenditure on oil consumption relative to aggregate consumption, this approximation performs extremely well. For the following calculations I will assume $V_t = C_t$ for parsimony. In calculating numerical results I will not impose this condition. I find that this assumption has very small effects on the results. This is an important point of differentiation between my model and the models of Pakos [2004], Yogo [2006], and Yang [2010], which rely on the degree of substitution between durable and nondurable consumption to generate asset pricing implications. The results in this model are driven by the growth rate dynamics of c_t and o_t, and the function V_t is merely a means to obtain the expression for the oil price in terms of consumption. The model is functionally equivalent to the standard long-run risk model with an exogenous specification of p_t, however describing the full model both confirms this equivalence and more generally motivates the use of a consumption based approach.

For the sake of exposition, here I also assume that there is a zero price of risk associated with shocks to volatility. The calibrated model solved in Appendix B includes these effects. Bansal et al. [2007] show that risks associated with shocks to a persistent stochastic volatility component can be important in explaining asset prices. In my calibration, the shocks to the latent expected growth of the aggregate consumption and the shocks to oil prices are the primary source of risk.

The representative agent has utility

$$U_t = \left[(1-\delta)C_t^{\frac{1-\gamma}{\Theta}} + \delta\left(E_t[U_{t+1}^{1-\gamma}]\right)^{\frac{1}{\Theta}}\right]^{\frac{\Theta}{1-\gamma}} \tag{7}$$

Following Epstein and Zin, the stochastic discount factor has the following form

$$M_t = \delta^{\Theta}\left(\frac{C_{t+1}}{C_t}\right)^{-\frac{\Theta}{\psi}} R_{W,t+1}^{\Theta-1} \tag{8}$$

To solve for the equilibrium return on wealth, I follow Bansal and Yaron [2004], and exploit the Campbell approximation for the log return

$$r_{W,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + \Delta c_{t+1} \tag{9}$$

I then assume, ignoring the contribution of stochastic volatility risk, that the log of the price-dividend ratio for consumption has the form,

$$z_t = A_0 + A_1 x_t + A_2(c_t - \eta o_t) \tag{10}$$

Exploiting the pricing equation

$$1 = E_t[exp(m_{t+1} + r_{g,t+1})] \tag{11}$$

Allows for solution of the coefficients. The coefficients for A_1 and A_2 are given by

$$A_1 = \frac{(1-\frac{1}{\psi}) + A_2\kappa_1(1-\eta\Phi_x)}{1-\kappa_1\rho_x} \tag{12}$$

$$A_2 = \pi^c \frac{(1-\frac{1}{\psi})}{1-\kappa_1(1+\pi^c-\eta\pi^o)} \tag{13}$$

$$\tag{14}$$

These values are very similar in flavor to the coefficient for the long-run risk shock, x_t in the standard formulation of Bansal and Yaron [2004]. The expression A_2 takes the sign of π^c, and represents the contribution of the predictable growth in consumption generated by the oil price to the expected consumption to wealth ratio. The A_1 term is the same as that of Bansal and Yaron [2004] with an additional term generated by the effect of x_t on the oil price. These values can then be used to calculate the log of the pricing kernel, with the innovation having the following expression.

$$m_{t+1} - E_t[m_{t+1}] = -\lambda_{m,c}\sigma_{c,t}e^c_{t+1} - \lambda_{m,x}\varphi_x\sigma_{c,t}e^x_{t+1} - \lambda_{m,o}\sigma_{o,t}e^o_{t+1} \tag{15}$$

Empirically I find that the correlation between the shocks e^c_{t+1} and e^o_{t+1} is such that innovations to e^c_{t+1} have little effect on the contemporaneous spot price. When I impose that the correlation is such that there is no effect, the prices of risk associated with each shock are given by

$$\lambda_{m,c} = \gamma \tag{16}$$

$$\lambda_{m,x} = (1-\Theta)A_1\kappa_1 \tag{17}$$

$$\lambda_{m,o} = -\eta(1-\Theta)A_2\kappa_1 \tag{18}$$

The first term in Equation (15) is the standard Breeden [1980] CCAPM term. The second term represents innovations to long run expectations in consumption growth as in Bansal and Yaron [2004]. The third is the innovation due to shocks to oil consumption, or equivalently oil price shocks.

2.2 Oil Futures Prices

The oil futures price[5] for a future with maturity j is described by the equation

$$0 = E_t\left[M_{t+1}(F_{t+1}^{j-1} - F_t^j)\right] \tag{19}$$

Exploiting the log-normality of both P_t and M_t and rearranging yields the following expression for the log of futures prices.

$$f_t^j = E_t[f_{t+1}^{j-1}] + \frac{1}{2}\text{var}_t(f_{t+1}^{j-1}) + \text{cov}_t(f_{t+1}^{j-1}, m_{t+1}) \tag{20}$$

That is the futures price is the log of the expected futures price for the same maturity one month from now, plus a covariance term that reflects the riskiness of the contract. While closed form expressions for various futures contracts are messy, they can be calculated through a simple recursive algorithm.

Futures prices can be expressed as linear function of the state variables

$$f_t^j = B_0^j + B_x^j x_t + B_p^j(c_t - \eta o_t) + B_{\sigma,c}^j \sigma_{c,t} + B_{\sigma,o}^j \sigma_{o,t} \tag{21}$$

Where the expressions are given in Appendix B. The initial value of the recursion represents the relation $f_t^0 = p_t$ so $B_p^0 = \frac{1}{\rho}$ while the other coefficients are zero.

These equations can also be used to calculate the expected returns on a futures contract. The expected return is

[5] I assume for simplicity that futures are marked to market on a monthly basis

$$E[r_t^j] = E_t[f_{t+1}^{j-1} - f_t^j] + \frac{1}{2}\text{var}_t(f_{t+1}^{j-1}) \tag{22}$$

The expected returns on futures depend on the loadings of futures prices different on the three state variables that describe the stochastic discount factor and the prices of risk of different prices of shocks. In the full model there are five shocks with associated prices of risk. As mentioned previously, the shocks to the two stochastic volatility components do not have a significant price of risk associated with them. Also in my calibration $\lambda_{m,c}$ is very small, so expected returns are driven mainly by two factors: shocks to expected growth, e_t^x and shocks to oil consumption, e_t^o, so that

$$E[r_t^{j+1}] \approx \frac{-\eta}{\rho} B_p^j(\lambda_{m,o}\sigma_{o,t}^2) + B_x^j \lambda_{m,x}\varphi_x \sigma_{c,t}^2 \tag{23}$$

This is the sense in which the long-run risk framework allows for a very intuitive linear factor model to explain the expected returns on oil prices. The return of future j depends on its loading on the two shocks, the B terms, and the prices of risk associated with exposure to each shock, the λ terms. The observed differences in the two parameters of consumption dynamics, π^o and Φ_x, have implications for both the loadings and prices prices of risk, and therefore change the expected return on futures prices.

2.3 Changing Φ_x and π^o

I will be focused on changing the values of two parameters as informed by the observed changes in consumption dynamics. Though empirical results will motivate these changes, it is worth discussing what they represent in an economic sense. The advantage of developing an endowment economy of consumption is that the economist may be agnostic to the sources of the shocks to consumption, while still being able to make inferences about their effect on asset prices. It is important to keep in mind however, that behind this model there is a real economy of production, supply, and demand which is generating the observed dynamics in consumption. I view the changes in parameters as a reflection of changes to the state of this economy, particularly in respect to the elasticity of crude oil production to respond to increases in the oil price.

In the model here, the parameters π^o and Φ_x are both intuitively related to the elasticity of oil supply, π^o as the speed with which oil consumption responds to an increase in price, and Φ_x as the expected increase in oil consumption corresponding to an expected increase in aggregate consumption. In a state of the world where production is highly elastic we expect π^o and Φ_x to be higher than in a state in which production is unable to respond, and indeed that is what I observe in the data.

Both of these changes in parameters, reflecting that oil consumption growth reacts differently to changes in oil prices or expected aggregate consumption growth, can be potentially explained by an inability of the oil industry to increase supply in response to changes in demand over the second half of the sample. There are many possible explanations for this, such as "Peak Oil" or a more temporary condition caused by increases in demand, such as from growth in developing countries, outstripping current production capacity, as is evidenced here by a quote from the International Energy Administration's Monthly Report in October of 2004:

> In response to rising prices, producers have increased supply to record levels. While this is a welcomed development, it reduces the amount of spare production capacity available to the market to offset supply disruptions associated with political and weather-related events. Consequently, prices have been subject to upward pressure. As prices shift upwards, the market has become more volatile and jittery and the demand for paper barrels has increased to offload risk. **-IEA Monthly Outlook, Oct 2004**

For the value of Φ_x, the estimate for the quarterly data prior to 2000 is positive, suggesting that oil consumption grows in response to expected growth. In fact, it is high enough to imply that expected growth in consumption implies negative growth in oil prices. This result seems economically unlikely, and since the observed value of Φ_x is significantly different from zero but not from $\frac{1}{\eta}$, I set the value of $\Phi_x = \frac{1}{\eta}$ in the first period so that a shock to x_t has no effect on future oil prices. For the second period the quarterly estimate of Φ_x is negative but not significantly different from zero, so I set $\Phi_x = 0$. With this value an increase in x_t has no effect on future oil consumption growth, and hence implies growth in oil prices.

The parameter π^o governs the rate with which oil consumption responds to a change in price to return prices to the long run stable oil price. The persistence of oil prices is simply the persistence of the cointegrating vector, $c_t - \eta o_t$, and has a value of $(1 + \pi^c - \eta\pi^o)$. Therefore, a high value of π^o will lead to low persistence of oil prices. I use monthly values of π^o that give the observed values from the quarterly data, with a higher value in the calibration for the first period, and a value close to zero for the second.

For the choice of the parameter π^c, I keep the values the same across the two calibrations of the model. Though the estimates in the data across the two periods are different, given the evidence that oil prices negatively predict future growth over longer time horizons, I keep π^c as a constant and focus on the effects of changes to π^o and Φ_x.

To further illustrate how changes to these parameters affect oil prices, Figure 1 shows the impulse responses to shocks to both oil consumption (an oil price shock), and the parameter x_t (an expected growth shock) under the two different parameterizations of the model. Plots (a) and (c) show the impulse response of c_t, o_t, and p_t to a negative innovation to e^o_t, which is equivalent to a positive oil price shock. As is evident in the first plot, a larger value of π^o means that the high price will induce growth in oil consumption in prior periods, which will result in a falling oil price. However, in the second period, the lower value of π^o means that the oil price will remain high, or that the shock to oil prices will be more persistent.

This change in π^o also has an effect on the response of c_t. Though the value of π^c is equal in the two figures, the continuing high oil price means that in the second period, the negative growth of oil prices persists longer than in the first period. This has an important effect on the magnitude

Figure 1: Model Impulse Response Functions: Basic Model

Period 1: $\pi_o = .1$ and $\Phi_x = \frac{1}{\eta}$

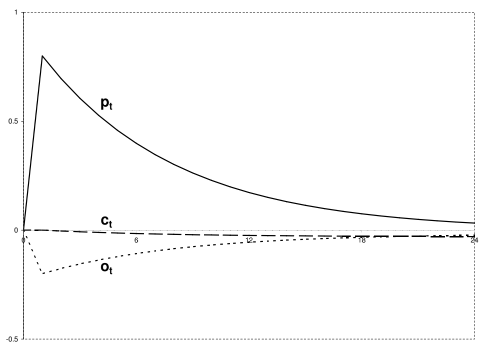

(a) Negative shock to e_t^o

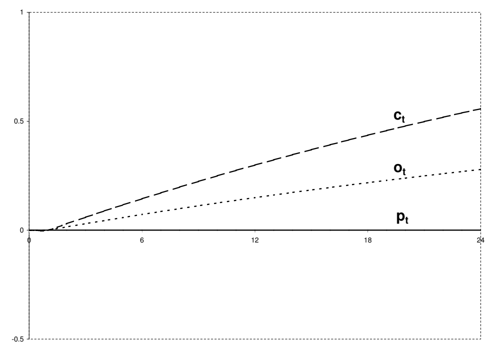

(b) Positive shock to e_t^x

Period 2: $\pi_o \approx 0$ and $\Phi_x = 0$

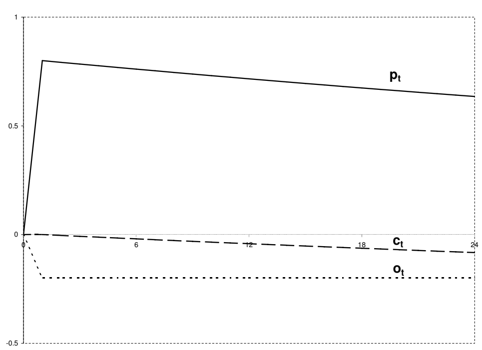

(c) Negative shock to e_t^o

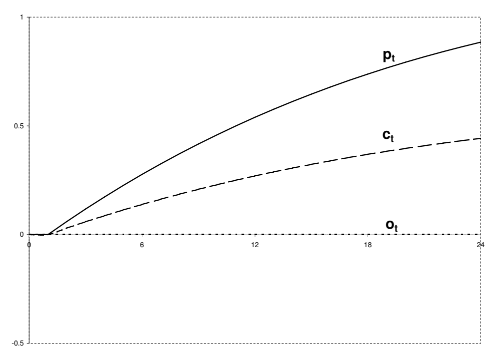

(d) Positive shock to e_t^x

Impulse response function of logs of aggregate consumption (c_t), oil consumption (o_t), and the oil price ($p_t = \frac{1}{\rho}(c_t - \eta o_t)$) to innovations to oil consumption and the expected growth of aggregate consumption.

of the risk premium associated with oil price shocks, since the persistence of expected growth is the primary determination of the price of growth risk in the Long-Run risk framework.

