

NBER WORKING PAPER SERIES

EQUILIBRIUM COMMODITY PRICES WITH
IRREVERSIBLE INVESTMENT AND NON-LINEAR TECHNOLOGY

Jaime Casassus
Pierre Collin-Dufresne
Bryan R. Routledge

Working Paper 11864
<http://www.nber.org/papers/w11864>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2005

We thank Luis Alvarez, Duane Seppi, Chester Spatt, Stan Zin, seminar participants at UC Berkeley, Carnegie Mellon University, the University of Madison-Wisconsin, the Q-group meeting in Key West 2005, the University of Utah, University of Southern California, and the 2003 LACEA Meeting. Correspondence: Jaime Casassus, Escuela de Ingeniería de la Pontificia Universidad Católica de Chile, Av. Vicuña Mackenna 4860, Santiago, Chile. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Jaime Casassus, Pierre Collin-Dufresne and Bryan R. Routledge. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technology
Jaime Casassus, Pierre Collin-Dufresne and Bryan R. Routledge
NBER Working Paper No. 11864
December 2005
JEL No. C0, G12, G13, D51, D81, E2

ABSTRACT

We model equilibrium spot and futures oil prices in a general equilibrium production economy. In our model production of the consumption good requires two inputs: the consumption good and a commodity, e.g., Oil. Oil is produced by wells whose flow rate is costly to adjust. Investment in new Oil wells is costly and irreversible. As a result in equilibrium, investment in Oil wells is infrequent and lumpy. Even though the state of the economy is fully described by a one-factor Markov process, the spot oil price is not Markov (in itself). Rather it is best described as a regime-switching process, the regime being an investment 'proximity' indicator. The resulting equilibrium oil price exhibits mean-reversion and heteroscedasticity. Further, the risk premium for exposure to commodity risk is time-varying, positive in the far-from-investment regime but negative in the near-investment regime. Further, our model captures many of the stylized facts of oil futures prices, such as backwardation and the 'Samuelson effect.' The futures curve exhibits backwardation as a result of a convenience yield, which arises endogenously. We estimate our model using the Simulated Method of Moments with economic aggregate data and crude oil futures prices. The model successfully captures the first two moments of the futures curves, the average non-durable consumption-output ratio, the average oil consumption-output and the average real interest rate. The estimation results suggest the presence of convex adjustment costs for the investment in new oil wells. We also propose and test a linear approximation of the equilibrium regime-shifting dynamics implied by our model, and test its empirical implication for time-varying risk-premia.

Jaime Casassus
Pontificia Universidad Catolica de Chile

Bryan R. Routledge
Carnegie Mellon University

Pierre Collin-Dufresne
Haas School of Business F628
University of California -Berkeley
545 Student Services Bldg #1900
Berkeley, CA 94720-1900
and NBER
dufresne@haas.berkeley.edu

1 Introduction

Empirical evidence suggests that commodity prices behave differently than standard financial asset prices. The evidence also suggests that there are marked differences across types of commodities. This paper presents an equilibrium model of commodity spot and futures prices for a commodity whose primary use is as an input to production, such as oil or copper. The model captures many stylized facts of the data, which we review below.

Empirical studies of time series of commodity prices have found evidence of mean-reversion and heteroscedasticity. Further, combining time series and cross-sectional data on futures prices provides evidence of time-variation in risk-premia as well as existence of a ‘convenience yield’ (Fama and French (1987), Bessembinder et al. (1995), Casassus and Collin-Dufresne (CC 2005)). Interestingly, the empirical evidence also suggests that there are marked differences across different types of commodities (e.g., Fama and French (1987)). CC (2005) use panel data of futures prices to disentangle the importance of convenience yield versus time-variation in risk-premia for various commodities. Their results suggest that ‘convenience yields’ are much larger and more volatile for commodities that serve as an input to production, such as copper and oil, as opposed to commodities that may also serve as a store of value, such as gold and silver. A casual look at a sample of futures curve for various commodities (reproduced in figure 1 below) clearly shows the differences in futures price behavior. Gold and silver markets exhibit mostly upward sloping futures curve with little variation in slope, whereas copper and especially oil futures curve exhibit more volatility. In particular, oil future curves are mostly downward-sloping (i.e., in backwardation), which, given the non-negligible storage costs¹ indicates the presence of a sizable ‘convenience yield.’ Further, casual empiricism suggests that the oil futures curves are not Markov in the spot oil price (as highlighted in figure 3, which shows that for the same oil spot price one can observe increasing or decreasing futures curves). Lastly, the volatility of oil futures prices tends to decrease with maturity (the ‘Samuelson effect’) much more dramatically than that of gold futures prices.

The commodity literature can be mainly divided into two approaches. The equilibrium (or structural) models of commodity prices focus on the implications

¹The annual storage cost are estimated to be around 20% of the spot price by Ross (1997).

of possible stockouts, which affects the no-arbitrage valuation because of the impossibility of carrying negative inventories (Gustafson (1958), Newbery and Stiglitz (1981), Wright and Williams (1982), Scheinkman and Schechtman (1983), Williams and Wright (1991), Deaton and Laroque (1992), Chambers and Bailey (1996), and Bobenrieth, Bobenrieth and Wright (2002)). These papers predict that in the presence of stock-outs, prices may rise above expected future spot prices net of cost of carry. The implications for futures prices have been studied in Routledge, Seppi and Spatt (2002). One of the drawbacks of this literature is that the models are highly stylized and thus cannot be used to make quantitative predictions about the dynamics of spot and futures prices. For example, these papers assume risk-neutrality which forces futures prices to equal expected future spot prices and thus rule out the existence of a risk premium. Further, these models in general predict that strong backwardation can occur only concurrently with stock-outs. Both seem contradicted by the data. Fama and French (1988), Casassus and Collin-Dufresne (2005) document the presence of substantial time variation in risk-premia for various commodities. Litzenberger and Rabinowitz (1995) find that strong backwardation occurs 77% of the time² in oil futures markets, whereas stock-outs are the exception rather than the rule. Litzenberger and Rabinowitz offer an alternative explanation for backwardation based on option pricing theory. They view oil in the ground as a call option written on the spot oil price with exercise price equal to the extraction cost. In equilibrium, a convenience yield (and backwardation) must exist for producers to have an incentive to extract (i.e, exercise their option). Backwardation is the price to pay for the flexibility producers have to refrain from producing at any time, and keep oil in the ground.

In contrast, reduced-form models exogenously specify the dynamics of the commodity spot price process, the convenience yield and interest rates to price futures contracts as derivatives following standard contingent claim pricing techniques (e.g., Gibson and Schwartz (1990), Brennan (1991), Ross (1997), Schwartz (1997), Schwartz and Smith (2000) and Casassus and Collin-Dufresne (2002)). The convenience yield is defined as an implicit dividend that accrues to the holder of the commodity (but not to the holder of the futures contract). This definition builds loosely on the insights of the original ‘theory of storage’ (Kaldor (1939), Working (1948, 1949), Telser (1958), Brennan (1958)) which argues that there are benefits

²And in fact, weak backwardation, when futures prices are less than the spot plus cost of carry, occurs 94% of the times.

for producers associated with holding inventories due to the flexibility in meeting unexpected demand and supply shocks without having to modify the production schedule. The reduced-form approach has gained widespread acceptance because of its analytical tractability (the models may be used to value sophisticated derivatives) as well as its flexibility in coping with the statistical properties of commodity processes (mean-reversion, heteroscedasticity, jumps). However, reduced-form models are by nature statistical and make no predictions about what are the appropriate specifications of the joint dynamics of spot, convenience yield and interest rates. The choices are mostly dictated by analytical convenience and data.

In this paper we propose a general equilibrium model of spot and futures prices of a commodity whose main use is as an input to production. Henceforth we assume that the commodity modeled is oil.

Three features distinguish our model from the equilibrium ‘stock-out’ models mentioned above. First, we consider that the primary use of the commodity is as an input to production. Commodity is valued because it is a necessary input to produce the (numeraire) consumption good. We assume a risky two-input constant returns to scale technology. Second, we assume that agents are risk-averse. This allows us to focus on the risk-premium associated with holding the commodity versus futures contracts. Finally, we assume that building oil wells and extracting oil out of the ground is a costly process. We assume these costs are irreversible in the sense that once built an oil well can hardly be used for anything else but producing oil. This last feature allows us to focus on the ‘precautionary’ benefits to holding enough commodity to avoid disruption in production.

We derive the equilibrium consumption and production of the numeraire good, as well as the demand for the commodity. Investment in oil wells is infrequent and ‘lumpy’ as a result of fixed adjustment costs and irreversibility. As a result there is a demand for a security ‘buffer’ of commodity. Further, the model generates mean-reversion and heteroscedasticity in spot commodity prices, a feature shared by real data. One of the main implications of our model is that even though uncertainty can be described by one single state variable (the ratio of capital to commodity stock), the spot commodity price is not a one-factor Markov process. Instead, the equilibrium commodity price process resembles a jump-diffusion regime switching process, where expected return (drift) and variance (diffusion) switch as the economy moves

from the ‘near-to-investment’ region to the ‘far-from-investment’ region. The equilibrium spot prices may also experience a jump when the switch occurs. The model generates an endogenous convenience yield which has two components, an absolutely continuous component in the no-investment region and a singular component at the investment boundary. This convenience yield reflects the benefit to smoothing the flow of oil used in production. It is decreasing in the outstanding stock of oil and increasing in the marginal productivity of oil in the economy.

When the economy is in the investment region, the fixed costs incurred induce a wealth effect which leads all security prices to jump. Since the investment time is perfectly predictable, all financial asset prices must jump by the same amount to rule out arbitrage. However, we find that in equilibrium, oil prices do not satisfy this no-arbitrage condition. Of course, the apparent ‘arbitrage opportunity’ which arises at investment dates, subsists in equilibrium, because oil is not a traded asset, but instead valued as an input to production.

We implement the Simulated Method of Moments of Duffie and Singleton (1993) to estimate the model. We use quarterly data of crude oil futures prices and aggregate macroeconomic variables of OECD countries from 1990 to 2004. In particular, we find parameters that best fit the futures curve, the volatility term structure of futures returns, the consumption-output ratio, the consumption of oil-output ratio and the real interest rate. We find strong evidence that supports the presence of fixed investment cost, and thus two regimes in prices. We further find that the futures curves can be in contango or in backwardation depending on the state of the economy. As observed in real data the frequency of backwardation dominates that of contango. The two-regimes which characterize the spot price also determine the shape of the futures curve. We find that futures curve reflect a high degree of mean-reversion (i.e., are more convex) when the economy is in the ‘near-to-investment’ region. This is partly due to the increased probability of an investment which announces a drop in the spot price. Finally, our model predicts that risk-premia on commodity prices are time varying, positive in the *far-from-investment* regime and negative in the *near-investment* regime, contributing to the mean-reversion in the spot price. Further, the systematic risk of the commodity price as measured with its beta relative to the market (defined as the present value of the capital stock) return is positive in the *far-from-investment* regime and negative in the *near-investment* regime. This is, at least in principle, consistent with the wildly different estimates

of the magnitude of the risk-premium on commodities obtained in recent empirical studies (e.g., Gorton and Rouwenhorst (2005), Erb and Harvey (2005)).