Plots (b) and (d) illustrate the differences under the two parameterizations of a positive shock to e_t^x. In Plot (b), oil consumption is expected to grow, so a shock to x_t has no impact on the price of oil. In Plot (d), with $\Phi_x = 0$, the shock to expected growth leads to an expected increase in oil prices as expected oil consumption growth is no longer higher. Therefore, a shock to expected growth in the second period has a large effect on expectations at long horizons, and relatively little effect at short horizons.

Both of these changes, the increase in growth risk from shocks to oil prices, and the increasing loading on expected growth shocks at longer horizons are important in generating the changes observed in the term structure of futures, and they are both reflected in the approximate analytic solutions to the prices of risk associated with each shock and the loadings of oil futures on the state variables of the model.

2.3.1 Changes in Prices of Risk

In order to examine how changes in parameters affect the prices of risk associated with shocks to future consumption growth, it is worthwhile to look more closely at how the coefficients for x_t and $c_t - \eta o_t$, A_1 and A_2 relate to the standard coefficient on x_t in the model of Bansal and Yaron [2004]. That coefficient is

$$A_1^{BY} = \frac{(1 - \frac{1}{\psi})}{1 - \kappa_1 \rho_x} \tag{24}$$

When $\psi > 1$, this coefficient is positive. Since $\kappa_1 \approx 1$, with a value of persistence, ρ_x, near one, this term can be very large implying a large magnitude for the price of risk of shocks to x_t. This coefficient is very similar to the coefficient associated with the relative level of aggregate consumption to oil consumption $c_t - \eta o_t$ in the model presented here, A_2

$$A_2 = \pi^c \frac{(1 - \frac{1}{\psi})}{1 - \kappa_1(1 + \pi^c - \eta\pi^o)} \tag{25}$$

Here, the value $(1 + \pi^c - \eta\pi^o)$ is the persistence of the oil price, and π^c is the effect the oil price has on consumption growth. Since π^c is negative, if the oil price is persistent then shocks to oil prices will have a large, negative price of risk associated with them. Therefore, the low value of π^o in the second period creates a higher persistence, which amplifies the price of risk associated with oil shocks. This price of risk for shocks to the oil price is also important in determining the price of risk for shocks to x_t, due to the extra term in the associated coefficient

$$A_1 = \frac{(1 - \frac{1}{\psi}) + A_2\kappa_1(1 - \eta\Phi_x)}{1 - \kappa_1\rho_x} \tag{26}$$

If $\Phi_x = 0$, shocks to x_t will also be shocks to future growth in oil prices, and if oil prices are persistent A_2 will have a large negative magnitude and the extra term will substantially reduce the price of risk for shocks to x_t. This is an algebraic representation of a very intuitive idea. In a world where oil prices are highly persistent and related to the level of consumption, they can act as a "counterweight" to shocks to expected growth. If high consumption growth is expected then a rise in oil prices is effected as well, which will reduce overall growth.

The highly persistent oil price also represents a new source of risk in the economy through shocks to the oil price, but in my calibrations I find that the reduction of risk from shocks to x_t is a stronger effect, and results in reduced systematic risk, a lower equity premium, and higher price-dividend ratios.

2.3.2 Changes in Loadings for Oil Futures

In order to consider how changes in the consumption parameters affect expected returns on oil futures, we also need to examine how they affect the loadings of oil futures prices on the two shocks. The values of B_x^j and B_p^j are determined by the following recursion.

$$B_x^j = B_x^{j-1}\rho_x + B_p^{j-1}(1 - \eta\Phi_x) \tag{27}$$

$$B_p^j = (1 + \pi^c - \eta\pi^o)B_p^{j-1} \tag{28}$$

With $B_x^0 = 0$, and $B_p^1 = \frac{1}{\rho}$. In the first period, with $\Phi_x = \frac{1}{\eta}$ and a large value of π^o, $B_x^j = 0$ for all maturities and B_p^j decays quickly at higher maturities. In the second period, B_p^j decays more slowly with the higher persistence, and $B_x^j \approx \left(\frac{1}{\rho}\right) j$. Therefore, exposure to shocks to x_t increases linearly across the futures curve.

Remembering that in the second period shocks to oil consumption command a significant, negative, price of risk, it is straightforward to see the source of the "hump" shape term structure in the model. In the second period, expected return is approximately.

$$E[r_t^{j+1}] \approx \frac{-\eta}{\rho}(\lambda_{m,o}\sigma_o^2) + j(\lambda_{m,x}\sigma_x^2) \tag{29}$$

For near term maturities, the first term dominates. This negative expected return from the negative exposure to shocks to o_t remains approximately constant across the term structure due to the slow decay of B_p^j, resulting in an upward sloping term structure at for short maturities. Meanwhile, the second term, representing the exposure to shocks to x_t, generates increasing positive expected returns across the term structure since B_x^j is approximately equal to j. This leads to an increasing downward slope in the term structure, which dominates at longer maturities. This change in slope from negative to positive gives the term structure its characteristic shape.

3 Consumption and Oil Price Dynamics

3.1 Data

Quarterly data for consumption come from the National Income and Product Account (NIPA) tables. Much of the analysis relies on a novel measure of oil consumption, the personal consumption of "Gasoline and other Energy Goods" from the NIPA survey. This measure includes personal consumption of of both gasoline and fuel oils, though in terms of expenditure over 90% of the total comes from expenditure on "Motor Vehicle Fuels, Lubricants, and Fluids" while the remaining 10% is attributed to "Fuel Oil and Other Fuels". Most importantly, this measure is constructed so as not to include consumption for government and corporate use, or consumption of gasoline for energy generation. In this sense it is different from the measure of "Product Supplied" provided by the Energy Information Administration (EIA), which is the typical measure of oil consumption. I divide my measure of personal oil consumption by the level of the population in order to obtain a measure of per capita consumption, as is consistent with literature.

Since gasoline is by far the most important good in this measure, and I am interested in quantifying the utility of consumption, I also adjust for efficiency gains in the use of gasoline, or namely the average miles per gallon. I calculate this using data from the Bureau of Transportation Safety for the average efficiency of the U.S. passenger car fleet. The relative price implied by the agents utility function is then a price for miles rather than a price for gasoline, so I convert it using the miles per gallon to the implied price for oil. For parsimony throughout the description of the model I refer to oil consumption as direct consumption of oil, but for the empirical work I perform these conversions. There is also the potential issue of changes in the efficiency of converting crude to gasoline, but I observe that the price of gasoline and oil have not deviated substantially over the period, and are nearly identical in their innovations, particularly at quarterly frequency.

In order to compare the relative levels of personal consumption of oil to total economic consumption, I construct a measure of total economic expenditure on gasoline and fuel oil using prices and quantities from the EIA. While these are not the only uses of petroleum in the economy, these two sources account for roughly 65% of total product supplied in terms of barrels. The lack of price availability for the remaining products in the EIA measure prevents quantifying the total dollar value, however in terms of expenditure these two components probably account for an even larger percentage since both of these products are more highly refined than many of the other petroleum products and thus command higher prices. Figure 2 shows the two level of expenditures from 1983 to 2010. Personal consumption expenditure of gasoline and fuel oil accounts for a relatively stable share of total economic consumption which varies from 60% to 70%. The fact that a very large portion of total gasoline and fuel oil consumption is accounted for by personal consumption suggests that considering oil as a consumption good rather than an input to production is not an unreasonable approach.

Figure 2: Personal Oil Consumption vs. Total Oil Consumption

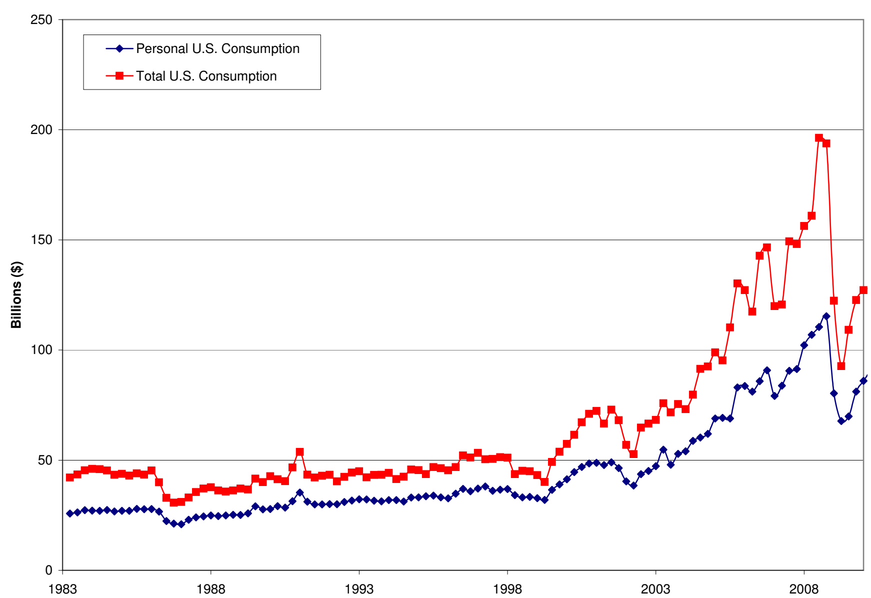

Personal consumption is nominal personal consumption expenditure on "Gasoline and Other Energy Goods" taken from NIPA data. Total oil consumption represents economy wide U.S. oil consumption calculated from prices and quantities of gasoline and fuel oil from the Energy Information Association's report of Product Supplied.

While the consumption based asset pricing literature traditionally relies on nondurables and services as the measure of consumption, recent work by Yogo [2006] and Yang [2010] emphasizes the importance of durable consumption for explaining asset prices. Yang in particular finds that in a long run risk setting, the high persistence of Durable consumption can explain much of the observed equity premium. I follow Yang [2010] and consider consumption as an equally weighted Cobb-Douglas aggregate of the stock of durable goods and expenditure on nondurables and services (excluding energy consumption). I find that this measure does a better job of explaining oil prices than nondurable consumption, and that the added persistence of consumption growth is important in explaining observed features of the futures curve. Following Yogo I construct a quarterly series for the stock of durable consumption using yearly data for the stock of consumer durables and quarterly data for expenditure on durable goods. Data for oil prices is historical data for futures contracts of horizons out to twelve months in Crude Light Sweet oil traded on the NYMEX, and the real spot price of oil is the West Texas Index deflated by a measure of the price of the aggregate consumption good. This price measure is constructed using price levels from the NIPA survey. Table 1 reports summary statistics and correlations for the growth rates of the pertinent data.

Table 1: Growth Rate Summary Statistics

1987 - 2010				Correlations			
Variable	Mean (%)	SD (%)	Autocorrelation	Spot Price	Nondurables	Durables	Aggregate
Real Spot Price of Oil	1.08	17.74	0.03				
Nondurables and Services	0.66	0.38	0.49	-0.10			
Stock of Durable Goods	1.34	0.51	0.95	0.18	0.45		
Cobb-Douglass Aggregate	1.01	0.39	0.85	0.07	0.80	0.90	
Personal Oil Consumption	0.22	1.33	-0.15	-0.42	0.35	-0.05	0.13
1952 - 2010				Correlations			
Variable	Mean (%)	SD (%)	Autocorrelation	Spot Price	Nondurables	Durables	Aggregate
Real Spot Price of Oil	0.51	14.54	0.11				
Nondurables and Services	0.80	0.51	0.31	-0.16			
Stock of Durable Goods	1.21	0.54	0.91	0.07	0.34		
Cobb-Douglass Aggregate	1.01	0.45	0.69	-0.04	0.77	0.86	
Personal Oil Consumption	0.49	1.65	0.01	-0.45	0.46	0.10	0.32

Summary statistics for quarterly growth rates of relevant variables. Cobb-Douglas aggregate is an equally weighted aggregate of the stock of durable goods and the sum of nondurables and services. Nondurable consumption excludes energy goods. The real spot price of oil is calculated as the WTI deflated by the CPI excluding energy.

3.2 Support for Model Specification

3.2.1 Intratemporal Utility

The model as written implies two cointegrating relations. The first, which I will refer to as the Intratemporal relation, arises from the functional form of V_t and implies that a linear combination the two types of consumption, $\frac{1}{\rho}(c_t - \eta o_t)$, will be cointegrated with p_t. This simple version of the model implies that they are in fact equal, but to test this empirically I will test that difference between $p_t - \frac{1}{\rho}(c_t - \eta o_t)$ is a stationary process. I find strong evidence that this is the case, and that not only is the difference a stationary process, but that the predicted spot price $\frac{1}{\rho}(c_t - \eta o_t)$ provides an excellent proxy for the real spot price of oil. This result is crucial for motivating the model, since the consumption dynamics can only have meaningful implications for oil prices if there exists a relation between levels of consumption and the spot price of oil. Documenting the existence and strength of this relation is one of the main empirical contributions of this paper, and provides a starting point for which to consider the relation between consumption and oil price risk.

Cointegration analysis is a common tool in the study oil or gasoline prices. Several studies such as Bentzen and Engsted [1993]and Ramanathan [1999] seek to estimate both long run and short run elasticities of consumption to prices using methods similar to those I use here. Typically these analyses begin by proposing a demand function for oil where the log of economy-wide oil or gasoline consumption is assumed to be a linear function of the logs of other economic variables, most often personal income and the price of oil. Since I am interested in pricing assets in the consumption based long run risk framework, the relation I focus on involves personal consumption, o_t, and is implied by the first order condition of an optimizing representative agent with utility over two goods.