To test some of the implications of our model for the shape of the term structure of futures and for the risk-premia across regimes we investigate a simple linear approximation of our regime switching spot price model. We use quasi-maximum likelihood technique of Hamilton (1989) to estimate the model with crude oil data from 1990 to 2003. We find strong support for the existence of two regimes with features consistent with those predicted by our model. There is an infrequent state that is characterized by high prices and negative return and a more frequent state that has lower average price and exhibits mean-reversion. To further test the model we estimate the smoothed inference about the state of the economy (Kim (1993)), i.e., we back out the inferred probability of being in one state or the other. We compare the shape of futures curves in both states of the economy and find that, as predicted by the theoretical model, futures curves are mostly convex in the near-to-investment region but concave in the far-from investment region, reflecting the high degree of mean-reversion when investment and a drop in prices is imminent.

We also find some evidence for time variation in the risk-premium on oil price returns that is related to the estimated regime. Indeed, regressing oil price return on the S&P 500 return we find that the beta is significantly negative in the estimated *near-investment* regime and positive (though not statistically significant) in the other regime. This significant time variation in beta is not driven out by conditioning on the slope of the futures term structure, which suggests that, as in the model, slope of the futures curve is not a perfect substitute for the investment regime.

This provides some validation for our equilibrium model and also suggests that a regime switching model may be a useful alternative to the standard reduced-form models studied in the literature.

In a sense our model formalizes many of the insights of the ‘theory of storage’ as presented in, for example, Brennan (1958). Interestingly, the model makes many predictions that are consistent with observed spot and futures data and that are consistent with the qualitative predictions made in the earlier papers on the theory of storage, and on which reduced-form models are based. Thus our model can provide a theoretical benchmark for functional form assumptions made in reduced-form models about the joint dynamics of spot and convenience yields.

Such a benchmark seems important for at least two reasons. First, it is well-known that most of the predictions of the real options literature hinge crucially on the specification of a convenience yield (e.g., Dixit and Pindyck (1994)).³ Indeed, following the standard intuition about the sub-optimality of early exercise of call options in the absence of dividends, if the convenience yield is negligible compared to storage costs, it may be optimal to not exercise real options. More generally, the functional form of the convenience yield can have important consequences on the valuation of real options (Schwartz (1997), Casassus and Collin-Dufresne (2005)). Second, equilibrium models deliver economically consistent long-term predictions. This may be a great advantage compared to reduced form models, which, due to the non-availability of data, may be hard to calibrate for long-term investment horizons.

The model presented here is related to existing literature and, in particular, builds upon Cox, Ingersoll and Ross (1985).⁴ Dumas (1992) follows CIR and sets up the grounds for analyzing dynamic GE models in two-sector economies with real frictions. He studies the real-exchange rate across two countries in the presence of shipping cost for transfers of capital.⁵ Recent applications of two-sector CIR economies along the lines of Dumas (1992) have been proposed by Kogan (2001) for studying irreversible investments and Mamaysky (2001) who studies interest rates in a durable and non-durable consumption goods economy. Richard and Sundaresan (1981) extends the CIR to a multi-good economy to study the theoretical relation between forward and futures prices. Unlike our paper, they do not allow for irreversible investment which produces most of the time variation in the economy. Similar non-linear production technologies to the one we use here have been proposed by Merton (1975) and Sundaresan (1984). Merton (1975) solves a one-sector stochastic growth model similar to the neoclassical Solow model where the two inputs are capital stock and labor force, while Sundaresan (1984) studies equilibrium interest rates with multiple consumption goods that are produced by technology that uses the consumption good and a capital good as inputs.⁶ Fixed adjustment

³Real Option Theory emphasizes the option-like characteristics of investment opportunities by including, in a natural way, managerial flexibilities such as postponement of investments, abandonment of ongoing projects, or expansions of production capacities (e.g. see the classical models of Brennan and Schwartz (1985), McDonald and Siegel (1986) and Paddock, Siegel and Smith (1988)).

⁴In fact, our model converges to a one -factor CIR production economy when oil is not relevant for the numeraire technology.

⁵Uppal (1993) presents a decentralized version of Dumas's economy.

⁶Surprisingly, there are not many models that use this type of production technologies in continuous time. Recently, Hartley and Rogers (2003) has extended the Arrow and Kurz (1970) two-sector

costs have been used in multiple research areas since the seminal (S,s) model of Scarf (1960) on inventory decisions. In the asset pricing literature, Grossman and Laroque (1990) uses fixed transaction costs to study prices and allocations in the presence of a durable consumption good.

There is an extensive literature that studies the effect of irreversibility and uncertainty on investments that is related to our model. Some examples of such contributions are Pindyck (1988), Bertola and Caballero (1994), Dixit and Pindyck (1994), Abel and Eberly (1994, 1996, 1997) and Baldursson and Karatzas (1997). More recently Kogan (2001, 2004) analyzes the effect of irreversible investment on asset prices. Some researchers have focused on the effect of fixed adjustment cost on investment behavior. Abel and Eberly (1994) incorporate fixed costs of investment and study the optimal investment rate as a function of the marginal value of a unit of installed capital (q). Caballero and Engel (1999) explains aggregate investment dynamics in a model that builds from the lumpy microeconomic behavior of firms facing stochastic fixed adjustment costs.

Our paper is also related to the work of Carlson, Khokher and Titman (2002), who propose an equilibrium model of natural resources. However, in contrast to our paper, they assume risk-neutrality, an exogenous demand function for commodity, and (the main friction in their model) that commodity is exhaustible, whereas in our paper commodity is essentially present in the ground in infinite supply but is costly to extract. Finally, Kogan, Livdan and Yaron (2005) identify a new pattern of futures volatility term structure that is inconsistent with standard storage models but can be explained within their model that exhibits investment constraints and irreversibility. Unlike our model, they take the demand side and risk-premia as exogenous and focus mainly on the implications for the volatility curve.

Section 2 presents the model. Section 3 characterizes equilibrium commodity prices in our benchmark model with irreversibility and costly oil production. Section 4 presents the empirical estimation of the model and discusses its economic implications. Finally, Section 5 concludes.

model to an stochastic framework and use this type of production technology with private and government capital as inputs.

2 The Model

We consider an infinite horizon production economy with two goods. The model extends the Cox, Ingersoll and Ross (CIR 1985a) production economy to the case where the production technology requires two inputs, which are complementary.

2.1 Representative Agent Characterization

There is a continuum of identical agents (i.e., a representative agent) which maximize their expected utility of intertemporal consumption, and have time separable constant relative risk-aversion utility given by

$$U(t, C) = \begin{cases} e^{-\rho t} \frac{C^{1-\gamma}}{1-\gamma} & \text{if } \gamma > 0, \gamma \neq 1 \\ e^{-\rho t} \log(C) & \text{if } \gamma = 1 \end{cases} \quad (1)$$

There is a single consumption good in our economy. Agents can consume the consumption good or invest it in a production technology. The production technology requires an additional input, the commodity, which is produced by a stock of oil wells. The dynamics of the stock of oil wells (Q_t) and the stock of consumption good (K_t) are described in equation (2) and (3) below:

$$dQ_t = -(\bar{i} + \delta)Q_t dt + \sigma_Q Q_t dw_{Q,t} + X_t dI_t \quad (2)$$

$$dK_t = (f(K_t, \bar{i}Q_t) - C_t) dt + \sigma_K K_t dw_{K,t} - \beta(X_t; Q_t, K_t) dI_t. \quad (3)$$

The oil ‘industry’ produces a flow of oil at rate \bar{i} and depreciates at rate δ .⁷ The representative agent can decide when and how many additional oil wells to build. We denote by I_t the investment time indicator, i.e., $dI_t = 1$ if investment occurs at date t and 0 else. Investment is assumed to be irreversible ($X_t \geq 0$) and costly in the sense that to build X_t new wells at t , the representative agents incurs a cost of $\beta(X_t; Q_t, K_t)$ of the numeraire good. We assume that the cost function has the following form:

$$\beta(X_t; Q_t, K_t) = \beta_K K_t + \beta_Q Q_t + \beta_X X_t \quad (4)$$

⁷For simplicity we assume that the extraction rate per unit time of each oil well is fixed. This is meant to capture the fact that it is very costly to increase or decrease the production flow of oil wells.

β_x is a variable cost paid per new oil well. $\beta_K K + \beta_Q Q$ represent the fixed costs incurred when investing. As is well-known, fixed costs ($\beta_K, \beta_Q > 0$) lead to an ‘impulse control’ optimization problem, where the optimal investment decision is likely to be lumpy (i.e., occurring at discrete dates).⁸ In contrast if only variable costs are present ($\beta_x > 0$ and $\beta_K = \beta_Q = 0$) then the optimal investment decision is an ‘instantaneous control’ which leads to a ‘local time,’ i.e., singular continuous, investment policy (e.g., Dumas (1991), Harrison (1990)). Below we assume that⁹

$$\beta_K, \beta_Q, \beta_x > 0.$$

Further, to insure that investment is feasible we assume that:¹⁰

$$\beta_K < 1 \text{ and } \beta_Q < \beta_x$$

We note that, while in our model investment immediately creates new oil wells (i.e., there is no time-to-build frictions in our model), one could potentially interpret the costs as a proxy for this friction.

The numeraire-good industry, equation (3), has a production technology that requires both the numeraire good and oil. Output is produced continuously at the mean rate

$$f(k, q) = \alpha k^{1-\eta} q^\eta.$$

As in Merton (1975) and Sundaresan (1984) we use the Cobb-Douglas production function (homogeneous of degree one and constant returns to scale). The parameter η represents the marginal productivity of oil in the economy. The output of this industry is allocated to consumption ($C_t \geq 0$), reinvested in numeraire good production, or used for investment to create more oil.¹¹

Uncertainty in our economy is captured by the Brownian motions $w_{Q,t}$ and $w_{K,t}$ which drive the diffusion term of the return of the technologies in equations (2)

⁸The assumption that the fixed component of the investment cost is scaled by the size of the economy, K_t and Q_t , ensures that the fixed cost does not vanish as the economy grows.

⁹The case where $\beta_K = \beta_Q = 0$ can be recovered by taking the appropriate limit as shown in Jeanblanc-Picque and Shiryaev (1995).

¹⁰We note that at the boundary when investment becomes optimal, the oil stock is valued at β_x . Thus for investment to be affordable we need $\beta_x X + \beta_K K + \beta_Q Q \leq K + \beta_x Q$ for some $X \geq 0$.

¹¹There is no storage of the numeraire good. Output that is not consumed, used in oil investment, or further production of the numeraire good depreciates fully.

and (3). We assume that there exists an underlying probability space (Ω, \mathbf{F}, P) satisfying the usual conditions, and where $\mathbf{F} = \{\mathcal{F}\}_{t \geq 0}$ is the natural filtration generated by the Brownian Motions.

Given our previous discussion it is natural to seek an investment policy of the form $\{(X_{T_i}, T_i)\}_{i=0,1,\dots}$ where $\{T_i\}_{i=0,\dots}$ are a sequence of stopping times of the filtration \mathbf{F} such that $I_t = \mathbf{1}_{\{T_i \leq t\}}$ and the X_{T_i} are \mathcal{F}_{T_i} -measurable random variables. Let us define the set of admissible strategies \mathcal{A} , as such strategies that lead to strictly positive consumption good stock process ($K_t > 0$ a.s.). Further, we restrict the set of allowable consumption policies \mathcal{C} to positive integrable \mathbf{F} adapted processes. Then the optimal consumption-investment policy of the representative agent is summarized by:

$$\sup_{C \in \mathcal{C}; \{(T_i, X_{T_i})\}_{i=0,\dots} \in \mathcal{A}} E_0 \left[\int_0^\infty e^{-\rho s} U(C_s) ds \right] \quad (5)$$

Let us denote by $J(t, K, Q) = \sup_{C; \mathcal{A}} E_t[\int_t^\infty e^{-\rho s} U(C_s) ds]$ the value function associated with this problem.

2.2 Sufficient Conditions for Existence of a Solution

Before characterizing the full problem 5 we establish sufficient conditions on the parameters for a solution to the problem to exist. We note that this is slightly different than in traditional models without fixed costs such as Dumas (1992) or Kogan (2002). Indeed, unlike in these models the no-transaction cost problem does not provide for a natural upper bound. Indeed, in our case, if we set $\beta_K = \beta_Q = \beta_X = 0$ the value function becomes infinite, since it is then optimal to build an infinite number of oil wells (at no cost). Thus unlike in these papers, it is natural to expect that sufficient conditions on the parameters for existence of the solution should depend on the marginal cost of building an oil well (as well as other parameters). Indeed, intuitively, if the marginal costs of an additional oil well is too low relative to the marginal productivity of oil in the K -technology one would expect the number of oil wells built (and thus the value function) to be unbounded. To establish reasonable conditions on the parameters we consider the case where there are only variable costs ($\beta_K = \beta_Q = 0$ and $\beta_X > 0$), but where the investment decision is perfectly

reversible. Let us denote $J_u(t, K, Q)$ the value function of the perfectly reversible investment/consumption problem. Clearly, the solution to that problem will be an upper bound to the value function of (5).

When the investment decision is perfectly reversible then it becomes optimal to adjust the stock of oil wells continuously so as to keep $\frac{J_{uQ}}{J_{uK}} = \beta_X$. This suggests that one can reduce the dimensionality of the problem, and consider as the unique state variable $W_t = K_t + \beta_X Q_t$ the ‘total wealth’ of the representative agent (at every point in time the agent can freely transform Q oil wells into $\beta_X Q$ units of consumption good and vice-versa). Indeed, the dynamics of W are:

$$dW_t = (\alpha K_t^{1-\eta} (\bar{i} Q_t)^\eta - C_t - \beta_X (\bar{i} + \delta) Q_t) dt + \sigma_K K_t dw_{K,t} + \beta_X \sigma_Q Q_t dw_{Q,t} \quad (6)$$

Since along each path, the agent can freely choose to adjust the ratio of oil to capital stock $Z_t = \frac{Q_t}{K_t}$, the Cobb-Douglas structure suggests that it will be optimal to maintain a constant ratio, $Z_t = Z^*$. We may rewrite the dynamics of W_t as

$$\frac{dW_t}{W_t} = (\mu_W^u(Z^*) - c_t^u) dt + \sigma_W^u(Z^*) dw_{W,t} \quad (7)$$

where $w_{W,t}$ is a standard Brownian motion and we define

$$C_t = c_t^u W_t, \quad (8)$$

$$\mu_W^u(Z) = \frac{\alpha (\bar{i} Z)^\eta - (\bar{i} + \delta) \beta_X Z}{1 + \beta_X Z} \quad (9)$$

and

$$\sigma_W^u(Z) = \frac{\sqrt{\sigma_K^2 + 2\rho_{KQ}\sigma_K\sigma_Q\beta_X Z + (\beta_X Z)^2\sigma_Q^2}}{1 + \beta_X Z}. \quad (10)$$

The proposition below verifies that if the function

$$f(Z) = \frac{\rho}{1-\gamma} - \mu_W^u(Z) + \gamma \frac{\sigma_W^u(Z)^2}{2} \quad (11)$$

admits a global minimum at Z^* such that

$$a^u := \frac{1}{\gamma} \left\{ \rho - (1-\gamma) \left(\mu_W^u(Z^*) - \gamma \frac{\sigma_W^u(Z^*)^2}{2} \right) \right\} > 0 \quad (12)$$

then the optimal strategy is indeed to consume a constant fraction of total wealth $c_t^u = a^u$ and to invest continuously so as to keep $Q_t/K_t = Z^*$.

Proposition 1 *Assume that there are no fixed costs ($\beta_K = \beta_Q = 0$), and that investment is costly ($\beta_X > 0$), but fully reversible. If the function $f(Z)$ defined in (11) admits a global minimum Z^* such that condition (12) holds then the optimal value function is given by*

$$J_u(t, K, Q) = e^{-\rho t} \frac{(a^u)^{-\gamma} (K + \beta_X Q)^{1-\gamma}}{1-\gamma} \quad (13)$$

The optimal consumption policy is

$$C_t^* = a^u (K_t^* + \beta_X Q_t^*) \quad (14)$$

and the investment policy is characterized by:

$$\frac{Q_t^*}{K_t^*} = Z^*. \quad (15)$$

Proof Applying Itô's lemma to the candidate value function we have:

$$\frac{dJ_u(t, K_t, Q_t) + U(t, C_t)dt}{J_u(t, K_t, Q_t)} = (1-\gamma) \{h(c_t) - f(Z_t)\} dt + (1-\gamma) \sigma_W^u(Z_t) dw_{W,t} \quad (16)$$

where we have set $C_t = c_t(K_t + \beta_X Q_t)$, $\sigma_W^u(Z)$ and $f(Z)$ are defined in equations (10) and (11), respectively, and we have defined:

$$h(c) = (a^u)^\gamma \frac{c^{1-\gamma}}{1-\gamma} - c.$$

Note that the function $h(c)$ is concave and admits a global maximum $c_t^* = a^u$ with $h(a^u) = \frac{a^u \gamma}{1-\gamma}$. Suppose the function $f(Z)$ is strictly convex and admits a global minimum at Z^* . Then, if we pick the constant a^u such that $h(a^u) = f(Z^*)$, we have for any c, Z :

$$h(c) - f(Z) \leq h(c^*) - f(Z^*) = 0$$

Thus integrating equation (16) we obtain:

$$J_u(T, K_T, Q_T) + \int_0^T U(t, C_t) dt \leq J_u(0, K_0, Q_0) + \int_0^T (1-\gamma) J_u(t, K_t, Q_t) \sigma_W^u(Z_t) dw_{W,t} \quad (17)$$

Taking expectation and using the fact that the stochastic integral is a positive local martingale we obtain:

$$E \left[J_u(T, K_T, Q_T) + \int_0^T U(t, C_t) dt \right] \leq J_u(0, K_0, Q_0) \quad (18)$$

Further we note that for when we choose the controls $c_t = a^u$ and $Z_t = Z^*$ then we obtain equality in equation (17) and further have:

$$\frac{dJ_u}{J_u} = -a^u dt + (1-\gamma) \sigma_W^u(Z^*) dw_{W,t} \quad (19)$$

which implies that the local martingale is a martingale and thus (18) obtains with equality. Further we have

$$\lim_{T \rightarrow \infty} E[J_u(T, K_T, Q_T)] = \lim_{T \rightarrow \infty} J_u(0, K_0, Q_0) e^{-a^u T} = 0$$

under the assumption (12). Letting $T \rightarrow \infty$ in (18) shows that our candidate value function indeed is the optimal value function and confirms that the chosen controls are optimal. \square

We note that in the case where $\eta = 0$, then Oil has no impact on the optimal decisions of the agent and the value function J_u is the typical solution one obtains in a standard Merton (1973) or Cox-Ingersoll-Ross (1985a) economy. In that case, the condition on the coefficient a^u becomes:

$$a_0 = \frac{1}{\gamma} \left\{ \rho - (1-\gamma) \left(\alpha - \gamma \frac{\sigma_K^2}{2} \right) \right\} > 0. \quad (20)$$

A lower bound to the value function is easily derived by choosing to never invest in oil wells (i.e., setting $dI_t = 0 \forall t$) and by choosing an arbitrary feasible consump-

tion policy $C_t^l = \alpha K_t^{1-\eta} (\bar{i} Q_t)^\eta$. Indeed, in that case we have:

$$\frac{dK_t}{K_t} = \sigma_K dw_{K,t} \quad (21)$$

It follows that if the following condition holds:

$$\begin{aligned} a^l := & \rho + (1 - \gamma) \left\{ (1 - \eta) \gamma \frac{\sigma_K^2}{2} + \eta \left(\bar{i} + \delta + \gamma \frac{\sigma_Q^2}{2} \right) \right\} \\ & + (1 - \eta) \eta (1 - \gamma)^2 \left\{ \frac{\sigma_K^2}{2} - \rho_{KQ} \sigma_K \sigma_Q + \frac{\sigma_Q^2}{2} \right\} > 0 \end{aligned} \quad (22)$$

then, we have

$$J_l(0, K_0, Q_0) := E \left[\int_0^\infty e^{-\rho t} \frac{(C_t^l)^{1-\gamma}}{1-\gamma} dt \right] = \frac{1}{a^l} \frac{(C_0^l)^{1-\gamma}}{1-\gamma} \quad (23)$$

We collect the previous results and a few standard properties of the the value function in the following proposition.

Proposition 2 *If $a^l, a^u > 0$, the value function of problem (5) has the following properties.*

1. $J_l(t, K, Q) \leq J(t, K, Q) \leq J_u(t, K, Q)$.
2. $J(t, K, Q)$ is increasing in K, Q .
3. $J(t, K, Q)$ is concave homogeneous of degree $(1 - \gamma)$ in Q and K .

For the following we shall assume conditions (12) and (22) are satisfied, i.e., that $a^l, a^u > 0$.

2.3 Optimal Consumption and Investment with Fixed Costs and Irreversibility

We first derive the HJB equation and appropriate boundary conditions, as well as the optimal consumption/investment policy based on a heuristic arguments due to the

nature of the optimization problem faced. Then we give a more formal verification argument.

First, since the solution depends on the time variable t only through the discounting effect in the expected utility function, we define the ‘discounted’ value function $J(K, Q)$, such that $J(K, Q, t) = e^{-\rho t} J(K, Q)$. Given that investment in new oil is irreversible ($X_t \geq 0$) and the presence of fixed costs, it is natural to expect that the optimal investment will be infrequent and ‘lumpy’ (e.g., Dumas (1991)) and defined by two zones of the state space $\{K_t, Q_t\}$: A *no-investment* region where $dI_t = 0$ and an *investment* region where $dI_t = 1$. This is analogous to the shipping cone in Dumas (1992), but with only one boundary because investment is irreversible.