I follow Yogo [2005] and estimate a cointegrating relation between the log of oil prices and measures of consumption and oil consumption. A simple method for doing this is the Dynamic OLS method described by Stock and Watson (1993), where equation (3) is estimated, including both leads and lags of the dependent variables, resulting in the following form for the regression.

$$p_t = \beta_0 + \beta_1 c_t + \beta_2 o_t + \sum_{t=-k}^{k} \Gamma_{1,k} \Delta c_{t+k} + \sum_{t=-k}^{k} \Gamma_{2,k} \Delta o_{t+k} \tag{30}$$

The coefficients are related to the parameters of the utility function V_t by $\beta_1 = \frac{1}{\rho}$ and $\beta_2 = \frac{\eta}{\rho}$ This regression model is identical to that of Bentzen, with personal aggregate consumption and personal oil consumption standing in for personal income and economy wide oil consumption. It is worthwhile to note here the implications of considering oil directly as a consumption good. While clearly consumers do not consume crude oil, and ultimately I will be concerned with pricing futures contracts for delivery of crude oil, there is a very tight relation between crude oil prices and the price of gasoline, which does directly enter the consumer's consumption basket. More importantly, I find that data on personal consumption of oil products taken from Bureau of Economic Analysis' NIPA tables, provide substantial improvement in explanatory power for oil prices over typical measures of crude oil and gasoline consumption taken from the Energy Information Association (EIA). Aggregate consumption performs equally as well as personal income in predicting oil prices.

When doing the regressions with consumption, I divide the consumption data by estimates of the U.S. population taken from census data. In order to account for changes in the efficiency of converting oil to consumption utility, I adjust the level of oil consumption by the multiplying it by average miles per gallon taken from the Bureau of Transportation Statistics. The assumption underlying this adjustment is that the consumption good is not actually gasoline, but rather miles driven. Therefore, I also adjust the price of oil by miles per gallon. Therefore, in the regression of Equation (30), I substitute p_t with $(p_t - \log(mpg_t))$, and o_t with $(o_t + \log(mpg_t))$. These adjustments do not add significant volatility to the series, but they do adjust the growth trends, which are important in determining the cointegrating relation. I do not perform these adjustments when using the data for economy wide oil consumption to be consistent with other studies, however performing this adjustment for this data does not significantly improve the estimates.

I estimate this regression using two different measures of aggregate consumption, both consumption of nondurable goods and services and a Cobb-Douglas aggregate of nondurable goods and services and the stock of durable goods constructed as in Yogo [2006]. I also include two different measures of the consumption of oil. The first, following Bentzen and others, is the economy-wide measure of product supplied from the EIA, the second is the measure of energy product consumption (including gasoline and heating oil) from NIPA consumption data. For comparison I also estimate the regression using personal income and GDP in place of consumption. Table 2 reports these regressions for 1987 to 2010, the period for which I have futures data, as well as regressions of the oil price on levels and leads and lags of each variable individually.

Table 2: Cointegration of Oil Prices and Economic Variables

Dep. Var.	C	c_t	n_t	$log(GDP_t)$	$log(I_t)$	$o_t^{Personal}$	o_t^{All}	*Adj.* R^2
Single Variable Regressions								
p_t	-4.615	0.709						0.703
	0.36	*7.96***						
p_t	-4.866		0.850					0.483
	0.73		*4.79***					
p_t	-9.190			0.831				0.471
	1.62			*4.73***				
p_t	-7.401				1.589			0.416
	0.75				*7.35***			
p_t	-2.781					1.072		0.468
	0.35					*3.02***		
p_t	-29.535						1.935	0.499
	*5.75***						*4.81***	
Two Variable Regressions: NIPA Personal Energy Good Consumption								
$p_t - \ln(mpg_t)$	-4.9995	4.04				-7.05		0.91
	0.31	*10.19***				*-8.96***		
$p_t - \ln(mpg_t)$	-16.988		6.575			-12.071		0.761
	1.42		*9.98***			*-8.66***		
$p_t - \ln(mpg_t)$	-51.383			6.788		-13.038		0.755
	*4.71***			*10.20***		*-8.82***		
$p_t - \ln(mpg_t)$	-12.189				5.114	-8.210		0.846
	0.67				*14.70***	*-12.12***		
Two Variable Regressions: EIA Product Supplied								
p_t	-28.824	0.392					1.796	0.699
	*36.20***	*0.73*					*0.66*	
p_t	-42.802		-0.455				3.014	0.507
	*23.30***		*-0.58*				*1.63*	
p_t	-35.572			-0.234			2.519	0.489
	*21.28***			*-0.28*			*1.26*	
p_t	57.692				3.173		-4.954	0.564
	*21.83***				*4.10***		*-2.89***	

Estimation of Stock and Watson (1993) regressions of log real spot price on logs of economic variables. Estimations are done with two leads and lags as well as contemporaneous differences. Coefficients on difference terms are suppressed. Standard errors are Newey-West with two lags. c_t is the log of the aggregation of durables and nondurables, n_t is log of nondurable consumption expenditure , log (I_t) is personal income taken from the NIPA tables. $o_t^{Personal}$ is personal oil consumption of energy goods taken from the NIPA tables adjusted for by miles per gallon, o_t^{All} is the measure of oil ”Product Supplied” taken from EIA data. All variables are measured per capita.

The two things to note in this table are that the measurement of oil consumption from NIPA data does a much better job of explaining oil prices than the measure of consumption obtained from the EIA. Secondly is that, in terms of R^2, the consumption aggregate of Durable and Non Durable goods explains prices equally as well as personal income and slightly better than Non Durable goods alone. Augmented Dickey-Fuller tests (not reported) of the residuals of these regressions strongly reject the presence of a Unit Root, indicating a cointegrating relation between oil prices and the economic variables. In order to illustrate the goodness of fit of this model Figure 3 graphs the predicted values from a simple regression of the log of the oil prices on the logs of aggregate consumption and energy consumption from 1965 to 2010. The estimates of these simple regressions on the longer sample are not statistically different from the estimates from the Stock and Watson regressions on the shorter sample. This strong evidence of a relation between oil prices and consumption suggests that it is reasonable to start with a model of consumption dynamics when considering the behavior of oil prices.

Figure 3: Approximating the Spot Price of Oil with Levels of Personal Consumption

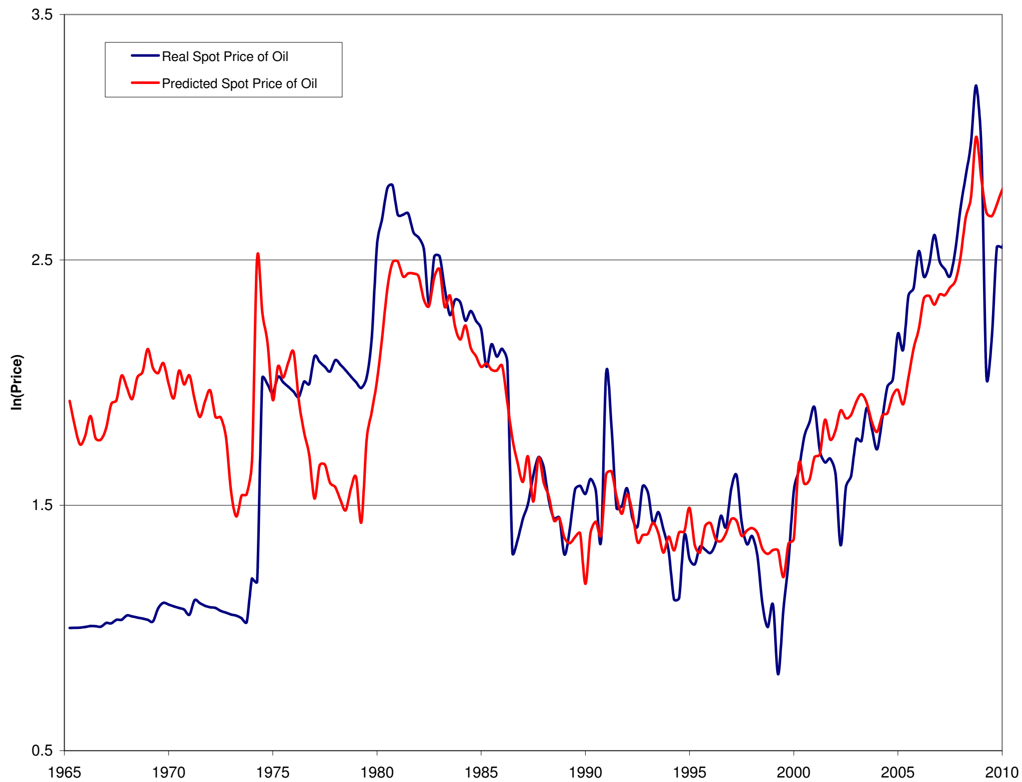

Predicted prices are the predicted value from the regression $(p_t - \log(mpg_t)) = \beta_0 + \beta_1 c_t + \beta_2 (o_t + \log(mpg_t)) + \epsilon_t$, where p_t is the log of the WTI spot price adjusted by CPI excluding energy costs, c_t is a CES aggregation of the stock of durable consumption and expenditure nondurable consumption (excluding energy goods), and o_t is the measure of energy good consumption from the NIPA survey. Consumption measures are adjusted by the U.S. Population. mpg_t is the average miles per gallon of the U.S. passenger car fleet taken from the Bureau of Transportation Statistics.

3.2.2 Consumption Dynamics

The first observation I make concerning consumption dynamics is the ability of oil shocks to predict future consumption growth. Hamilton [2008] shows that regressing GDP growth on lagged innovations to oil prices from 1972 - 2005 indicates that positive oil price increases negatively predict future GDP growth. I perform identical regressions using my measure of aggregate consumption in place for 1972 - 2010 and confirm this result. Results are reported in Table 3. Therefore, the result that consumption growth predicts negative aggregate consumption is not unique to my choice of sample period.

Table 3: Oil Price Shocks and Consumption Growth: 1972 - 2010

Δc_t	μ_c	Δc_{t-1}	Δp_{t-1}	Δp_{t-2}	Δp_{t-3}	Δp_{t-4}	R^2
c_t	0.001	0.915**	-0.0020*	-0.0004	-0.0020**	-0.0007	86.3%
	0.005	*0.037*	*0.001*	*0.001*	*0.001*	*0.001*	

Regressions of growth of Cobb-Douglas aggregate of durable and nondurable consumption on lagged aggregate consumption growth and oil price innovations. p_t is the log of the real spot price, as measured by the WTI spot price deflated by the CPI. Data is quarterly frequency. Standard Errors are Newey-West with three lags.

I next estimate the parameters which govern the dynamics of consumption in the model. In order to do this, I first need a value of expected consumption growth x_t. I estimate expected consumption growth as the predicted value implied by a regression of aggregate consumption on the three lags of durable and nondurable consumption growth.

$$\Delta c_t = \sum_{i=1}^{3} \beta_i^d \Delta d_{t-i} + \sum_{i=1}^{3} \beta_i^n \Delta n_{t-i} + \epsilon_t \tag{31}$$

I use the estimate of η from the Stock and Watson regressions and estimate the following two regressions to obtain estimates of Φ_x^c, Φ_x^o, π^c, and π^o. In the model x_t will be normalized so that $\Phi_x^c = 1$, leaving only a single parameter Φ_x.

$$\Delta c_{t+1} = \pi^c (c_t - \eta o_t) + \Phi_x^c x_t + \sigma_c e_t^c \tag{32}$$

$$\Delta o_{t+1} = \pi^o (c_t - \eta o_t) + \Phi_x^o x_t + \sigma_o e_t^o \tag{33}$$

Estimates for the two periods are in Table 4. The estimates illustrate how the changes in price dynamics are reflected in the changes in consumption dynamics. The estimate for π^c is positive by not significant in the first period, and negative and significant in the second period. Given the significant negative predictive power of the oil price for consumption in the second period, and the results from the regression of lagged oil price innovations on consumption growth for the longer time horizon, I will set the value of π^c to be equal and negative across the two calibrations, and focus on the impact of the changes in the parameters governing oil consumption. The estimate for π^o is significantly positive in the first period while not significantly different from zero in the second period. The estimate for Φ_x^o is significantly positive in the first period, and not significantly different from zero in the second.

Table 4: Consumption Dynamics

Parameter	**1987-1999**	**2000-2010**
Regression:	$\Delta c_{t+1} = \pi^c(c_t - \eta o_t) + \Phi_x^c x_t + e_{t+1}^c$	
π^c	0.002	-0.006^*
	(0.002)	*(0.001)*
Φ_x^c	1	1
	(0.08)	*(0.02)*
Regression:	$\Delta o_{t+1} = \pi^o(c_t - \eta o_t) + \Phi_x^o x_t + e_{t+1}^o$	
π^o	0.081^*	-0.008
	(0.01)	*(0.01)*
Φ_x^o	$1.36^\dagger$	-0.31
	(0.81)	*(0.88)*

Estimates of regressions of consumption growth on a predictable consumption growth x_t and the error correction term $c_t - \eta o_t$. x_t is estimated as the predicted value from a regression of aggregate consumption growth on three lags of durable and nondurable consumption growth. Data is quarterly frequency. Standard Errors are Newey-West with three lags.

Expected log oil prices can be expressed in the model as

$$E_t[p_{t+1}] = (1 + \pi^c - \eta\pi^o)p_t + (1 - \Phi_x\eta)x_t \tag{34}$$

As is shown by this relation, the changes in consumption dynamics reflect the changes in spot price dynamics. The decrease in π^o leads to a larger value for the AR(1) coefficient for spot prices, $(1 + \pi^c - \eta\pi^o)$, and hence less mean reversion in oil prices. The change in Φ_x leads to differences in how expected spot prices respond to changes in x_t. In the context of the model, both of these changes have significant effects on the expected returns to oil futures and are discussed in more detail in the following section.