2.3.1 Optimal Consumption Strategy in the No-Investment Region

When the state variables $\{K_t, Q_t\}$ are in the *no-investment* region, the numeraire good K can be consumed or invested in numeraire-good production. In this region, it is never transformed into new oil ($dI_t = 0$). That is; $J(K_t - \beta(X_t), Q_t + X) < J(K_t, Q_t)$ and it is not optimal to make any new investment in oil. The solution of the problem in equation (5) is determined by the following the Hamilton-Jacobi-Bellman (HJB) equation:

$$\sup_{\{C \geq 0\}} \{-\rho J + U(C) + \mathcal{D}J\} = 0 \quad (24)$$

where \mathcal{D} is the Itô operator

$$\begin{aligned} \mathcal{D}J(K, Q) \equiv & (f(K, \bar{i}Q) - C) J_K - (\bar{i} + \delta) Q J_Q \\ & + \frac{1}{2} \sigma_K^2 K^2 J_{KK} + \frac{1}{2} \sigma_Q^2 Q^2 J_{QQ} + \rho_{KQ} \sigma_K \sigma_Q K Q J_{KQ} \end{aligned} \quad (25)$$

with J_K and J_Q representing the marginal value of an additional unit of numeraire good and oil respectively. J_{KK} is the second derivative with respect to K .

The first order conditions for equation (24) characterize optimal consumption. At the optimum, the marginal value of consumption is equal to the marginal value of an additional unit of the numeraire good; that is

$$C_t^* = J_K^{-\frac{1}{\gamma}}. \quad (26)$$

Similarly, at the optimum, the marginal value of an additional unit of oil determines the representative agent's shadow price for that unit and we denote S_t as the equilibrium oil price. Define the marginal price of oil, S_t . That is, S_t solves $J(K_t, Q_t) = J(K_t + S_t\epsilon, Q_t - \epsilon)$. With a Taylor expansion, this implies

$$S_t = \frac{J_Q}{J_K}. \quad (27)$$

2.3.2 Optimal Investment Strategy

We assume in equation (4) that there is a fixed cost when investing in new oil. This increasing-returns-to-scale technology implies that the investment in new oil decision faced by the representative agent is an *Impulse Control* problem (see Harrison, Sellke, and Taylor (1983)). As is well known, these problems have the characteristic that whenever investment is optimal, the optimal size of the investment is non-infinitesimal and the state variables jump back into the *no-investment* region. Optimal investment is infrequent and lumpy.

The *investment* region is defined by $J(K_t - \beta(X_t), Q_t + X_t) \geq J(K_t, Q_t)$; that is when the value of additional oil exceeds its cost. Of course, along the optimal path, the only time when this inequality could be strict is at the initial date $t = 0$ with stocks $\{K_0, Q_0\}$.¹² Without loss of generality we assume that the initial capital stocks $\{K_0, Q_0\}$ are in the *no-investment* region. Let $J_1 = J(K_t^*, Q_t^*)$ be the value function before investment and $J_2 = J(K_t^* - \beta(X_t^*), Q_t^* + X_t^*)$ be the value function right after the investment is made. The *investment* zone is defined by the value matching condition.

$$J_1 = J_2 \quad (28)$$

There are three optimality conditions that determine the level of numeraire good K_t^* , the amount of oil Q_t^* , and the size of the optimal oil investment X_t^* at the investment boundary. We follow Dumas (1991) to determine these super-contact

¹²If this is the case, there is an initial lumpy investment that takes the state variables into the *no-investment* zone.

(smooth pasting) conditions.¹³

$$J_{1K} = (1 - \beta_K)J_{2K} \quad (29)$$

$$J_{1Q} = -\beta_Q J_{2K} + J_{2Q} \quad (30)$$

$$0 = -\beta_X J_{2K} + J_{2Q} \quad (31)$$

These equations imply that

$$(\beta_X - \beta_Q)J_{1K} - (1 - \beta_K)J_{1Q} = 0. \quad (32)$$

2.3.3 Reduction of number of state variables

Because the numeraire good production function is homogeneous of degree one ($f(k, q) = \alpha k^{1-\eta} q^\eta$) and the utility function is homogeneous of degree $(1 - \gamma)$, the value function inherits that property. This implies that the ratio of oil to the numeraire good is sufficient to characterize the economy. Indeed, let us define $j(z)$ as

$$J(K, Q) = \frac{K^{1-\gamma}}{1-\gamma} j(z) \quad (33)$$

where z is the log of the oil wells to numeraire-good ratio

$$z = \log\left(\frac{Q}{K}\right) \quad (34)$$

The dynamic process for z_t is obtained using a generalized version of Itô's Lemma.

$$dz_t = \mu_{zt} dt + \sigma_z dw_{z,t} + \Lambda_z dI_t^* \quad (35)$$

where $w_{z,t}$ is a standard Brownian motion,

$$\mu_{zt} = \left(-(\bar{i} + \delta) - \frac{1}{2}\sigma_Q^2\right) - \left(f(1, \bar{i}e^{z_t}) - c_t^* - \frac{1}{2}\sigma_K^2\right), \quad (36)$$

¹³For a discussion of value-matching and super-contact (smooth-pasting) conditions, see Dumas (1991), Dixit (1991) and Dixit (1993). If $\beta_K = \beta_Q = 0$ in equation (4) then we face an *Infinitesimal Control* problem. In this case, the optimal investment is a continuous regulator (Harrison (1990)), so that oil stock before and after investment are the same. In this case, equations (29) to (32) result directly from equation (28) as can be checked via a Taylor series expansion (as shown in Dumas (1991)). To solve this case we consider two additional 'super-contact' conditions $-J_{1QK} + \beta_X J_{1KK} = 0$ and $-J_{1QQ} + \beta_X J_{1KQ} = 0$.

$$\sigma_z = \sqrt{\sigma_K^2 - 2\rho_{KQ}\sigma_K\sigma_Q + \sigma_Q^2}, \quad (37)$$

$$\Lambda_z = z_2 - z_1, \quad (38)$$

and the consumption rate, $c_t^* = C_t^*/K_t^*$, is a function of z_t .

The *no-investment* and *investment* regions are also characterized solely by z_t . Using the same subscripts as in equation (28), define $z_1 = \log(Q_t^*) - \log(K_t^*)$ as the log oil to numeraire-good ratio just prior to investment. Similarly, define $z_2 = \log(Q_t^* + X_t^*) - \log(K_t^* - \beta(X_t^*))$ as the log ratio immediately after the optimal investment in oil occurs. z_1 defines the *no-investment* and *investment* region. When $z_t > z_1$ it is optimal to postpone investment in new oil. If the state variable z_t reaches z_1 , an investment to increase oil stocks by X_t^* is made. The result is that the state variable jumps to z_2 which is inside the *no-investment* region. Given the investment cost structure in equation (4), the proportional addition to oil, x_t , is just a function of z_1 and z_2 .

$$x_t^* = \frac{X_t^*}{Q_t^*} = \frac{e^{-z_1} - e^{-z_2} - (\beta_K e^{-z_1} + \beta_Q)}{e^{-z_2} + \beta_X} \quad (39)$$

The jump in oil wells is

$$\frac{Q_2}{Q_1} = 1 + x^* \quad (40)$$

and, we can express the jump in the consumption good stock simply as:

$$\frac{K_2}{K_1} = \frac{1 - \beta_K + e^{z_1}(\beta_X - \beta_Q)}{1 + \beta_X e^{z_2}} \quad (41)$$

Finally, the optimal consumption from (26) can be rewritten in terms of j as:

$$c_t^* = \frac{C_t^*}{K_t^*} = \left(j(z_t) - \frac{j'(z_t)}{(1-\gamma)} \right)^{-\frac{1}{\gamma}} \quad (42)$$

Plugging this into the Hamilton-Jacobi-Bellman in equation (24) we obtain one-dimensional ODE for the function j .

$$\begin{aligned} \theta_0 j(z) + \theta_1 j'(z) + \theta_2 j''(z) + \gamma \left(j(z) - \frac{j'(z)}{1-\gamma} \right)^{1-\frac{1}{\gamma}} \\ + \alpha(\bar{i} e^z)^\eta ((1-\gamma)j(z) - j'(z)) = 0 \end{aligned} \quad (43)$$

where

$$\theta_0 = -\rho - \gamma(1 - \gamma)\frac{\sigma_K^2}{2}, \quad \theta_1 = -(\bar{i} + \delta) + \gamma\sigma_K(\sigma_K - \rho_{KQ}\sigma_Q) - \frac{\sigma_z^2}{2}, \quad \theta_2 = \frac{\sigma_z^2}{2} \quad (44)$$

To determine the investment policy, $\{z_1, z_2\}$, the value-matching condition of equation (28) becomes:

$$(1 + e^{z_2}\beta_X)^{1-\gamma}j(z_1) - (1 - \beta_K + e^{z_1}(\beta_X - \beta_Q))^{1-\gamma}j(z_2) = 0 \quad (45)$$

Lastly, using the homogeneity there are only two super-contact conditions to determine that capture equations (29), (30), and (31).¹⁴ They are

$$(1 - \gamma)e^{z_1}(\beta_X - \beta_Q)j(z_1) - (1 - \beta_K + e^{z_1}(\beta_X - \beta_Q))j'(z_1) = 0 \quad (46)$$

$$(1 - \gamma)e^{z_2}\beta_Xj(z_2) - (1 + e^{z_2}\beta_X)j'(z_2) = 0 \quad (47)$$

The following proposition summarizes the above discussion and offers a verification argument. Let us define the functions:

$$a(z) := j(z) - \frac{j'(z)}{1 - \gamma} \quad (48)$$

$$F(x, y) := \left(\frac{1 - \beta_K + e^x(\beta_X - \beta_Q)}{1 + \beta_X e^y} \right)^{1-\gamma} \frac{j(y)}{1 - \gamma} - \frac{j(x)}{1 - \gamma} \quad (49)$$

$$(50)$$

Proposition 3 *Suppose that we can find two constants z_1, z_2 ($0 \leq z_1 \leq z_2$) and a function $j(\cdot)$ defined on $[z_1, \infty)$, which solve the ODE given in equation (43) with boundary conditions (45), (46), and (47), such that the following holds:*

$$0 < a(z)^{-1/\gamma} < M_1 \quad (51)$$

$$0 < \frac{a(z)}{j(z)} < M_2 \quad (52)$$

¹⁴In a similar way, if $\beta_K = \beta_Q = 0$ the two super-contact conditions presented in footnote (13) become the same condition $(1 + (1 - \gamma)e^{z_1}\beta_X)j'(z_1) - (1 + e^{z_1}\beta_X)j''(z_1) = 0$.

$$F(x, y) \leq 0, \quad \forall y \geq x \geq z_1 \quad (53)$$

$$0 = F(z_1, z_2) \geq F(z_1, y), \quad \forall y \geq z_1 \quad (54)$$

where M_1, M_2 are constants.