3.3 Changes in Oil Prices

The time period I am focusing on for this analysis is the 23 year period from 1987 to 2010 for which I have data on futures prices out to 12-month horizons. Figure 4 graphs the real oil price as well as my measures of aggregate consumption and oil consumption from 1947 to 2010. Following the Oil Price Crash of 1986, there is roughly 12 year period of remarkable stability in oil prices. Around the end of the 1990s, prices began to rise, and though they fell again in the early part of the next decade, they continued their rise again, increasing by 400% over the next 8 years, before falling sharply following the financial crisis of 2008. Therefore, though this period is dictated by the availability of data, even in the absence of this constraint this time period is a potentially interesting one.

Figure 4: Oil Prices, Aggregate Consumption, and Oil Consumption

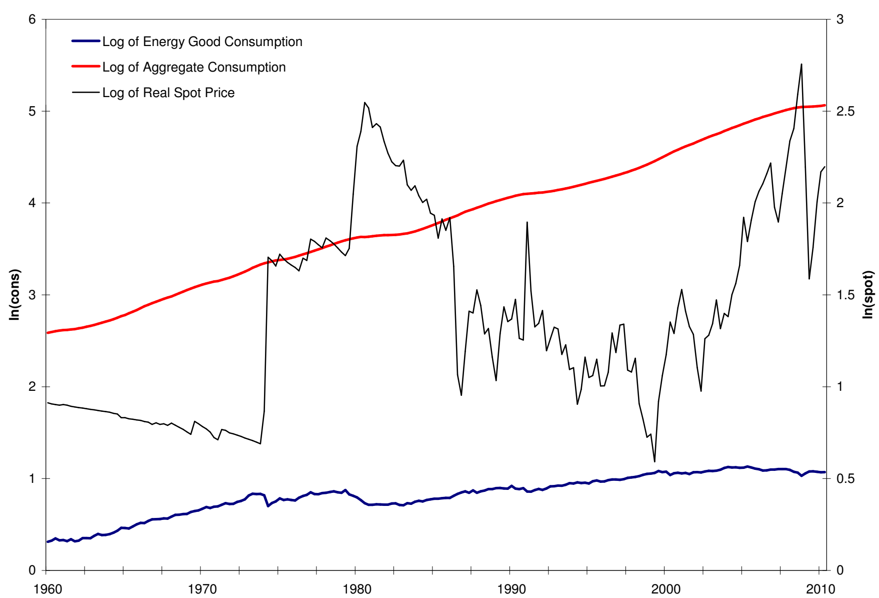

Oil consumption is personal consumption of "Gasoline and Other Energy Goods" taken from NIPA data. Aggregate consumption is an aggregate of durable and nondurable consumption excluding oil consumption from NIPA data. The real price of oil is the WTI spot price of oil adjusted using CPI index of All Items Less Energy.

3.3.1 Changes in the Persistence of Oil Prices

The existence of mean reversion in oil prices is a topic which has received substantial attention in the macroeconomic literature. Some studies, such as Routledge et al. [2000] and Schwartz [1997], find evidence of mean reversion in oil using high frequency data in the 1990s. However, Hamilton [2008] describes oil prices over the period of 1973 to 2008 as a pure random walk based on the results of an Augmented Dickey-Fuller test using quarterly data. More recently, Dvir and Rogoff [2009] employ the test of Harvey et al. [2006] to detect structural changes in the price of oil from an I(0) to an I(1) process and vice versa. They examine a much longer horizon, and test for a single change in behavior from 1881 to 2008 and find evidence of a change of oil from an I(0) to an I(1) process in 1973.

There are two ways which I test for changes in persistence. The first is by directly testing the log of the monthly spot price for the existence of a unit root. The second is to employ a regression technique similar to that of Bessembinder et al. [1995], whereby the changes in long term futures prices are regressed on the innovations in the spot price. High mean reversion should imply the longer term contract moves less in response to a change in the short term contract, as in the long term prices are expected to mean revert.

To formally test for a change in the behavior of spot prices, I use the test of Busetti and Taylor [2004] and find evidence for a switch from I(0) to I(1) at the beginning of 2000. I also use the test of Bai and Perron [1998] to test for the change in the exposure of returns on the 12 month futures to changes in the spot price, and find evidence for a structural break around the end of 2002. I split the sample at the beginning of 2000 to consider differences in the dynamics of the spot price of oil and the dynamics of consumption. Since my model will be calibrated for two regimes without describing the dynamics around a regime switch, for futures prices I consider two slightly different subperiods. I consider the period prior to the first structural break, 1987 - 1999, and then after the second structural break, 2002 - 2010. Simply splitting the data at either point yields the same qualitative results, but this method allows for asset markets to fully incorporate the changes in consumption dynamics. Rather than reporting the results of the structural break tests, for each subperiod I report the the standard unit root tests and my regression results which are more easily interpretable.

Table 5 shows standard unit root tests of unit root tests for the log of the real spot price. For the first sample both Augmented Dickey Fuller Tests and Phillips-Perron tests reject the null that oil prices contain a unit root, while for the second there is no evidence to reject a unit root. The table also reports the first order autocorrelation of the log of the spot price, estimated from an AR(1) regression. The autocorrelation is significantly higher in the second period.

Table 5: Unit Root Tests of Oil Spot Prices

Statistic		**1987 - 1999**	**2000 - 2010**	**1987-2010**
Augmented Dickey-Fuller	Z(t)	-2.94^{*}	-1.55	-1.23
	P value	0.04	0.51	0.63
Phillips-Perron	Z(rho)	$-14.78^{\dagger}$	-3.73	-3.41
	P value	0.08	0.57	0.75
AC(1) p_t	Est	0.72	0.91	0.93
	Std Err	0.10	0.09	0.06

Phillips-Perron tests and Augmented Dickey Fuller tests for a unit root in p_t, the log of the WTI spot using monthly data. P-values are in parentheses. A ($*$) or ($\dagger$) denote rejection of a unit root at the 5% and 10% significance levels.

From the perspective of an asset pricing model, such as the one presented in this paper, it is agents expectations of the persistence of oil that is the important determinant for prices of risk. In order to examine changes in expected persistence, I employ a regression technique similar to that of Bessembinder et al. [1995]. They observe that for commodities futures, higher mean reversion implies that futures of longer maturities move less in response to changes in the spot price, and test this implication for several commodities by regressing contemporaneous movements of futures horizons at different maturities on the spot price. I modify this slightly and consider the returns of longer term futures contracts and their comovement with shorter term contracts.

I follow convention and define the excess return on a futures contract with j months to maturity as.

$$r_{t+1}^{j} = f_{t+1}^{t+j-1} - f_{t}^{t+j} \tag{35}$$

Given my data for futures out to 12 months, I have observations for returns of futures with horizons from two months to 12 months. I ignore the return on the nearest term futures price to avoid issues of high volatility as the contract gets close to delivery. I perform a simple regression of the futures return at each maturity on the contemporaneous change in spot price.

$$r_{t+1}^{j} = \gamma_0^j + \gamma_1^j \Delta p_{t+1} + \epsilon_{t+1} \tag{36}$$

Results of this regression for each maturity in the two periods are reported in Table 6. These regressions confirm that the realized differences in persistence also lead to changes in expected persistence, with a coefficient of the longest term contract changing from 0.5 in the first period to 0.7 in the second.

3.3.2 Changes in the Term Structure of Futures Returns

While the run up in prices over the second half of the sample was well publicized, what has not been as closely studied is the difference in the term structure of futures over these two periods. Panel A

Table 6: Regressions of Returns on Changes in Spot Price

	1987 - 1999		2002 - 2010	
	γ_0	γ_1	γ_0	γ_1
r_t^2	0.006	0.97	-0.004	0.98
	(0.00)	*(0.02)*	*(0.00)*	*(0.02)*
r_t^3	0.006	0.88	0.000	0.93
	(0.00)	*(0.02)*	*(0.00)*	*(0.02)*
r_t^4	0.005	0.81	0.003	0.89
	(0.00)	*(0.02)*	*(0.00)*	*(0.02)*
r_t^5	0.005	0.75	0.004	0.85
	(0.00)	*(0.02)*	*(0.00)*	*(0.02)*
r_t^6	0.005	0.69	0.005	0.82
	(0.00)	*(0.02)*	*(0.00)*	*(0.03)*
r_t^7	0.004	0.65	0.006	0.79
	(0.00)	*(0.02)*	*(0.00)*	*(0.03)*
r_t^8	0.004	0.60	0.006	0.76
	(0.00)	*(0.02)*	*(0.00)*	*(0.03)*
r_t^9	0.004	0.57	0.007	0.74
	(0.00)	*(0.02)*	*(0.00)*	*(0.03)*
r_t^{10}	0.003	0.53	0.007	0.71
	(0.00)	*(0.02)*	*(0.00)*	*(0.03)*
r_t^{11}	0.003	0.51	0.007	0.69
	(0.00)	*(0.02)*	*(0.00)*	*(0.03)*
r_t^{12}	0.003	0.48	0.007	0.67
	(0.00)	*(0.03)*	*(0.00)*	*(0.03)*

Results for regressions of the form $r_{t+1}^j = \gamma_0^j + \gamma_1^j \Delta p_{t+1} + \epsilon_{t+1}$. Returns are monthly returns on the NYMEX futures for different horizons. Standard errors are in parentheses.

of Figure 5 graphs the average term structure of futures for each subperiod. What is noteworthy here is the development of a "hump" shape in the term structure of futures over the latter half of the sample. While this change may seem of little significance, the change in curvature from a concave curve to a convex curve has important implications for expected returns and hence risk premia. The difference in futures prices for contracts at adjacent months can be expressed as

$$f_t^{t+j} - f_t^{t+j-1} = -E_t[r_{t+1}^j] + E_t[f_{t+1}^{t+j-1} - f_t^{t+j-1}] \tag{37}$$

This difference is decomposed into two pieces, the expected return, and the expected change in the futures price for a contract maturing at date $t+j-1$. Therefore, one possible candidate for explaining changes in the term structure of prices is changes in expected returns. Panel B in Figure 5 graphs average log returns over the two subperiods. Returns are increasing across the term structure of futures in the second period, as opposed to decreasing in the first. Holding the second term of Equation (37) constant, the decreasing expected returns of the first period imply a convex term structure of future prices, while the increasing expected returns of the second period imply a concave term structure, which is precisely what we see.

Figure 5: The Term Structure of Crude Oil Futures

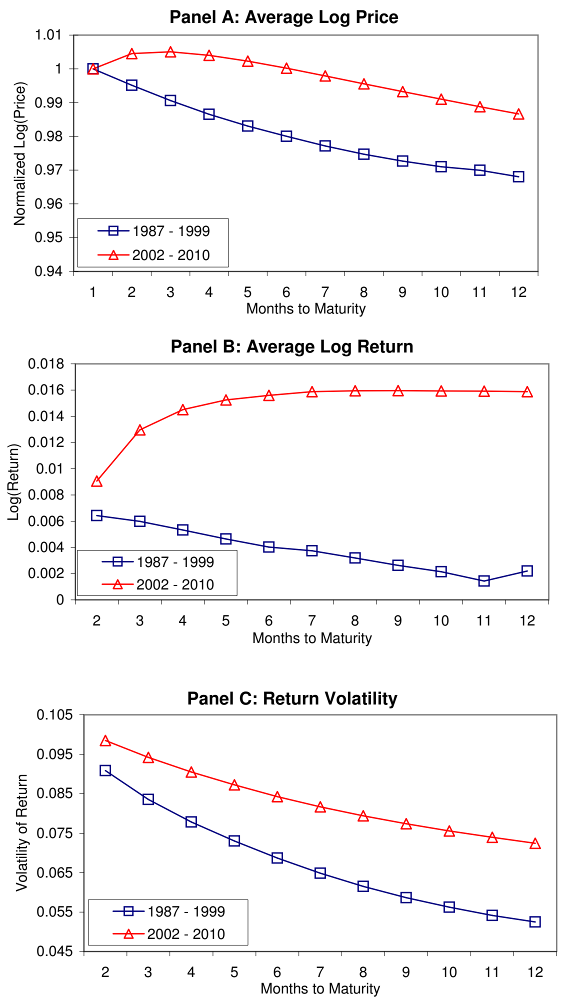

Panel A reports the average log of futures prices for the two halves of the sample. Both curves are normalized so that $\bar{f}^1 = 1$. Panel B reports the averages of monthly returns for futures prices. Panel C reports monthly volatility of oil returns. Data for NYMEX futures prices on Crude Light Sweet Oil of up to 12 months to maturity.

While the returns to commodity futures are highly volatile like any asset, they are also highly correlated with futures prices at other maturities. Therefore, when examining differences between levels of expected return across the term structure, the relative returns are considerably less volatile than the absolute returns, and inference can be made at much shorter time horizons than would normally be required when considering the return on a single asset. This is especially important in this setting, as I am interested in making statements about changes in risk premia using merely 10 years of data.

This feature of futures prices, that the added dimension of returns across the term structure gives extra power in identifying changes in the pattern of expected returns, has been mostly overlooked in the literature. Many studies, such as Fama and French [1987] and Gorton et al. [2007b], examine the futures basis, or the "slope" of the futures term structure, as a possible predictor of either changes in spot price or returns on the nearest futures contract. While these are obviously related issues to this analysis, they are focused on explaining the return to the contract of a single maturity, rather than studying the term structure of expected returns.

In order to assign statistical significance to the observed differences of returns across I estimate the following simple regression of expected returns on the maturity of the futures contract.

$$E[r_t^j] = \beta_0 + \beta_j j + \epsilon_j \tag{38}$$

I estimate the coefficients using the Fama and MacBeth [1973] procedure. While this procedure capitalizes on the comovement in returns by essentially allowing for a time fixed effect, it does not account for the fact that when prices are rising, the short end of the futures curve tends to increase more than the longer term contracts as evidenced by Table 6. This effect creates larger standard errors in this setting. In order to control for it, I define the following return

$$\tilde{r}_t^j = r_t^j - \tilde{\gamma}_j r_t^{12} \tag{39}$$

This is the return to a strategy of going long on a short term contract j, and short a proportional position in the 12 month contract. The proportion, $\tilde{\gamma}_j$ is determined by a 3 year rolling regression of r_t^j on r_t^{12}. I then repeat the regression of equation 38. Results for the two regressions are reported in Table 7. For the basic regression the positive slope in the second period is significant with a p-value of 6%. For the regression using $\tilde{r}_t^j$, this positive slope is highly significant with a p-value of 1%. The negative slope in the first period is not statistically significant at any conventional level for either regression.

Table 7: Fama-MacBeth Regressions of Futures Returns

	Dep Var	Constant	Maturity
1987 - 1999	r_t^j	0.920	-0.051
		(9.24)	*(0.05)*
	$\tilde{r}_t^j$	0.202	-0.030
		(0.34)	*(0.05)*
2002 - 2010	r_t^j	0.682	$0.080^{\dagger}$
		(9.82)	*(0.04)*
	$\tilde{r}_t^j$	-0.903^{*}	0.112^{*}
		(0.35)	*(0.04)*

Regressions of expected return on maturity. r_t^j is the return on future of maturity j. $\tilde{r}_t^j$ is the return on maturity j controlling for the return on the 12 month maturity. Data is monthly. Errors are computed using the Fama-Macbeth procedure.

3.4 Volatility and the Term Structure

Across these two periods there are also changes in the conditional relation of volatility and the slope of the futures curve. The model presented here will rely on changes in the riskiness of futures from both expected aggregate consumption growth and oil consumption growth to explain the hump shape in the term structure of futures. Exposure to oil consumption shocks will generate an upward sloping term structure, and exposure to expected aggregate consumption growth will generate a downward sloping term structure. Since, in the model, these effects are conditionally stronger in periods of higher volatility, it implies that changes in dynamics across the two periods will imply changes in the conditional relation of the volatilities of these two shocks and the slope of the futures curve[6].

Here another potential benefit of having the futures curve to measure changes in risk premia associated with changes in volatility is that the comparison of futures prices of different maturities controls for changes in the level of prices that may accompany shocks to volatility. For example, the tendency for equity prices to decline when option implied volatility increases might be explained by either a positive shock to volatility causing an increase in the required rate of return on equities, or by a negative shock to expected future cash flows that also results in an increase in volatility, possibly due to a leverage effect.[7]

In the model, the state variables will be volatilities which are difficult to directly observe, since consumption growth available at only quarterly frequency, I instead consider volatilities of aggregate equity returns and spot prices for which I have analogs in the model. In each month I calculate the monthly return volatility implied by the volatility of daily returns on the S&P index and daily changes in spot prices, $\sigma_{S\&P,t}$ and $\sigma_{spot,t}$. I then perform predictive regressions to estimate an expected volatility under the physical measure in each month. Following Drechsler and Yaron [2009] I use the lag of the CBOE VIX index and one lag of $\sigma_{S\&P,t}$ to calculate expected market return volatility. I use three lags of $\sigma_{spot,t}$ to estimate an expected volatility of spot prices. This gives me a time series of expected volatilities, $E_t[\sigma_{S\&P,t+1}]$ and $E_t[\sigma_{spot,t+1}]$.

I then perform the following regression in each half the sample to examine the relation of these expected volatilities and the slope of the term structure, which is defined as the difference between the log of the 12-month future and the 1-month future.

[6]One of these effects has been noted by Singleton [2008] who observes that over the recent period, times of high volatility tend to coincide with a futures curve in "contango", that is having a positive slope.

[7]Eraker [2008] provides evidence that suggests the observed negative correlation between equity prices and implied volatilities can be explained by changes in required rates of return.

$$f_t^{12} - f_t^1 = \beta_0 + \beta_{\sigma,S\&P} E_t[\sigma_{S\&P,t+1}] + \beta_{\sigma,spot} E_t[\sigma_{spot,t+1}] + \sum_{i=0}^{L} \beta_p^i \Delta p_{t-i} \tag{40}$$

Lagged changes in spot prices are included to capture variation in the slope caused by movements in spot prices and the mean reverting nature of oil. Results are reported in Panel A of Table 8. Not surprisingly, given the decrease in mean reversion, the lagged price movements have less effect on the slope of the futures curve in the second period. More important is the change in the relation between the slope of the futures curve and volatility. In the first period, volatility has little impact on the slope of the term structure. In the second period, both expected volatility of stock prices and expected volatility of spot prices are significant in explaining the slope of the futures curve. High expected equity volatility coincides with a more upward sloping term structure, and while high oil price volatility coincides with a downward sloping term structure.

Table 8: Regressions of Volatility and the Futures Curve

Panel A: Expected Equity Volatility, Spot Volatility, and the Slope

Period	**Dep. Var.**	$E_t[\sigma_{spot,t+1}]$	$E_t[\sigma_{S\&P,t+1}]$	Δp_t	Δp_{t-1}	Δp_{t-2}	**Constant**	R^2
1987	$f_t^{12} - f_t^1$	-0.72	0.11	-0.40^*	-0.32^*	-0.3^*	0.04	0.43
- 1999		*(0.44)*	*(0.45)*	*(0.10)*	*(0.07)*	*(0.07)*	*(0.03)*	
2002	$f_t^{12} - f_t^1$	2.21*	-1.66^*	-0.23^*	-0.10	-0.04	-0.17	0.33
- 2010		*(0.56)*	*(0.60)*	*(0.10)*	*(0.10)*	*(0.12)*	*(0.06)*	

Panel B: The Absolute Value of the Slope and Spot Volatility

	Dep. Var	$\lvert f_{t-1}^{12} - f_{t-1}^1\rvert$	$\sigma_{p,t-1}$	VIX_{t-1}	**Constant**	R^2
1987 - 1999	$\sigma_{p,t}$	4.633*	0.32*	0.09*	0.54*	0.38
		(1.22)	*(0.05)*	*(0.04)*	*(0.13)*	
2002 - 2010	$\sigma_{p,t}$	0.033	0.38*	0.38*	0.53	0.12
		(1.21)	*(0.11)*	*(0.14)*	*(0.30)*	

Panel A are regressions of the slope of the futures curve on the expected spot price volatility and the expected volatility of equity prices. Expected volatility of equity prices is calculated following Drechsler and Yaron [2009] using a regression of realized daily volatility of the returns of the S&P 500 on the lag of realized volatility and the CBOE VIX index. Expected spot price volatility is calculated using the lag of the volatility of daily changes in the WTI index. The slope is the log difference between the twelve month futures price and the 1-month futures price. Panel B are regressions of the realized daily volatility of spot prices on its lag, the lag of the VIX index, and the lag of the absolute value of the slope of the futures curve.

Kogan et al. [2009] also consider the conditional relation of spot price volatility to the absolute value of the slope of the futures curve. They find that, over the the period 1985 - 2001, when the futures curve has either a large positive slope or a large negative slope, it predicts high volatility of spot prices. They explain this effect with a production model with constraints on the adjustment of supply in each period. When the producer is adjusting supply to respond to a price shock, the adjustment constraint is binding and they are unable to respond to further changes in prices. Though this mechanism is is not present in my model, it is worth noting that if production is no longer able to respond to prices at all, there will be no changing elasticity of supply and this effect will disappear. Panel B of Table 8 repeats their regression for each half of the sample. Their results are confirmed in the first half of the sample, but are no longer present in the second half of the sample. The fact that it no longer exists in the second period is both consistent with their explanation of this result, and evidence for a lack of production response as a potential explanation for the changes in consumption dynamics that I focus on here.

4 Extensions to the Model

Before taking the model to the data, I extend it in two ways to help match observed behavior of oil prices. The model as given does not generate a downward sloping curve in the first period. I therefore add a constant drift in the long term price of oil $\bar{p}_t$. This captures the notion of an average convenience yield for oil prices. When calibrating the model I fix this value to be constant across the two periods. Though there is theoretical evidence that convenience yields depend on the level of storage, and storage levels are indeed lower during the second period, I find that the changes in riskiness of futures contracts are sufficient to explain observed changes in futures curves. Moreover, changes in convenience yields can not explain the differences in expected return across these two periods, so I hold the convenience yield constant and focus the effects from the changes in consumption dynamics.

I also note that the price of oil tends to be above the model predicted price when prices are rising, and vice versa. I define ξ_t to be the difference between the observed price of oil $\hat{p}_t$ and the price implied by the agents F.O.C., $\frac{1}{\rho}(c_t - \eta o_t)$, where ρ and η and taken from the original Stock and Watson regression. I perform the following regression.

$$\xi_t = \alpha_\xi + \sum_{i=0}^{n} \beta_{\xi,i} \Delta(c_{t-i} - \eta o_{t-i}) + \epsilon_{\xi,t} \tag{41}$$

The results of this regression are shown in Table 9. Adding changes in the relative level of consumption provides significant extra explanatory power to explain prices. To reflect the results of this regression, I redefine price to be

$$p_t = (c_t + \eta o_t) + \xi_t \tag{42}$$

$$\xi_{t+1} = \rho_\xi \xi_t + \frac{1-\epsilon}{\epsilon} \Delta(c_{t+1} - \eta o_{t+1}) \tag{43}$$

Table 9: Observed Oil Price and Innovations to Implied Oil Price

	ξ_0	$\Delta(c_t - \eta o_t)$	$\Delta(c_{t-1} - \eta o_{t-1})$	$\Delta(c_{t-2} - \eta o_{t-2})$	R^2
ξ_t	-0.019	1.55*	1.42*	0.97	12.2%
	(0.03)	*(0.72)*	*(0.70)*	*(0.63)*	

ξ_t is the observed difference between the spot price of oil and the value $\frac{1}{\rho}(c_t - \eta o_t)$. Results are reported for a regression of ξ_t on lags of innovations to the value of $c_t - \eta o_t$. Data is quarterly frequency. Standard Errors are Newey-West with three lags.

Economically this empirical result potentially reflects a habit in oil consumption. This behavior in prices can be created in the model when ξ_t is the inverse of a relative external habit for oil consumption, similar to Ravn et al. [2005], so that utility is given by.

$$V_t(C_t, O_t) = \left[(1-a)C_t^{1-\frac{1}{\rho}} + a\left(\frac{O_t}{X_t}\right)^{1-\frac{\eta}{\rho}}\right]^{\frac{\rho}{\rho-1}} \tag{44}$$

Where

$$X_t = e^{-\xi_t} \tag{45}$$

Though habits are usually thought to evolve according to only the innovations of a single consumption good, given the high complementarity of oil to aggregate consumption, it is reasonable to assume that innovations to aggregate consumption also will affect the level of habit. Or equivalently that the habit is in effect a habit relating to the level of oil consumption relative to aggregate consumption rather than the absolute level of oil consumption.

When solving the model, including even an external habit potentially complicates the solution for the agent's pricing kernel. However, the approximation utilized to cope with the generalized CES intratemporal utility is also a suitable approximation to an intratemporal utility with a habit for oil due to the small ratio of oil consumption expenditure to aggregate consumption expenditure. Again this highlights the point that the interesting dynamics of risk in this model come from the relation of oil to aggregate consumption growth, rather than the fact that oil consumption directly enters the utility function. It also means that this extension can be thought of as a "reduced form" model, where consumer utility comes only from the aggregate consumption good and the price of oil is given exogenously by Equation 42. Setting the parameter $\tau = 0$ gives this interpretation, and provides nearly identical calibrations to the ones shown.

This specification of price has little qualitative effect on patterns of returns in the model. To show how this extension affects the exposure of oil prices to the underlying shocks, Figure 6 shows impulse response functions with the extended formulation of prices. The patterns are qualitatively the same as in Figure 1, but the magnitudes of exposure are larger. Empirically, adding this allows for a better fit of prices as evidenced by the regression in Table 9, and also allows the model to better match magnitudes of observed expected returns.

5 Model Calibration

I calibrate two different scenarios for the model to match observed moments from the two halves of the sample. Following convention in the literature I set the risk premia γ equal to 10, and the IES ψ equal to 1.5. I constrain Φ_x to be zero in the second period, while setting $\Phi_x = \frac{1}{\eta}$ in the first period so that long run consumption growth has no impact on the expectation of the long run oil price. I set $\frac{1}{\rho} = 4$ and $\eta = 1.75$ to match the values from the intratemporal cointegrating regression. Remaining parameters are chosen to match important model moments of oil prices, most notably the persistence and volatility of aggregate consumption and oil consumption, and the volatility of prices.

Table 10 provides parameters for the estimated models. Table 11 gives sample moments and model moments of consumption dynamics. Data moments are reported for both 1957 - 2010, the period over which I have available quarterly data, and 1987 - 2010, the data sample period. Table 12 reports sample and model moments for oil prices. For this table oil price data moments are for the two haves the sample for which I have futures data, 1987 - 1999, and 2000 - 2010. Table 13 reports asset moments for the both 1957 - 2010, and 1987 - 2010.