Then the value function is given by

$$J(t, K, Q) = e^{-\rho t} \frac{K^{1-\gamma}}{1-\gamma} j(z) \quad (55)$$

where $z = \log \frac{Q}{K}$. Further the optimal consumption policy is to set

$$c(z_t) = a(z_t)^{-1/\gamma}.$$

The optimal investment policy consists of a sequence of stopping times and investment amounts, $\{(T_i, X_{T_i})\}_{i=0,2,\dots}$ given by $T_0 = 0$ and:

- If $z_0 \leq z_1$ then invest (to move z_0 to z_2):

$$X_0^* = Q_0 \frac{e^{-z_0}(1 - \beta_K) - e^{-z_2} - \beta_Q}{e^{-z_2} + \beta_X} \quad (56)$$

Then start with new initial values for the stock of consumption good $K_0 - \beta(X_0^*, K_0, Q_0)$ and stock of oil wells $Q_0 + X_0^*$.

- If $z_0 > z_1$ then set $X_0^* = 0$ and define the sequence of \mathbf{F} -stopping times:

$$T_i = \inf \{t > T_{i-1} : z_{t-} = z_1\} \quad i = 1, 2, \dots \quad (57)$$

and corresponding \mathcal{F}_{T_i} -measurable investments in oil wells:

$$X_{T_i}^* = Q_{T_i} \frac{e^{-z_1}(1 - \beta_K) - e^{-z_2} - \beta_Q}{e^{-z_2} + \beta_X}. \quad (58)$$

Proof We define our candidate value function as $J(K, Q, t) = e^{-\rho t} \frac{K^{1-\gamma}}{(1-\gamma)} j(z)$, where $z = \log(Q/K)$ as before and where we define $j(z)$ as in the proposition for $z \geq z_1$ and where we set

$$j(z) = \left(\frac{1 - \beta_K + e^z(\beta_X - \beta_Q)}{1 + \beta_X e^{z_2}} \right)^{1-\gamma} j(z_2), \quad \forall z < z_1.$$

Applying the generalized Itô's lemma to our candidate value function for some arbitrary controls we find:

$$\begin{aligned}
dJ(t, K_t, Q_t) + U(t, C_t)dt &= e^{-\rho t} K_t^{1-\gamma} \left\{ \left[\frac{\hat{\theta}_0(z_t)j(z_t) + \hat{\theta}_1(z_t)j'(z_t) + \theta_2 j''(z_t)}{1-\gamma} + \frac{(c_t)^{1-\gamma}}{1-\gamma} - a(z_t)c_t \right] dt \right. \\
&\quad \left. + a(z_t)\sigma_K dw_{K,t} + \{j(z_t) - a(z_t)\}\sigma_Q dw_{Q,t} + F(z_t^-, z_t) \right\} \quad (59)
\end{aligned}$$

where for simplicity we have defined $\hat{\theta}_0(z) = \theta_0 + (1-\gamma)\alpha(\bar{i}e^z)^\eta$ and $\hat{\theta}_1(z) = \theta_1 - \alpha(\bar{i}e^z)^\eta$ and $C_t = c_t K_t$.

Now the definition of the function $j(z)$ implies that

$$\frac{\hat{\theta}_0(z)j(z) + \hat{\theta}_1(z)j'(z) + \theta_2 j''(z)}{1-\gamma} + \sup_c \left[\frac{(c)^{1-\gamma}}{1-\gamma} - a(z)c \right] \begin{cases} = 0 & \forall z \geq z_1 \\ < 0 & \forall z < z_1 \end{cases}$$

Further, $F(x, y) \leq 0 \forall x \leq y$ with equality only if $x \leq z_1$ and $y = z_2$. Thus we have that for arbitrary controls

$$\begin{aligned}
J(T, K_T, Q_T) + \int_0^T U(t, C_t)dt &\leq J(0, K_0, Q_0) + \int_0^T e^{-\rho t} K_t^{1-\gamma} a(z_t)\sigma_K dw_{K,t} + \\
&\quad \int_0^T e^{-\rho t} K_t^{1-\gamma} \{j(z_t) - a(z_t)\}\sigma_Q dw_{Q,t}. \quad (60)
\end{aligned}$$

Taking expectation (using the fact that the stochastic integral is a positive local martingale hence a supermartingale) we obtain that for arbitrary controls

$$E \left[J(T, K_T, Q_T) + \int_0^T U(t, C_t)dt \right] \leq J(0, K_0, Q_0) \quad (61)$$

For the controls proposed in the proposition equation (60) holds with equality. Further, we have for these particular controls:

$$\frac{dJ(t, K_t, Q_t)}{J(t, K_t, Q_t)} = -a(z_t)^{-1/\gamma} \frac{a(z_t)}{j(z_t)} dt + \sigma_J \left(\frac{a(z)}{j(z)} \right) dw_{J,t} \quad (62)$$

where $w_{J,t}$ is a standard Brownian motion and

$$\sigma_J(x) = (1-\gamma) \sqrt{x^2 \sigma_K^2 + 2x(1-x)\rho_{KQ}\sigma_K\sigma_Q + (1-x)^2 \sigma_Q^2}. \quad (63)$$

This implies that (using the assumptions that $\frac{a(z)}{j(z)} \in (0, M_1)$ and $a(z_t)^{-1/\gamma} \in (0, M_2)$) the stochastic integral in (60) is a martingale and that

$$\lim_{T \rightarrow \infty} E[J(T, K_T, Q_T)] = \lim_{T \rightarrow \infty} J(0, K_0, Q_0) \tilde{E} \left[e^{-\int_0^T a(z_t)^{-1/\gamma} \frac{a(z_t)}{j(z_t)} dt} \right] = 0.$$

where we have defined a new measure $\tilde{P} \sim P$ by the Radon-Nikodym derivative $\frac{d\tilde{P}}{dP} = e^{-\int_0^T \frac{1}{2} \sigma_J \left(\frac{a(z_t)}{j(z_t)} \right)^2 dt + \int_0^T \sigma_J \left(\frac{a(z_t)}{j(z_t)} \right) dw_{J,t}}$.

□

The Hamilton-Jacobi-Bellman equation with boundary conditions does not have (to the best of our knowledge) a closed-form solution. In Appendix A we sketch the numerical technique used to solve this system of equations.

In the following we characterize the equilibrium asset prices and oil prices.

3 Equilibrium Prices

The solution to the representative agent's problem of equation (5) is used to characterize equilibrium prices.¹⁵ We first describe the pricing kernel and financial asset prices. Next, we use the marginal value of a unit of oil, as in equation (27), to characterize the equilibrium spot-price of oil. Finally, we characterize the structure of oil futures' prices. Interestingly, with only a single source of diffusion risk, the model produces prices that can have both jumps and a regime-shift pattern.

3.1 Asset Prices and the Pricing Kernel

Since in our model markets are dynamically complete, the pricing kernel is characterized by the representative agent's marginal utility (see Duffie (1996)). First,

¹⁵We do not consider conditions under which the representative agent's problem we solve corresponds to the outcome of a decentralized competitive equilibrium with multiple agents. For the case where there are no fixed costs the structure of our framework is similar to Dumas (1992) and Uppal (1993) so we conjecture their results apply. For the case with fixed costs, the problem is complicated by 'local' non-convexity of the production function (e.g., Guesnerie (1975)). We leave the problem for future research and proceed under the assumption of a unique maximizing agent.

define the risk-free money-market account whose price is B_t . The process for the money market price is

$$\frac{dB_t}{B_t} = r_t dt + \Lambda_B dI_t \quad (64)$$

where r_t is the instantaneous risk-free rate in the *no-investment* region. Λ_B is a jump in financial market prices that can occur when the lumpy investment in the oil industry occurs. Note that the jumps, $\Lambda_B dI_t$, occur at stochastic times, but since they occur based on the oil-investment decision, they are predictable.

The pricing kernel for our economy satisfies

$$\frac{d\xi_t}{\xi_t} = -\frac{dB_t}{B_t} - \lambda_{K,t} dw_{K,t} - \lambda_{Q,t} dw_{Q,t} \quad (65)$$

with $\xi_0 = 1$. In the *no-investment* region ($dI_t = 0$), the pricing kernel is standard. However, when investment occurs ($dI_t = 1$), there is a singularity in the pricing kernel (through the $\Lambda_B dI_t$ term in dB_t). This is consistent with Karatzas and Shreve (1998), who show that in order to rule out arbitrage opportunities, all financial assets in the economy must jump by the same amount Λ_B .¹⁶

Proposition 4 *In equilibrium, financial assets are characterized by:*

$$\xi_t = e^{-\rho t} \frac{J_K(K_t, Q_t)}{J_K(K_0, Q_0)} \quad (66)$$

$$r_t = f_1(K_t, \bar{Q}_t) - \sigma_K (\lambda_{K,t} + \rho_{KQ} \lambda_{Q,t}) \quad (67)$$

$$\lambda_{K,t} = -\sigma_K \frac{K_t J_{KK}}{J_K} \quad (68)$$

$$\lambda_{Q,t} = -\sigma_Q \frac{Q_t J_{KQ}}{J_K} \quad (69)$$

$$\Lambda_B = -\frac{\beta_K}{1 - \beta_K} \quad (70)$$

where $f_1(\cdot, \cdot)$ is the first derivative of the production function with respect its first argument. Moreover, the equilibrium interest rate and market prices of risk are only functions of the state variable z_t , i.e., $r_t = r(z_t)$, $\lambda_{K,t} = \lambda_K(z_t)$ and $\lambda_{Q,t} = \lambda_Q(z_t)$.

¹⁷

¹⁶The oil commodity price, S_t , is not a financial asset and may, as is described later, jump by a different amount at the point of oil-industry investment.

¹⁷We decide to present these variables under $\{K_t, Q_t\}$ rather than under z_t to show that these expressions are similar to the standard results in a CIR economy.

Proof Using that $\xi_t \propto U_C(t, C_t)$ and the first order condition of equation (24) with respect to consumption (and setting $\xi_0 = 1$), we obtain equation (66). To get the interest rate, market prices of risk dynamics, we apply the generalized Itô's lemma to the pricing kernel equation . \square

The interest rate in the *no-investment* region is the marginal productivity of the numeraire good adjusted by the risk of the technology as in Cox, Ingersoll Jr., and Ross (1985) (CIR). The only difference in our model is the effect of the non-linear technology $f(k, q)$. Similarly, the price of risk in equations (68) and (69) is driven by the shape of the productivity of the numeraire good. Interestingly, there can be a jump (predictable) in asset prices that occurs each time investment in oil is optimal ($dI_t = 1$). From equation (29) we can calculate the size of the jump in the stochastic discount factor and note that it depends only on the oil investment cost structure. In particular, note that since $0 \leq \beta_K < 1$, financial asset prices jump down $\Lambda_B \leq 0$ if $\beta_K \neq 0$. Effectively, the fixed investment costs create a wealth effect, which increases marginal utility of the representative agent. Since financial asset prices normalized by marginal utility must be martingales to avoid arbitrage opportunities, prices must jump down to offset the jump in marginal utility. In the case where $\beta_K = 0$ both the state price density and financial asset prices are continuous ($\Lambda_B = 0$).