Figure 6: Model Impulse Response Functions: Extended Model

Period 1: $\pi_o = .1$ and $\Phi_x = \frac{1}{\eta}$

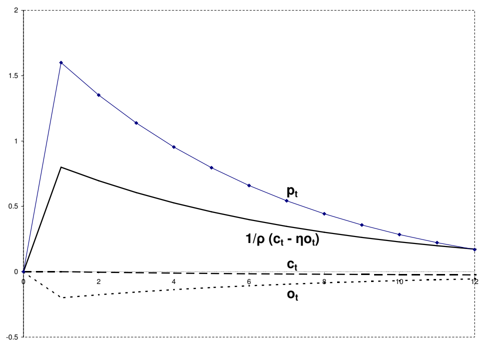

(a) Negative shock to e_t^o

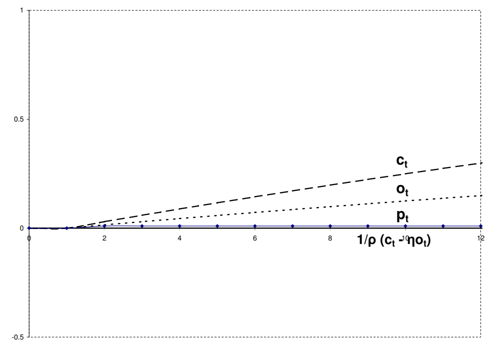

(b) Positive shock to e_t^x

Period 2: $\pi_o \approx 0$ and $\Phi_x = 0$

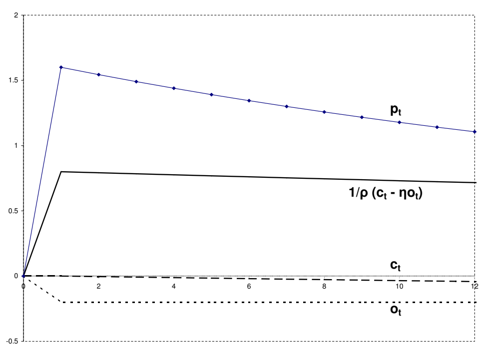

(c) Negative shock to e_t^o

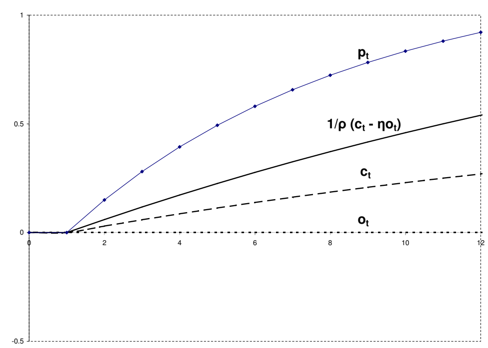

(d) Positive shock to e_t^x

Impulse response function of logs of aggregate consumption (c_t), oil consumption (o_t), and the oil price ($p_t = \frac{1}{\rho}(c_t - \eta o_t) + \xi_t$) to innovations to oil consumption and the expected growth of aggregate consumption.

Table 10: Model Parameters

Parameter	**Period 1**	**Period 2**
Utility		
ψ	1.5	1.5
γ	10	10
δ	0.998	0.998
ρ	0.25	0.25
η	1.75	1.75
τ	0.03	0.03
α	0.5	0.5
Consumption		
σ_c	0.0007	0.0007
σ_o	0.0065	0.0065
σ_x	0.00	0.00
$\sigma_{c,w}$	0.0000001	0.0000001
$\sigma_{o,w}$	0.000009	0.0000092
π^c	-0.002	-0.002
ρ_x	0.995	0.995
π^o	$\underline{0.030}$	$\underline{0.007}$
Φ_x	$\underline{0.57}$	$\underline{0.00}$
Oil Price		
ρ_ξ	0.95	0.95
ϵ	0.56	0.56

Parameters for model one and model two.

Table 11: Data and Model Sample Moments: Aggregate Consumption

	Data				**Model 1**			**Model 2**		
	1952 - 2010		**1987 - 2010**							
	Estimate	Std Error	Estimate	Std Error	**Mean**	**5%**	**95%**	**Mean**	**5%**	**95%**
Cobb-Douglas Aggregate										
$E[\Delta c_t]$	1.01	*(0.03)*	1.01	*(0.04)*	0.96	0.45	1.47	0.94	0.42	1.42
$\sigma(\Delta c_t)$	0.41	*(0.02)*	0.37	*(0.03)*	0.43	0.38	0.71	0.43	0.38	0.67
$AC(1)\Delta c_t$	0.68	*(0.05)*	0.85	*(0.06)*	0.86	0.73	0.95	0.88	0.73	0.94

Calibrated model and data moments for consumption. c_t is a Cobb-Douglass aggregate of durable and nondurable consumption. Durable consumption growth is calculated as the growth in stock of durable goods from the NIPA consumption survey. Nondurable consumption is the sum of nondurables and services excluding energy goods.

Figure 7 provides graphs of expected futures curves for the two different calibrations of the model each containing three panels. These graphs show the success the model has in generating the change

Table 12: Data and Model Sample Moments: Oil Prices and Oil Consumption

	Data 1987 - 1999 Estimate	Std Error	Model 1 Mean	5%	95%	Data 2000 - 2010 Estimate	Std Error	Model 2 Mean	5%	95%
				A. Oil Consumption						
$E[\Delta o_t]$	0.36	*(0.19)*	0.54	0.14	0.89	0.03	*(0.20)*	0.56	0.39	0.74
$\sigma(\Delta o_t)$	1.42	*(0.14)*	1.16	1.03	1.29	1.20	*(0.14)*	1.16	1.03	1.28
$AC(1)\Delta o_t$	-0.31	*(0.13)*	-0.04	-0.17	0.09	0.13	*(0.17)*	-0.01	-0.12	0.12
				B. Spot Price of Oil						
$E[\Delta p_t]$	0.40	*(2.30)*	0.62	-0.11	-0.10	1.80	*(2.63)*	0.10	-2.80	3.20
$\sigma(\Delta p_t)$	16.90	*(1.63)*	18.28	16.25	20.24	19.30	*(1.86)*	18.41	16.34	20.43
$AC(1)p_t$	0.73	*(0.09)*	0.68	0.44	0.84	0.91	*(0.07)*	0.85	0.67	0.97
				C. Oil Futures Prices						
	Data 1987 - 1999 Estimate	Std Error	**Model 1** Mean	5%	95%	**Data 2002 - 2010** Estimate	Std Error	**Model 2** Mean	5%	95%
$\sigma(f_1)$	9.48	*(0.41)*	10.82	9.06	12.72	9.79	*(0.56)*	10.73	8.89	12.66
$\sigma(f_{12})$	5.04	*(0.29)*	4.02	3.38	4.71	6.91	*(0.40)*	7.60	6.40	8.93
$\sigma(slope)$	10.04	*(0.97)*	16.59	11.52	22.81	10.71	*(1.03)*	10.95	7.30	15.66

Calibrated model and data moments for oil consumption and prices. Data for oil consumption is from consumption of "Energy Goods" in the NIPA survey. The spot price of oil is the WTI spot price deflated by the CPI excluding energy goods. Future prices are NYMEX futures for Crude Light Sweet oil.

Table 13: Data and Model Sample Moments: Dividends and Returns

	Data 1952 - 2010 Estimate	Std Error	Data 1987 - 2010 Estimate	Std Error	Model 1 Mean	5%	95%	Model 2 Mean	5%	95%
				Dividend Growth (Quarterly)						
$E[\Delta y_t]$	1.33	*(0.03)*	0.10	*(0.04)*	1.23	0.15	2.24	1.29	0.41	2.11
$\sigma(\Delta y_t)$	5.80	*(0.29)*	5.70	*(0.03)*	5.04	3.99	6.11	5.08	4.21	6.05
				Price Dividend Ratio (Quarterly)						
$E[P^m/Y]$	35.80	*(1.03)*	46.20	*(0.06)*	22.92	19.13	26.50	38.23	34.44	42.13
$\sigma(p^m - y)$	11.00	*(0.72)*	8.90	*(0.04)*	11.65	7.13	18.76	9.65	6.25	14.33
				Returns (Annual)						
$E[r_f]$	1.01	*(0.03)*	1.01	*(0.04)*	1.93	-0.53	4.36	3.09	1.10	4.89
$\sigma[r_f]$	0.49	*(0.02)*	0.39	*(0.03)*	1.11	0.57	1.91	0.97	0.52	1.64
$E[r - r_f]$	5.80	*(1.79)*	4.23	*(0.48)*	7.40	3.37	12.29	4.63	0.52	8.95
$\sigma[r - r_f]$	15.32	*(0.05)*	2.13	*(0.06)*	12.10	7.67	17.38	11.53	7.37	16.31

Equity return, price, and dividend data are from the Standard and Poor's Composite Index. The risk free rate is the one-month Treasury Bill.

in the shape of the term structure of futures as well as the term structure of returns, and creating the distinctive hump shaped pattern of observed futures in the second period. I am primarily concerned with matching the relative changes across the term structure. I do not view the aggregate return on on oil spot prices over such short periods as true indicators of risk. I therefore normalize the return on the one month future to be equal to the observed return and then observe the relative pattern of returns at different term structures.

The tables and figure show that for reasonable values of consumption volatility and autocorrelation, the model is able to match the change in behavior of the oil futures curve. Table 14 gives the result of the volatility regressions in the model corresponding to the regressions Panel A of Table 8. The regressions in the model have the same qualitative pattern as those in the data. The upward

Figure 7: The Term Structure of Crude Oil Futures: Model and Data

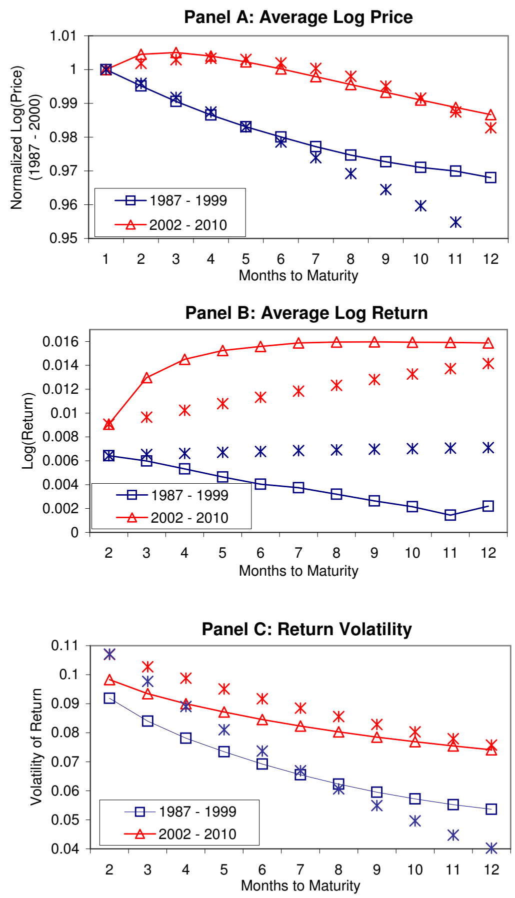

Observed curves and average model generated curves. Lines represent data and stars represent the model generated curves. For both futures and returns, the nearest maturity model moment is normalized to equal the observed moment in the data.

slope in the futures curve from exposure to oil price shocks is larger when oil price volatility is high, and likewise the downward slope is more pronounced when aggregate stock market volatility is high. This gives further support of the general result that two separate sources of risk have opposite effects on the slope of the futures curve. Finally Table 15 shows the increase in price-dividend ratio over the second half of the sample. Though there are obviously other factors that could be at play in generating this effect, such as the regulations relating to stock repurchases, it is generally consistent with the result in the model, which is generated by the muting effect of oil prices on long run consumption growth.

Table 14: Volatility Regressions in Model

Period	Dep. Var.	$\sigma_{spot,t}$	$\sigma_{market,t}$	Δp_t	Δp_{t-1}	Δp_{t-2}	Δp_{t-3}	Constant
1987 - 2000	$f_t^{12} - f_t^1$	0.44	0.01	-0.31	-0.29	-0.28	-0.27	0.00
2001 - 2010	$f_t^{12} - f_t^1$	1.53	-0.83	-0.11	-0.10	-0.10	-0.09	-0.01

$\sigma_{market,t}$ and $\sigma_{spot,t}$ are conditional volatilities of the stock market return and spot price for time $t+1$ based on observed values based on observed values of $\sigma_{c,t}$ and $\sigma_{o,t}$.

Table 15: Data and Model Sample Moments: P/D Ratio for Split Sample

	Data 1987 - 1999		Model 1			Data 2002-2010		Model 2		
	Estimate	Std Error	Mean	5%	95%	Estimate	Std Error	Mean	5%	95%
					Price Dividend Ratio (Quarterly)					
$E[P^m/Y]$	42.27	*(2.55)*	32.98	28.59	37.34	56.86	1.86	53.77	51.09	58.96
$\sigma(p^m - y)$	18.97	*(1.24)*	11.84	7.36	18.12	11.50	0.91	8.56	5.78	12.31

Equity return and dividend data are from the Standard and Poor's Composite Index.

6 Conclusion

This paper highlights the importance of oil prices as a factor in the dynamics of expected consumption growth, yielding a rich set of implications for asset prices. The risks associated with predicted consumption growth provide an explanation for observed changes in the term structure of oil futures prices. The changes in consumption dynamics which generate these changes also have much broader implications for risk in the overall economy. The decreased response of levels of oil consumption to high oil prices lead to a highly persistent oil price, which leads to increase in risk from oil price shocks, but has a counterbalancing effect on shocks to expected consumption growth. In models of Long-Run Risk, shocks to growth are the primary force behind generating levels of risk sufficient to explain observed returns in asset prices. In the model presented here, the effect of oil prices is to reduce this risk, and with it reduce the equity premium.