3.2 Oil Spot Prices

The market-clearing spot price of oil is determined by the marginal value of a unit of oil along the representative agent's optimal path. This shadow price, from equation (27), is a function of the ratio of oil to numeraire good state variable, z_t :

$$S_t = \frac{J_Q}{J_K} = \frac{e^{-z_t} j'(z_t)}{(1 - \gamma)j(z_t) - j'(z_t)} \quad (71)$$

To characterize the oil spot price behavior, consider the spot price at the investment boundary, z_1 . From the smooth-pasting condition in equation (31), the oil price immediately after new investment is

$$S_{2,t} = \beta_X \quad (72)$$

That is, oil's value is equal to the marginal cost of new oil at the time of investment. Immediately prior to new investment, the condition in equation (32) implies that

$$S_{1,t} = \frac{\beta_X - \beta_Q}{1 - \beta_K} \quad (73)$$

which depends on both the fixed and marginal cost of acquiring new oil. Therefore, at the point of investment, the oil price jumps by the constant $\Lambda_S S_{1,t}$ where

$$\Lambda_S = \frac{\beta_Q - \beta_K \beta_X}{\beta_X - \beta_Q} \quad (74)$$

Since oil is not a traded financial asset, the jump in the price of oil can be different than the Λ_B jump in financial prices. Only when there are no fixed costs (i.e., when investment is not lumpy) to investing in oil ($\beta_K = \beta_Q = 0$) are both prices continuous. In general, oil prices jump by a different amount than financial asset prices. It is possible to generate continuous asset prices and discontinuous oil prices ($\beta_K = 0, \beta_Q > 0$). In that case, note that the oil prices jumps up at the time of investment $\lambda_S = \beta_Q > 0$. Alternatively, if $\beta_Q = \beta_K \beta_X$, then oil prices have no jump. In this case, the cost of oil investment from equation (4) is $\beta(X_t; Q_t, K_t) = \beta_K(K_t + \beta_X Q_t) + \beta_X X_t$. Since $S_{2,t} = \beta_X$, this implies that the fixed cost component of investing in new oil wells is proportional to aggregate wealth in the economy at the time of investment. The simulations that follow illustrate this case.

3.3 Oil Futures Prices

Given the equilibrium processes for spot prices and the pricing kernel, we can characterize the behavior of oil futures prices in our model. Define $F(z, t, T)$ as the date- t futures contract that delivers one unit of oil at date T given that the state of the economy is z .¹⁸ The stochastic process for the futures price is

$$\frac{dF_t}{F_t} = \mu_{F,t} dt + \sigma_{FK,t} dw_{K,t} + \sigma_{FQ,t} dw_{Q,t} + \Lambda_F dI_t \quad (75)$$

where $\mu_{F,t}$, $\sigma_{FK,t}$, $\sigma_{FQ,t}$ and Λ_F are determined in equilibrium following Cox, Ingersoll Jr., and Ross (1985).

¹⁸Since the futures contracts are continuously market-to-market, the value of the futures contract is zero.

Proposition 5 *The equilibrium futures price $F(z, t, T)$ in equation (75) satisfies $\mu_{F,t} = \sigma_{FK,t}(\lambda_{K,t} + \rho_{KQ}\lambda_{Q,t}) + \sigma_{FQ,t}(\lambda_{Q,t} + \rho_{KQ}\lambda_{K,t})$ and $F(z_1, t, T) = F(z_2, t, T)$, implying $\Lambda_F = 0$ and the following partial differential equation*

$$\frac{1}{2}\sigma_K^2 F_{zz} + (\mu_z + \sigma_K(\lambda_{K,t} + \rho_{KQ}\lambda_{Q,t}) - \sigma_Q(\lambda_{Q,t} + \rho_{KQ}\lambda_{K,t}))F_z + F_t = 0 \quad (76)$$

with boundary condition

$$F(z, T, T) = S(z). \quad (77)$$

In many commodity pricing models the second factor used to describe futures prices is the *net convenience yield* (see Gibson and Schwartz (1990)). Typically, this assumption is motivated as a benefit for holding stocks (net of any storage or depreciation costs). In these models, backwardation (downward sloping forward curve) is implied by the convenience yield. For example, Casassus and Collin-Dufresne (2005) present a reduced-form model with mean reversion in commodity prices. When the spot price is high, the convenience yield is high and pushes the spot price back toward a long-term mean (under the risk-neutral measure).¹⁹

The convenience yield is defined as the implicit return to the holder of the commodity, but not to the owner of a futures contract. If the commodity S_t were a traded financial asset, then the convenience yield would be the monetary dividend flow that would have to accrue to its holder to guarantee the absence of arbitrage. This is analogous to calculating the implicit convenience yield from the “cost-of-carry” and the slope of the futures curve as in Routledge, Seppi, and Spatt (2000).

The following proposition presents the equilibrium cumulative convenience yield in our economy:

Proposition 6 *The implicit cumulative net convenience yield Y_t has the following dynamics*²⁰:

$$dY_t = y_t S_t dt + \Lambda_Y S_t dI_t \quad (78)$$

¹⁹Mean-reversion in prices under the historical measure can also be due to time-variation in risk premia.

²⁰The continuous component of the convenience yield y_t is a function only of z_t , but as before, we prefer to present this variable under $\{K_t, Q_t\}$ rather than under z_t to deliver better economic intuition from the result. In fact, the variable f_2 would be expressed in terms of f_z which has a less clear economic meaning.

where

$$y_t = \frac{\bar{i}}{S_t}(f_2(K_t, \bar{i}Q_t) - S_t) - \delta - \sigma_Q(\theta_{Q,t} + \rho_{KQ}\theta_{K,t}) \quad (79)$$

$$\theta_{Q,t} = -\sigma_Q \frac{QJ_{QQ}}{J_Q} \quad (80)$$

$$\theta_{K,t} = -\sigma_K \frac{KJ_{KQ}}{J_Q} \quad (81)$$

$$\Lambda_Y = \Lambda_B - \Lambda_S. \quad (82)$$

An alternative representation of the continuous part of the convenience yield is:

$$y_t = -E_t \left[\frac{d(e^{-\rho t} J_Q)^c}{e^{-\rho t} J_Q} \right] \quad (83)$$

where X^c denotes the continuous component of the process X .

Proof The convenience yield is determined implicitly from equilibrium prices using the no-arbitrage condition for tradable assets

$$E_t^* \left[\frac{dS_t}{S_t} \right] = \frac{dB_t}{B_t} - \frac{dY_t}{S_t} \quad (84)$$

where E_t^* is the expectation under the equivalent martingale measure. The relation between this expectation and the expectation under physical measure is:

$$E_t \left[\frac{dS_t}{S_t} \right] = E_t^* \left[\frac{dS_t}{S_t} \right] - \frac{d\xi_t}{\xi_t} \frac{dS_t}{S_t}. \quad (85)$$

Applying Itô's lemma to equation (27) and using equation (85) we can determine dY_t from equation (84). For the second part we observe that the spot price is proportional to $e^{-\rho t} J_Q / \xi_t$. Applying Itô's lemma to this expression and using equation (85) we can obtain that

$$E_t^* \left[\frac{dS_t^c}{S_t} \right] = E_t \left[\frac{d(e^{-\rho t} J_Q)^c}{e^{-\rho t} J_Q} - \frac{d\xi_t^c}{\xi_t} \right] \quad (86)$$

Using equation (65) and a continuous version of equation (84) we obtain equation (83). \square

Equation (78) shows two components of the convenience yield. The first is the continuous component y_t which accrues continuously. It depends on the marginal productivity of oil in production. The endogenous convenience yield is increasing in f_2 and, hence, is increasing in the oil's importance as a productive input, η . Also, y_t is decreasing in the commodity inventories, Q_t . This implies that the convenience yield is higher near the investment region.

The second component of convenience yield is the predictable jump that occurs in prices at the time of oil investment. If oil were a traded asset then Λ_Y would represent pure arbitrage profits that can be locked in by trading oil prices against any other financial asset. Instead, the commodity is not a financial asset, and its 'price' is the shadow value to the consumers of using it as an input to production.

Finally, the second part of Proposition 6 gives a clear interpretation of the convenience yield in terms of the marginal productivity of a unit of oil in excess of its financial cost S_t , its physical depreciation δ and an adjustment for supply shock risk. Comparing equation (79) with that for the short rate r in equation (67) we see a strong resemblance. Effectively, the convenience yield y can be interpreted as an interest rate in an economy where we switch numeraire and use the commodity instead of the consumption good. In that economy, r_t would become a 'convenience yield' on the consumption good.²¹

4 Model Estimation

In this section we want to understand the empirical properties of the model in Sections 2 and 3. First, we use crude oil derivatives data, interest rates and economic aggregates to estimate our model using a simulation based technique. Then we discuss the implications of the model for commodity prices. In particular, we find that two regimes arise in our economy due to the fixed cost components of the investment. Finally, we do a simple estimation of a regime-switching model that supports our findings.

²¹This isomorphism between convenience yield and interest rates is made by Richard and Sundaresan (1981) in a multi-good economy.

4.1 Moments and SMM Estimation

We implement the Simulated Method of Moments (SMM) of Duffie and Singleton (1993). The main idea is to pick parameters that minimize the *weighted* distance between a set of model implied unconditional moments, $G_Z(\psi)$, and their corresponding moment conditions from the data, G_T .

We are mainly interested in the behavior of crude oil prices and how they are related to macro variables such as interest rates, oil production, output (GDP) and consumption. For this reason, we consider a vector g_t of micro and macro variables in a sample of size T . The set of unconditional moments are the sample averages in our dataset, i.e., $G_T = \frac{1}{T} \sum_{t=1}^T g_t$. In particular, our dataset is composed by the following series: (i) crude oil futures prices for different maturities, (ii) volatility of futures returns, (iii) aggregate consumption-output of capital ratio, (iv) aggregate consumption of oil-output ratio and (v) real interest rates. It is important to note that we include futures prices for different maturities to match a full term structure of prices and volatilities.

To obtain the model implied moment conditions, we simulate our economy for a given set of parameters $\hat{\psi}$. Recall that the economy is uniquely determined by the state variable z_t defined in equation (34). The dynamics of z_t is endogenous since it depends on the optimal consumption and investment strategies. For this reason, we first need to solve the Hamilton-Jacobi-Bellman equation in (43) with the numerical technique described in the Appendix. Then, we simulate to obtain the implied density function of z , $f(z; \hat{\psi})$.²² We calculate the implied variables used for the moment conditions as functions of z , $g(z; \hat{\psi})$, and using the simulated density of z_t we compute the model implied moments as $G_Z(\hat{\psi}) = E^Z[g(z; \hat{\psi})] \approx \int g(z) f(z) dz$.²³

Due to the high computational burden of the simulation approach and numerical solution of the HJB, we estimate only a subset the parameters $\hat{\psi}$ (the remaining ones are calibrated using available studies). The SMM parameter estimates solve

²²To ensure convergence, we discretize the state space of $z_t \in [-20, 10]$ in a grid of 15,000 points and then simulate weekly samples of the state variable for 10^5 years.

²³In our model, the aggregate consumption-output of capital ratio is defined as $C_t/f(K_t, \bar{i} Q_t)$ and the aggregate consumption of oil-output ratio is $\bar{i} Q_t S_t/f(K_t, \bar{i} Q_t)$. Both ratios are only functions of of the state variable z_t .

the following problem²⁴

$$\psi^* = \underset{\psi \in \Psi}{\operatorname{argmin}} [G_Z(\psi) - G_T]' W_T [G_Z(\psi) - G_T] \quad (87)$$

where W_T is weighting or distance matrix. We choose W_T to be the inverse of the diagonal of the unbiased estimate covariance matrix of the sample averages. This weighting matrix ensures that the scale of each moment condition is the same, and gives more weight to less volatile moments.²⁵

4.2 Data

For the SMM estimation we use quarterly time series from Q4/1990 to Q4/2004. We build the series of crude oil futures prices and interest rates, private consumption, GDP and petroleum consumption from OECD countries. Crude oil futures prices are obtained from the New York Mercantile Exchange (NYMEX). We use contracts with maturities of 1, 3, 6, 9, 12, 18, 24, 30 and 36 months. If a specific contract is missing, we select the one with the nearest maturity. For the quarterly figures we use the average prices within that period. To get the (annualized) volatility of futures returns, we sample quarterly observations of a GARCH(1,1) estimated separately for each (log) futures series using weekly prices. The volatilities time series are necessary to use the weighting matrix described above. Consumption and output data is from www.oecd.org. The aggregate data is available from Q1/1995 for all OECD members (30 countries). For the initial years we build a proxy for the series with the G7 countries data available from the same site. We assume that the GDP ratio of the G7 countries and all OECD members was constant from Q1/1990 to Q1/1995. Petroleum consumption data for OECD countries is from the U.S. Energy Information Administration site (www.eia.doe.gov). Finally, the interest rate data is from Federal Reserve FRED site (research.stlouisfed.org/fred2). To build the real interest rate time series we also use the CPI series, which are obtained from the same site.

The “Historical data” group in Table 2 shows the statistics of our sample. For

²⁴The feasible set of parameters Ψ are all ψ such that the existence of the value function is guaranteed (see equations (12) and (22)).

²⁵Cochrane (2001) discusses the pros and cons of using different weighting matrices for the estimation.

the period considered the average futures curve is downward sloping, implying a high degree of backwardation in crude oil prices (in our dataset 70% of the times the 6-months maturity contract is below the 1-month maturity contract). The average volatility term structure of futures returns is also downward sloping, implying high degrees of mean-reversion (“Samuelson effect”). The average annual consumption-GDP ratio in the data is 61.7% and is stable with a slightly increasing tendency. The average annual consumption of oil-output ratio in the data is 1.6% and its volatility is 0.004. This ratio is very stable most of the time, but it peaked in 1981 at almost 6%. The average annual real interest rate for this period was 1.5% with a volatility of 1.5%. The interest rate becomes negative (in real terms) in the last two years of the sample period. The low standard deviation of the macro variables compared to the ones from crude oil prices and volatilities, yield a higher weight for the macro moments in the SMM estimation.

4.3 Parameter Estimates

The complete set of parameters in our economy is given by $\hat{\psi} = \{\alpha, \eta, \bar{i}, \delta, \sigma_K, \sigma_Q, \rho_{KQ}, \beta_K, \beta_Q, \beta_X, \rho, \gamma\}$. These are too many parameters for the simulation-based estimation technique. We choose the productivity factor α , the input ratio \bar{i} , the volatility of capital σ_K , the investment fixed cost β_K and the risk aversion γ to be the free parameters and set the others to reasonable numbers.²⁶ This leaves the parameter *search* space as $\psi = \{\alpha, \bar{i}, \sigma_K, \beta_K, \gamma\} \subset \hat{\psi}$.

The oil share of income η is set to 0.04 which is consistent with recent RBC studies that include energy as a production factor (see Finn (1995), Finn (2000) and Wei (2003)). The depreciation rate of the commodity stock δ is set to 0.2, which implies an average storage costs of around \$4 per barrel. This figure is similar to the one used in Ross (1997). The marginal production cost of oil is fixed at \$12.5 per barrel. The fixed cost component β_Q is chosen such that there is no jump in prices at the investment boundary. For the volatility of sector Q, we calculate the standard deviation of annual changes of petroleum consumption in our dataset

²⁶While the selection of the parameters is somewhat arbitrary, it was driven by the extent to which we could find existing studies that help with the calibration, and by the fact that we are mostly interested in estimating the cost parameters $\beta_K, \beta_Q, \beta_X$ which are crucial for the predictions of the model.

(i.e., $\sigma_Q = 0.013$). Finally, we assume that the shocks to capital and oil stocks are independent ($\rho_{KQ} = 0$) and that the patience factor ρ is 0.05.

The parameter estimates ψ^* are marked with an asterisk in Table 1. The historical moments and their implied value using the SMM estimates are shown in Table 2. Figure 3 shows the plots for the mean and volatility of futures prices. We can see that the model successfully matches the unconditional moments of the futures data. Specifically, it generates reasonable average futures prices and average volatilities of futures returns across maturities. The model implies a decreasing and convex average futures curve, but with a smaller average degree of backwardation than the one from our sample. Also, the model-implied standard deviations (SD) of futures prices are very close to their sample counterparts.²⁷ The GARCH volatility term structure has a better fit than the futures curve, mainly because of the implied degrees of mean reversion in our model. The macro moments are matched almost perfectly, because of their higher weight in the SMM estimation technique. The expected consumption/GDP ratio is 61.7%, the expected petroleum consumption/GDP ratio is 1.6% and the expected real interest rate is 1.5%.

4.4 Oil Spot Prices

Figure 4 plots the equilibrium oil price as a function of the state variable, z_t , the log ratio of oil stocks to the numeraire good. The oil price is driven by both current and anticipated oil stocks. In the *no-investment* region, the supply of oil depletes as oil is used in the production of the numeraire good. Far from the investment trigger, the decreased supply of oil increases the price. The marginal cost of adding new oil is β_x (equation (4)). The fixed cost involved in adding new oil stocks implies that it is not optimal to make a new investment as soon as the spot price (marginal benefit of oil) reaches β_x . Therefore the spot price rises above β_x as oil is depleted. However, closer to the investment threshold, the oil price reflects the expected lumpy investment in new oil (i.e., the probability of hitting the investment threshold is high) and the price decreases. The parameters in this example are such that $\Lambda_S = 0$ so the price is continuous at the investment threshold, i.e., $S(z_1) = S(z_2)$.

The maximum price S_{\max} in Figure 4 partitions the state space into two regimes.

²⁷These sample SDs were not included as moments, but as weights in the SMM estimation.

On the right side of the figure, where $z_t \geq z_{Smax}$, is the *far-from-investment* zone. In this region, investment in new oil is sufficiently unlikely in the short term, and the oil price is decreasing in z_t . On the left side of the figure, where $z_1 < z_t \leq z_{Smax}$, is the *near-investment* zone. In this region, the likelihood of investing in new oil dominates and the economy anticipates an increase in the supply of oil. This implies that as z_t declines the spot price decreases as well, because the probability of an increase in oil stocks increases towards the investment boundary. Figure 5 shows the probability of investing at least one time for different horizons. Since the state variable is continuous inside the *no-investment* region, the probability in the *near-investment* zone is higher than the one in the *far-from-investment* region. Of course, the likelihood of investment is increasing in the horizon.

The fact that the oil price S_t is a non-monotonic function of the state variable z_t is an important feature of our model. Since the inverse function $z(S)$ does not exist, the oil price process is non-Markov in S_t . This is a feature found in the data. Typically, more than one factor is required to match oil futures prices (see, for example, Schwartz (1997)). Note in Figure 2 that two futures curves with the same spot price are not identical. In our model, the “second factor” that is needed in addition to the current spot price is whether the economy is in the *near-investment* or *far-from-investment* region.

We state the equilibrium process for the oil price in terms of S_t and ε_t where ε_t is an indicator that is one if z_t is in the *far-from-investment* region, and two if z_t is in the *near-investment* region.

Proposition 7 *The oil price in equation (71) is governed by the following two-regime stochastic process*

$$\frac{dS_t}{S_t} = \mu_S(S_t, \varepsilon_t)dt + \sigma_{SK}(S_t, \varepsilon_t) dw_{K,t} + \sigma_{SQ}(S_t, \varepsilon_t) dw_{Q,t} + \Lambda_S dI_t \quad (88)$$

$$\begin{aligned} \mu_S(S_t, \varepsilon_t) = & r(S_t, \varepsilon_t) - y(S_t, \varepsilon_t) + \\ & \sigma_{SK}(S_t, \varepsilon_t) \{ \lambda_K(S_t, \varepsilon_t) + \rho_{KQ} \lambda_Q(S_t, \varepsilon_t) \} + \\ & \sigma_{SQ}(S_t, \varepsilon_t) \{ \lambda_Q(S_t, \varepsilon_t) + \rho_{KQ} \lambda_K(S_t, \varepsilon_t) \} \end{aligned} \quad (89)$$

$$\sigma_{SK}(S_t, \varepsilon_t) = \lambda_K(S_t, \varepsilon_t) - \theta_K(S_t, \varepsilon_t) \quad (90)$$

$$\sigma_{SQ}(S_t, \varepsilon_t) = \lambda_Q(S_t, \varepsilon_t) - \theta_Q(S_t, \varepsilon_t) \quad (91)$$

where

$$\varepsilon = \begin{cases} 1 & \text{if } z > z_{Smax} \\ 2 & \text{if } z_1 < z \leq z_{Smax} \end{cases} \quad (92)$$

and where $r(S_t, \varepsilon_t) = r_t$, $\lambda_K(S_t, \varepsilon_t) = \lambda_{K,t}$ and $\lambda_Q(S_t, \varepsilon_t) = \lambda_{Q,t}$ as in Proposition 4, $y(S_t, \varepsilon_t) = y_t$, $\theta_K(S_t, \varepsilon_t) = \theta_{K,t}$ and $\theta_Q(S_t, \varepsilon_t) = \theta_{Q,t}$ as in Proposition 6, and Λ_S is defined in equation (74).

Proof First, we apply Itô's lemma to the definition of the spot price S_t in equation (27). The dynamics of S_t depends on the third order terms J_{KKK} , J_{KKQ} , J_{KQQ} and J_{QQQ} . We differentiate the HJB equation in (24) to simplify the resulting sde for S_t and obtain equations (88) to (91). To show that the dynamics of S_t depends only on $\{S_t, \varepsilon_t\}$, we note that there is a one-to-one mapping between $\{S_t, \varepsilon_t\}$ and z_t . Using the reduction of states variables presented in subsection 2.3.3, we obtain that all the variables in Propositions 4 and 6 are only a function of z_t and thus of $\{S_t, \varepsilon_t\}$. \square

Figure 6 shows a typical path for the state variable z_t (bottom plot) and the oil price S_t (top plot). The horizontal lines below show the optimal investment strategy (z_1, z_2) and the boundary between the two regimes z_{Max} . Whenever z_t hits the investment boundary z_1 , it jumps back to z_2 inside the *no-investment* region. The process for z_t is only bounded by below and exhibits mean reversion. When z_t is far from the investment trigger (z_t is high) the drift of z_t is negative, because the production function $f(k, q)$ uses oil to produce capital (and because of depreciation), i.e., Q decreases while K increases. The simulated oil price is shown in the upper part of the figure. The price is non-negative, bounded at S_{Max} , and mean reverting.