These changes also have many implications outside of those considered here. For example, if companies' stock returns have differential exposure to shocks to oil prices or shocks to expected growth, the changes in prices of risks associated with these shocks will have implications in the cross-section of expected equity returns. With the ongoing concerns about oil supply in the coming decades, understanding how changes in the state of the oil market affect asset markets as a whole is crucially important. To my knowledge this paper is the first to explore these issues in detail, and will hopefully encourage further research in this area.

7 Appendix

Appendix A - Cointegration of Oil and Aggregate Consumption

This section provides empirical evidence for the cointegrating relation between oil consumption and aggregate consumption, which I will refer to as the Intertemporal relation. This relation also implies the stationarity of the price of oil. If p_t is an $I(0)$ variable, then the Intratremporal relation implies that $\frac{1}{\rho}(c_t - \eta o_t)$ will likewise be $I(0)$. This is equivalent to saying that there exists a cointegrating relation between oil consumption and aggregate consumption. Therefore, alternative way to test for the existence of this unit root arises from my novel formulation of log oil prices as linear combination of oil consumption and total consumption.

The system of consumption dynamics implied by the equations (5) along with the intratemporal cointegrating relation of equation (1) imply that the real oil price itself be a stationary variable, which is at odds with findings by Hamilton [2008] and Maslyuk and Smyth [2008]. I reconcile this fact by noting that for the period from 1987 to 2000, augmented Dickey-Fuller and Phillips-Perron tests can reject the hypothesis of the unit root, though not for the whole sample (Table 5). In addition, given the relation between oil prices and consumption, I can also approach this issue by looking for the existence of a cointegrating vector between oil consumption and consumption, or likewise by checking for the existence of two cointegrating vectors amongst the system of oil prices and both types of consumption. Johansen (1991) tests provide a method of testing a null hypothesis H_0: (m cointegrating vectors) versus an alternative hypothesis H_1: ($m+1$ cointegrating vectors). I perform tests for both the existence of a single cointegrating variable between oil consumption and aggregate consumption, as well as two cointegrating variables between aggregate consumption, oil consumption, and the real spot price. Results are reported in Table 16.

Table 16: Johansen Tests of Cointegration for Consumption, Oil Consumption, and Oil Prices

Variables	**Max Rank**	**1987 - 1999**	**2000 - 2010**	**1987 - 2010**	**5% Critical Values**
c_t, o_t	0	16.29*	16.00*	16.31*	15.41
	1	0.94	3.73	3.14	3.76
c_t, o_t, p_t	0	42.03*	44.74*	49.89*	29.68
	1	15.58*	16.21*	21.54*	15.41
	2	2.59	2.88	2.26	3.76

Johansen tests of cointegration are conducted with two lags. o_t is the measure of personal energy good consumption from the NIPA survey, c_t an aggregation of nondurables and services and the stock of durable goods. Consumption is real consumption per capita.

The tests generally support the existence of a cointegrating vector between oil consumption and aggregate consumption, and the existence of two cointegrating variables in the trivariate system. Given the results of these estimates, I estimate a Vector Error Correction Model (VECM) of oil consumption and aggregate consumption in each subperiod of the following form, with results reported in Table 17.

$$\begin{bmatrix} \Delta c_t \\ \Delta o_t \end{bmatrix} = \begin{bmatrix} \mu_c \\ \mu_o \end{bmatrix} + \begin{bmatrix} \pi^c \\ \pi^o \end{bmatrix} \begin{bmatrix} c_{t-1} - \eta o_{t-1} \end{bmatrix} + \Gamma_1 \begin{bmatrix} \Delta c_{t-1} \\ \Delta o_{t-1} \end{bmatrix} \tag{46}$$

This estimation supports the change in dynamics used in the model, namely the changes to π_o and Φ_x. In the second period, expected consumption growth (here represented by lagged innovation to growth) has no impact on future oil consumption growth. Likewise, the cointegrating vector generates negative future aggregate consumption growth but little expected growth in oil consumption in the second period. In the first period, both the cointegrating vector and lagged consumption growth predict positive future oil consumption growth, consistent with the model parameters. The one worrying result from these estimations is that the parameter in the cointegrating vector governing the weight on the value of oil consumption is estimated to be significantly different than the analogous parameter, η, representing the weight of o_t relative to c_t in the oil price approximation. The standard error on this variable is wide however, and if this variable is constrained to equal η,the result is very similar to the simple regressions reported in the text.

As a further robustness check I also perform a VAR in first differences (not reported) which again yields the same qualitative results. I therefor choose the cointegrating framework in the model since it provides the maximum amount of parsimony while still capturing the pertinent effects.

Appendix B - Model and Solutions

In this section I derive approximate analytical solutions for the long run risk model with oil consumption. Lowercase variables represent logs.

Table 17: VECM for Aggregate Consumption and Oil Consumption

(a) Panel A: 1987 - 1999

Equation	**Parameter**	**Estimate**	**Std. Error**	**t−Value**	**Pr $> \|t\|$**	**Variable**
Δc_{t+1}	μ^c	0.001	0.00	1.13	0.26	
	π^c	0.014	0.02	0.84	0.03	$c_t - \eta o_t$
	$\Gamma_{1,1}$	$0.824^\star$	0.08	9.69	0.00	Δc_t
	$\Gamma_{1,2}$	0.006	0.02	0.29	0.77	Δo_t
Δo_{t+1}	μ^o	0.000	0.01	0.02	0.99	
	π^o	$0.273^\star$	0.08	3.24	0.03	$c_t - \eta o_t$
	$\Gamma_{2,1}$	$1.110^\star$	0.51	2.16	0.03	Δc_t
	$\Gamma_{2,2}$	-0.070	0.16	-0.43	0.67	Δo_t

	Cointegrating Vector			
Equation	**Estimate**	**Std. Error**	**t−Value**	**Pr $> \|t\|$**
c_t	1			
o_t	-1.75	0.16	-18.33	0
constant	3.50			

(b) Panel B: 2000 - 2010

Equation	**Parameter**	**Estimate**	**Std. Error**	**t−Value**	**Pr $> \|t\|$**	**Variable**
Δc_{t+1}	μ^c	.001	0.00	.85	.398	
	π^c	$-0.012^\star$	0.00	-3.10	0.00	$c_t - \eta o_t$
	$\Gamma_{1,1}$	$0.37^\star$	0.18	1.98	0.05	Δc_t
	$\Gamma_{1,2}$	0.001	0.00	1.34	0.18	Δo_t
Δo_{t+1}	μ^o	0.033	0.03	1.13	0.26	
	π^o	0.033	0.03	0.68	0.50	$c_t - \eta o_t$
	$\Gamma_{2,1}$	0.88	1.31	0.42	0.68	Δc_t
	$\Gamma_{2,2}$	0.000	0.00	0.06	0.95	Δo_t

	Cointegrating Vector			
Equation	**Estimate**	**Std. Error**	**t−Value**	**Pr $> \|t\|$**
c_t	1			
o_t	-5.15	1.12	-3.88	0
constant	-1.37			

Vector error correction methods estimated using Johansen's MLE estimation. c_t is the aggregation of nondurable and durable consumption, o_t is the NIPA measure of energy consumption. Consumption is real consumption per capita.

Intratemporal Utility

Define

$$C_t \equiv N_t^{1-\alpha} D_t^{\alpha} \tag{47}$$

Where N_t is nondurable consumption expenditure excluding energy goods, and D_t is the stock of durable consumption goods. Define intratemporal utility as

$$V_t(C_t, O_t) = \left[(1-a)C_t^{1-\frac{1}{\rho}} + aO_t^{1-\frac{\eta}{\rho}}\right]^{\frac{\rho}{\rho-1}} \tag{48}$$

Cobb-Douglas Approximation

In order to allow for analytical solutions to the model, I approximate the generalized CES utility function with a Cobb-Douglas utility function.

Let $H_t = \frac{P_t O_t}{C_t}$ be the ratio of expenditure on oil to expenditure on the aggregate consumption good. $\bar{H}$ is the sample average of H_t.

Given the intratemporal first order condition. The generalized CES function can be rewritten as[8]:

$$V_t = C_t \left((1+a)(1+\frac{1-\frac{1}{\rho}}{1-\frac{\eta}{\rho}}H_t)\right)^{\frac{1}{1-\frac{1}{\rho}}} \tag{49}$$

Taking a first order Taylor approximation of the log of intratemporal utility around the sample average ratio of expenditure gives

$$v_t = c_t + \frac{1}{1-\frac{1}{\rho}}\left(log(1-a) + (1+\frac{1-\frac{1}{\rho}}{1-\frac{\eta}{\rho}}\bar{H}_t) + \frac{\frac{1-\frac{1}{\rho}}{1-\frac{\eta}{\rho}}\bar{H}_t}{1+\frac{1-\frac{1}{\rho}}{1-\frac{\eta}{\rho}}\bar{H}_t}(h_t - \bar{h}_t)\right) \tag{50}$$

Since empirically the average value of of H_t is roughly .025, the higher order terms are extremely small. Therefore I focus on the ability of the approximation in explaining the first order terms.

The Cobb - Douglas approximation is

$$\tilde{V}_t = C_t^{1-\tau}O_t^{\tau} \tag{51}$$

where $\tau = \frac{\bar{H}}{1+\bar{H}}$

The approximation error to the first order terms is

[8]This holds for both Equations 1 and 44

$$v_t - \tilde{v}_t = \text{constant} + (h_t - \bar{h})\frac{\bar{H}^2}{1+\bar{H}}\frac{(1-\eta)(\rho-1)}{(\rho-\eta)^2} \tag{52}$$

Again, since $\bar{H}$ is observed to be very small, this approximation error is negligible. The marginal utilities of consumption and oil consumption under the generalized CES specification are

$$V_{c,t} = \frac{V_t}{C_t}\frac{1}{1+\left(\frac{1-\frac{1}{\rho}}{1-\frac{\eta}{\rho}}\right)H_t} \tag{53}$$

$$V_{o,t} = \frac{V_t}{O_t}\frac{H_t}{1+\left(\frac{1-\frac{1}{\rho}}{1-\frac{\eta}{\rho}}\right)H_t} \tag{54}$$

The marginal utilities under the approximation are

$$\tilde{V}_{c,t} = \frac{\tilde{V}_t}{C_t}\frac{1}{1+\bar{H}} \tag{55}$$

$$\tilde{V}_{o,t} = \frac{\tilde{V}_t}{O_t}\frac{\bar{H}}{1+\bar{H}} \tag{56}$$

Due to the small values of H_t observed in the sample, and the low variance of H_t, this approximation performs well in terms of relative changes in marginal utility.

Intertemporal Utility

I consider an agent with Epstein-Zin Preferences and intratemporal utility $\tilde{V}_t$.

Following Yogo [2005] and Yang [2010] the log of the pricing kernel is

$$m_{t+1} = \Theta\log\delta + -\frac{\Theta}{\psi}\Delta c_{t+1} + \tau\Theta(1-\frac{1}{\psi})(\Delta o_{t+1} - \Delta c_{t+1}) + (\Theta-1)r_{W,t+1}$$

where $r_{W,t+1}$ is the return on total wealth.

Solving for the Return on Wealth

I consider the following dynamics for aggregate consumption and oil consumption.

$$
\begin{aligned}
\Delta c_{t+1} =& \mu^c + \pi^c \left[c_t - \eta o_t - \bar{p}_t\right] + x_t + \sigma_{c,t} e^c_{t+1} \\
\Delta o_{t+1} =& \mu^o + \pi^o \left[c_t - \eta o_t - \bar{p}_t\right] + \Phi_x x_t + \sigma_{o,t} e^o_{t+1} \\
x_{t+1} =& \rho_x x_t + \varphi_x \sigma_{c,t} e^x_{t+1} \\
\bar{p}_{t+1} =& \bar{p} - \mu^p \\
\sigma_{c,t+1} =& \nu_c [\sigma^2_{c,t} - \bar{\sigma_c}^2] + \bar{\sigma_c}^2 + \sigma_{c,w} w^c_{t+1} \\
\sigma_{o,t+1} =& \nu_o [\sigma^2_{o,t} - \bar{\sigma_o}^2] + \bar{\sigma_o}^2 + \sigma_{o,w} w^o_{t+1} \\
p_t =& \frac{1}{\rho}(c_t - \eta o_t) + \xi_t \\
\xi_{t+1} =& \frac{1-\epsilon}{\epsilon}(\Delta c_{t+1} + \eta \Delta o_{t+1}) + \rho_\xi \xi_t
\end{aligned} \tag{57}
$$

For exposition I will assume that the shock terms are uncorrelated, though it is easy to allow for correlation between the shock terms. In practice, when estimating the model, I will set correlations to match the observed correlations in the data.

I follow Bansal and Yaron [2004] and utilize the Campbell approximation for the return on total wealth.

$$
r_{g,t+1} = \kappa_0 + \kappa_1 z_{t+1} - z_t + \Delta v_{t+1}
$$

Given these approximations, z_t is affine in four state variables, the predictable term x_t and the oil price error correction term $c_t - \eta o_t - \bar{p}_t$, and the two stochastic volatility components, $\sigma^2_{c,t}$ and $\sigma^2_{o,t}$.

$$
z_t = A_0 + A_1 x_t + A_2(c_t - \eta o_t - \bar{p}_t) + A_3 \sigma^2_{c,t} + A_4 \sigma^2_{o,t} \tag{58}
$$

To solve for the values of the A coefficients, I utilize the pricing equation.

$$1 = E_t[exp(m_{t+1} + r_{g,t+1})] \tag{59}$$

$$0 = E_t[m_{t+1} + r_{g,t+1}] + \frac{1}{2}\text{var}_t[m_{t+1} + r_{g,t+1}] \tag{60}$$

I collect terms and obtain the following values for the coefficients on the state variables.

Define

$$M^c = (1-\tau) - \frac{1}{\psi} - \tau(1 - \frac{1}{\psi}) \tag{61}$$

$$M^o = \tau(1 - \frac{1}{\psi}) + \tau \tag{62}$$

Then the solutions for the coefficients are

$$A_0 = \frac{log(\delta) + M^c\mu^c + M^o\mu^o - A_2\kappa_1\mu^p + \kappa_1 A_3(1-\nu_c)(\bar{\sigma_c})^2 + \kappa_1 A_4(1-\nu_o)(\bar{\sigma_o})^2 + \frac{1}{2}\Theta[A_3\sigma^2_{c,w} + A_4\sigma^2_{o,w}]}{1-\kappa_1} \tag{63}$$

$$A_1 = \frac{M^c + M^o\Phi_x + A_2\kappa_1(1-\eta\Phi_x) - \frac{1}{2}\varphi[\sigma_c^2(M^c + \phi^c M^o + A_2(1-\eta\Phi_c)\kappa_1)]}{1-\kappa_1(\rho_x - \sigma_x^2\varphi\Theta\frac{1}{2})} \tag{64}$$

$$A_2 = \frac{\pi^c M^c + \pi^o M^o}{1-\kappa_1(1+\pi^c-\eta\pi^o)} \tag{65}$$

$$A_3 = \frac{\frac{1}{2}\Theta\left[(M^c + A_2\kappa_1)^2 + A_1^2\varphi_x^2\right]}{1-\nu_c\kappa_1} \tag{66}$$

$$A_4 = \frac{\frac{1}{2}\Theta(M^o - \eta A_2\kappa_1)^2}{1-\nu_o\kappa_1} \tag{67}$$

$$\tag{68}$$

Innovations to the pricing kernel are

$$m_{t+1} - E_t[m_{t+1}] = \left(-\theta(\frac{1}{\psi} - \tau(1 - \frac{1}{\psi})\right) + (\theta - 1)A_2\kappa_1\sigma_{c,t}e^c_{t+1} \tag{69}$$

$$+ \left(-\gamma\tau + \theta\tau - \eta(\theta - 1)\right) A_2\kappa_1\sigma_{o,t}e^o_{t+1}$$

$$+ (\theta - 1)\kappa_1 A_1\varphi_x\sigma_{c,t}e^x_{t+1}$$

$$+ (\theta - 1)\kappa_1 A_3\sigma_{w,c}w^c_{t+1}$$

$$+ (\theta - 1)\kappa_1 A_4\sigma_{w,o}w^o_{t+1}$$

$$\tag{70}$$

Equivalently

$$m_{t+1} - E_t[m_{t+1}] = -\lambda_c\sigma_{c,t}e^c_{t+1} - \lambda_o\sigma_{o,t}e^o_{t+1} - \lambda_x\varphi_x\sigma_{c,t}e^x_{t+1} - \lambda_{\sigma,c}\sigma_{w,c}w^c_{t+1} - \lambda_{\sigma,c}\sigma_{w,o}w^o_{t+1} \tag{71}$$

Oil Prices

Oil futures prices are linear in the state variables.

$$f^j_t = \bar{p}_t + B^j_0 + B^j_x x_t + B^j_p(c_t - \eta o_t - \bar{p}) + B^j_{\sigma,c}\sigma_{c,t} + B^j_{\sigma,o}\sigma_{o,t} + B^j_\xi\xi_t \tag{72}$$

The coefficients can be calculated by the following recursions.

$$B_0^j = B_0^{j-1} + \left(B_p^{j-1} + \frac{1-\epsilon}{\epsilon} B_\xi^{j-1}\right)(\mu_c - \eta\mu^o) + \mu_p \tag{73}$$

$$+ B_{\sigma,o}^{j-1}(1-\nu_o)\bar{\sigma_o} + B_{\sigma,c}^{j-1}(1-\nu_c)\bar{\sigma_c} \tag{74}$$

$$+ \frac{1}{2} B_{w,o}^{j-1}\sigma_{w,o}^2 + \lambda_{w,o} B_{\sigma,c}^{j-1}\sigma_{w,c}^2$$

$$+ \frac{1}{2} B_{w,c}^{j-1}\sigma_{w,c}^2 + \lambda_{w,c} B_{\sigma,o}^{j-1}\sigma_{w,o}^2$$

$$B_x^j = B_x^{j-1}\rho_x + (B_p^{j-1} + \frac{1-\epsilon}{\epsilon} B_\xi^{j-1})(1 - \eta\phi_x)$$

$$B_p^j = (1 + \pi^c - \eta\pi^o) B_p^{j-1}$$

$$B_{\sigma,c}^j = \nu_c B_{\sigma,c}^{j-1} + \frac{1}{2}\left[(B_p^{j-1} + \frac{1-\epsilon}{\epsilon} B_\xi^{j-1})^2 + (B_x^{j-1}\varphi_x)^2\right] \tag{75}$$

$$- \left(B_p^{j-1} + \frac{1-\epsilon}{\epsilon} B_\xi^{j-1}\right)\lambda_c - B_x^{j-1}\varphi_x\lambda_x$$

$$B_{\sigma,o}^j = \nu_o B_{\sigma,o}^{j-1} + \frac{1}{2}(B_p^{j-1} + \frac{1-\epsilon}{\epsilon} B_\xi^{j-1})^2\eta^2 + \left(B_p^{j-1} + \frac{1-\epsilon}{\epsilon} B_\xi^{j-1}\right)\lambda_o\eta \tag{76}$$

$$B_\xi^j = \rho_\xi B_\xi^{j-1} \tag{77}$$

Since $f_t^0 = p_t$. The initial values for the recursion are given by

$$B_0^0 = 0 \tag{78}$$

$$B_x^0 = 0$$

$$B_p^0 = \frac{1}{\rho}$$

$$B_{\sigma,c}^0 = 0$$

$$B_{\sigma,o}^0 = 0$$

$$B_\xi^0 = 1$$

Equity Returns

The innovation to the market dividend y_t is represented by.

$$\Delta y_{t+1} = \mu^y + \chi \left(x_t + \pi^c(c_t - \eta o_t - \bar{p}_t)\right) + \varphi^y e^y_{t+1} \tag{79}$$

The return on the market portfolio, r_{t+1}, solves

$$E_t[exp(m_{t+1} + r_{t+1})] = 1 \tag{80}$$

Again exploiting the Campbell approximation, $r_{t+1} = \kappa^y_0 + \kappa^y_1 z^y_{t+1} - z^y_t + \Delta y_{t+1}$ and assume a linear form

$$z_t = A^y_0 + A^y_1 x_t + A^y_2(c_t - \eta o_t - \bar{p}_t) + A^y_3 \sigma_{c,t} + A^y_4 \sigma_{o,t} \tag{81}$$

The coefficients for are solved for in the same manner as the consumption coefficient, by expanding the the expect pricing equation and collecting terms in each state variable. The amount of terms make the closed form solutions extremely complicated, so they are not reported here.

Expected return can then be calculated as in Bansal-Yaron 2004.

References

Jushan Bai and Pierre Perron. Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):pp. 47–78, 1998.

Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

Ravi Bansal, Dana Kiku, and Amir Yaron. Risks for the long run: Estimation and inference. Working paper. 2007.

Jan Bentzen and Tom Engsted. Short- and long-run elasticities in energy demand: A cointegration approach. *Energy Economics*, 15(1):9–16, 1993.

Hendrik Bessembinder, Jay Coughenour, Paul Seguin, and Margaret Monroe Smoller. Mean reversion in equilibrium asset prices: Evidence from the futures term structure. *Journal of Finance*, 50(1):361–75, March 1995.

Jules H. van Binsbergen, Michael W. Brandt, and Ralph S.J. Koijen. On the timing and pricing of dividends. NBER Working Papers 16455, National Bureau of Economic Research, Inc, 2010.

Douglas T Breeden. Consumption risk in futures markets. *Journal of Finance*, 35(2):503–20, May 1980.

Michael J Brennan and Eduardo S Schwartz. Evaluating natural resource investments. *Journal of Business*, 58(2):135–57, April 1985.

Fabio Busetti and A. M. Robert Taylor. Tests of stationarity against a change in persistence. *Journal of Econometrics*, 123(1):33 – 66, 2004. ISSN 0304-4076.

Murray Carlson, Zeigham Khokher, and Sheridan Titman. Equilibrium exhaustible resource price dynamics. *Journal of Finance*, 62(4):1663–1703, 08 2007.

Jaime Casassus, Pierre Collin-Dufresne, and Bryan R. Routledge. Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technology. *SSRN eLibrary*, 2005.

Angus Deaton and Guy Laroque. On the behaviour of commodity prices. *The Review of Economic Studies*, 59(1):1–23, 1992. ISSN 00346527.

Angus Deaton and Guy Laroque. Competitive storage and commodity price dynamics. *The Journal of Political Economy*, 104(5):896–923, 1996. ISSN 00223808.

Itamar Drechsler and Amir Yaron. What's Vol Got to Do With It. *SSRN eLibrary*, 2009.

Katherine Dusak. Futures trading and investor returns: An investigation of commodity market risk premiums. *Journal of Political Economy*, 81(6):1387–1406, 1973.

Eyal Dvir and Kenneth S. Rogoff. Three epochs of oil. Working Paper 14927, National Bureau of Economic Research, April 2009.

Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4).

Bjorn Eraker. The volatility premium. Working paper, 2008.

Eugene F Fama and Kenneth R French. Commodity futures prices: Some evidence on forecast power, premiums,and the theory of storage. *Journal of Business*, 60(1):55–73, January 1987.

Eugene F. Fama and James D. MacBeth. Risk, return, and equilibrium: Empirical tests. *The Journal of Political Economy*, 81(3):pp. 607–636, 1973.

Rajna Gibson and Eduardo S. Schwartz. Stochastic convenience yield and the pricing of oil contingent claims. *The Journal of Finance*, 45(3):959–976, 1990. ISSN 00221082.

Gary B. Gorton, Fumio Hayashi, and K. Geert Rouwenhorst. The fundamentals of commodity futures returns. Working Paper 13249, National Bureau of Economic Research, July 2007a.

Gary B. Gorton, Fumio Hayashi, and K. Geert Rouwenhorst. The Fundamentals of Commodity Futures Returns. *SSRN eLibrary*, 2007b.

James D Hamilton. Oil and the macroeconomy. *in S. Durlauf and L. Blume, eds., The New Palgrave Dictionary of Economics, 2nd edition*, 2005.

James D. Hamilton. Understanding crude oil prices. NBER Working Papers 14492, National Bureau of Economic Research, Inc, November 2008.

Lars Peter Hansen, John C. Heaton, and Nan Li. Consumption strikes back? Measuring long-run risk. *Journal of Political Economy*, 116(2):260–302, 2008.

David I. Harvey, Stephen J. Leybourne, and A.M. Robert Taylor. Modified tests for a change in persistence. *Journal of Econometrics*, 134(2):441 – 469, 2006.

Ravi Jagannathan. An investigation of commodity futures prices using the consumption-based intertemporal capital asset pricing model. *Journal of Finance*, 40(1):175–91, March 1985.

Sren Johansen. Estimation and hypothesis testing of cointegration vectors in gaussian vector autoregressive models. *Econometrica*, 59(6):pp. 1551–1580, 1991. ISSN 00129682.

Nicholas Kaldor. Speculation and economic stability. *The Review of Economic Studies*, 7(1):1–27, 1939. ISSN 00346527.

John M. Keynes. *A Treatise on Money*, volume 2. Macmillan, London, 1930.

Leonid Kogan, Dmitry Livdan, and Amir Yaron. Oil futures prices in a production economy with investment constraints. *The Journal of Finance*, 64(3):pp. 1345–1375, 2009. ISSN 00221082.

Svetlana Maslyuk and Russell Smyth. Unit root properties of crude oil spot and futures prices. *Energy Policy*, 36(7):2591–2600, July 2008.

Michal Pakos. Asset pricing with durable goods and non-homothetic preferences. Working paper. 2004.

R. Ramanathan. Short- and long-run elasticities of gasoline demand in india: An empirical analysis using cointegration techniques. *Energy Economics*, 21(4):321 – 330, 1999. ISSN 0140-9883. doi: DOI: 10.1016/S0140-9883(99)00011-0.

Morten Ravn, Stephanie Schmitt-Grohe, and Martin Uribe. Relative deep habits. Technical note, 2005.

Bryan R. Routledge, Duane J. Seppi, and Chester S. Spatt. Equilibrium forward curves for commodities. *The Journal of Finance*, 55(3):1297–1338, 2000. ISSN 00221082.

Eduardo S. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997. ISSN 00221082.

Kenneth J. Singleton. The 2008 boom/bust in oil prices. Working paper, 2008.

Lester G. Telser. Futures trading and the storage of cotton and wheat. *The Journal of Political Economy*, 66(3):233–255, 1958. ISSN 00223808.

J.C. Williams and B.D. Wright. *Storage and Commodity Markets.* Cambridge University Press, Cambridge, UK, 1991.

Holbrook Working. The theory of price of storage. *The American Economic Review*, 39(6):1254–1262, 1949. ISSN 00028282.

Wei Yang. Asset Pricing with Left-Skewed Long-Run Risk in Durable Consumption. *SSRN eLibrary*, 2010.

Motohiro Yogo. A Consumption-Based Explanation of Expected Stock Returns. *Journal of Finance, Vol. 61, No. 2, 2006*, 2005.