Central to commodity derivative pricing are the conditional moments for the spot-price process. Figure 7 plots the conditional instantaneous return and conditional instantaneous volatility of return as a function of S_t . The second factor ε_t , indicating if z_t is in the *far-from-investment* or *near-investment* region, is 1 above the dashed-line and 2 below this line. From the conditional drift, note that the oil price is mean-reverting however, the rate of mean reversion (negative drift) is much higher in the *near-investment* region. Similarly, the conditional volatility behaves differently across the two regions. The sign of the volatility in the figure measures

the correlation of the oil price with the shocks in numeraire good production (see equation (3)). A positive shock to K_t means a negative change in z_t (less oil relative to the numeraire good). Recall from Figure 4, the decrease in z_t implies an increase in the spot price in the *far-from-investment*, hence a positive correlation. However, in the *near-investment* region the spot price decreases implying a negative correlation. At the endogenously determined maximum price, S_{Max} , the volatility is zero and the drift is negative, which means that the price will decrease almost surely. The volatility of $z_{S_{max}}$ is non-zero, so there is uncertainty to which direction is the state variable moving after being at this point.

In order for the regime shifting behavior of the spot price to be detectable (and economically important), the unconditional distribution for the state variable, z_t needs to place some weight near the boundary of the *near-investment* and *far-from-investment* regions. Figure 8 plots the probability density function (simulated) for the state variable z_t . This variable is bounded from below by z_1 . The distribution has positive skewness. For our calibration, 91.1% of the time the oil price is above the marginal cost (that is $z_1 < z_t < z_2$) and 21.5% of the time the economy is in the *near-investment* region ($z_t < Z_{S_{max}}$).²⁸

4.4.1 Fixed Costs Effect

The fixed cost components β_K and β_Q of the irreversible investment play an important role in the economy. Figure 9 makes a comparison of the probability density function of z_t and S_t for the cases without fixed costs and with large fixed costs (10 times the SMM estimates of β_K and β_Q). The upper left plot shows the PDFs of z_t and the commodity spot price S_t as a function of z_t . Without fixed costs investment is infinitesimal ($z_1 = z_2$) and the state variable stays most of the time near the boundary. Also, the price is always below the marginal cost β_X . Given that the production cost for non-OPEC countries is between \$10 and \$15 per barrel, we can see that fixed costs are crucial to generate observed prices. The effect of the fixed costs is also important for higher moments of commodity prices. Due to the infinitesimal investment, crude oil prices are typically near the maximum price and

²⁸Recall that for this example, we are assuming that the price is continuous, so $S_1 = S_2 = \beta_X$. This implies that S_t is above β_X when $z_1 < z_t < z_2$.

have low volatility.²⁹ This creates a lot of negative skewness in the distribution of the prices (see the upper right plot in this figure), implying that the futures curve is backwardated almost 100% of the times. If we consider fixed costs, investment is lumpy and the state variable jumps from the investment trigger to the optimal oil-capital ratio. This creates an extra source of variability in the economy. In this example, the fixed costs are so high that the oil price can achieve an unrealistic price of almost \$1500 per barrel. The volatility of price is low when the price is equal to the marginal costs (returning point z_2 after investment is made), so the price stays *low* most of the times.³⁰ This implies that high fixed costs could potentially generate positive skewness (see the lower right plot in this figure) and very infrequent large investments.

4.4.2 Investment Policy

The non-monotonic relation between the state variable, z_t and the spot price, S_t , is crucial for the regime shifting behavior of the spot price. The size of the hump in Figure 4 is determined by the optimal investment policy z_1 and z_2 . Alternatively, we can see the investment policy as the investment trigger z_1 and the relative size of the investment, X_t/Q_t . In order for the hump to be large, investment in new oil wells needs to be large; that is the ratio X_t/Q_t . To understand how investment policy is affected by our model parameters, Figure 10 shows the investment strategy under various parameters. The graph on the upper-left corner shows the effect of economies of scale in the strategy. The bigger is the fixed cost component β_K , the bigger are the investment delay (z_1 is decreasing in β_K) and the size of the investment. When the fixed cost component is small, the number of new oil wells is low (in the limiting case, investment is infinitesimal). The graph to the right shows that a higher marginal cost delays investments and but the size of the investment increases. The lower-left graph of figure 10 shows the investment strategy as a function of the oil share η . If the oil share is very low, then investment is postponed indefinitely. As long as oil becomes relevant for the production function, the investment trigger increases, which means that investment is made earlier. The graph on the lower-right corner shows the investment sensitivity to the risk aversion degree of the individuals. The

²⁹As we show later, the volatility of prices is related to the drift of the price function, which in this case is zero at the boundary.

³⁰We discuss the non-monotonic price function in the next subsection.

higher the degree of risk aversion the earlier is the investment undertaken (z_1 is increasing with γ). The intuition for this is that agents care more about smoothing consumption, so they make investment decisions to stabilize the state variable z_t . These decisions are to invest a less amount more frequently.

4.5 Oil Futures prices

Figure 11 shows the futures prices for different spot prices and maturities. As with the process for spot prices in Proposition 7, we can use the $\{S_t, \varepsilon_t\}$ characterization of the state variable z_t with futures prices. The thick futures curves are for spot prices in the *far-from-investment* region ($\varepsilon_t = 1$) while the thin lines are for spot prices in the *near-investment* region ($\varepsilon_t = 2$). The mean-reversion in futures prices is inherited from the bounded equilibrium oil price. When the oil price is low, the state variable is far from the investment trigger. This means that the supply of oil decreases on average, so the expected price in the future is above the current price. In these situations the futures curves are upward-sloping or in contango (for example, see the curve when $S_t = 10$ in figure 11). When the price is at the maximum price the futures curves are downward-sloping, i.e., backwardation (see the curves when $S_t = S_{Max}$). The expected price is below the current price, because of a high probability of an increase in oil supply. Figure 11 also shows that the spot price is not sufficient to characterize the futures curve. For higher prices there are two different futures curves that share the same spot price. One for the case of S_t in *far-from-investment* and one for S_t in the *near-investment* region. The futures curve are steeper when the spot price is in the *near-investment* region. This is a direct implication of a likely sooner investment to create new oil. Our model also generates non-monotonic curves (see the humped curve when $S_t = 20$ and the economy is in the *far-from-investment* region). In these situations, there is an expected shortage of oil in the short-run, but in the medium-run some new oil will likely be created through investment. The case when $S_t = 20$ and the economy is in the *near-investment* region has the opposite situation. Today the price is above the marginal cost, but with a high probability there will be new investments, which drops the expected price in the short-run and price is likely to rise in the medium range.

Another way of understanding the behavior of the futures curves is in terms of

interest rates and convenience yields. The no-arbitrage condition in equation (84), shows that the instantaneous slope of the futures curve is related to the difference of interest rates and convenience yields. Figure 12 shows these two variables in terms of the state variable z_t . Interest rates in our model are fairly flat because of a small marginal productivity of oil in the economy, $\eta = 0.04$ (see the plot in the left). Thus, most of the variation in the instantaneous slope of the futures curve is due to the variation in convenience yields (plot in the right of Figure 12).³¹ When inventories of oil are low (*near-investment* region), the convenience yield is very high, implying that the futures curves presents high degrees of backwardation in the short term. When inventories are high (*far-from-investment* region), the convenience yield is low and mostly dominated by the interest rates, implying an upward sloping futures curve (contango).

Recall from equations (70) and (74) that both asset prices and Oil spot prices may jump at the (predictable) investment in oil. However, as shown in Proposition 5, futures prices are continuous and $\Lambda_F = 0$. This is not surprising since a futures price is a martingale (expectations under the equivalent measure of the future spot price) and perfectly anticipate the spot price jump.

The volatility of the futures contract are shown in figure 13. To compare the futures volatility for different oil spot prices we show the *relative volatility* which we define as $\sigma_F(S_t, \varepsilon_t; T - t) / \sigma_S(S_t, \varepsilon)$. This ratio corresponds to the inverse of the optimal *hedge ratio*, which is the number of futures contracts in a portfolio that minimizes the risk exposure of one unit of oil. This ratio is 1 when $t = T$, because the futures price with zero maturity is the spot price. The thick lines show the relative volatility for oil spot prices in the *far-from-investment* region and the thin lines when the spot is in the *near-investment* zone. In general, the volatilities are much lower for higher maturities, which is a consequence from the mean reverting behavior of risk-adjusted prices (often called the Samuelson Effect). The figure also demonstrates the non-linearity in equilibrium futures prices. First, the volatility curves depend on the spot price. In affine reduce-form models for commodity prices the (log) futures prices is linear on the (log) spot price (see for example Schwartz (1997)). This implies that the hedge ratio is independent from the price level.³²

³¹This is consistent with the evidence found for crude oil and interest rates in Casassus and Collin-Dufresne (2005).

³²In fact, the volatility ratio can also be expressed as $F_S S / F$. In linear models this expression can be a function of maturity, but not of the spot price S .

Second, the curves are non-monotonic in the maturity horizon. For high prices, the expected investment in oil (rise in supply) is reflected in the futures contract and also in the volatility. For short maturities and very high prices the *relative volatility* has an abrupt behavior because the volatility of the spot price is very low (recall that $\sigma_S(S_{Max}, \varepsilon, t) = 0$). A negative volatility in figure 13 implies a negative correlation between the spot price and the futures price. This is something possible for high prices. For example, consider that the price is high and the economy is in the *far-from-investment* region, say, $S_t = 25$ and $\varepsilon_t = 1$ (thick line in the plot). Here, the spot price is negatively correlated with shocks in z_t (see figure 7). In the near future, the price is expected to be in the *near-investment* region and to be positively correlated with shocks in z_t . This implies that the spot and futures price can have negative correlation, which is shown with negative relative volatility values in the figure.

4.6 Regime-Switching Estimation

In this section we estimate a linear approximated version of the commodity pricing model in Proposition 7. This model has two regimes that corresponds to the *near-investment* and *far-from-investment* regions. The model for the price is exponentially affine conditional on any given regime. Despite the fact that we are linearizing the conditional moments with our approximation, the model is non-linear because of its regime switching characteristic. Estimating the linear approximation version of the model has several advantages. First, the estimation is much simpler because we can get an approximation of the likelihood in closed form, while in the “exact” model everything has to be calculated numerically. Second, it is easier to extend the exponentially affine model with regime shifts for derivative pricing and risk-management applications. Finally, structural estimations typically need information about the state variables, which in our case is difficult to observe. By considering the approximated model we can base our estimation solely on observed oil prices.

The main prediction of our model is that there are two different regions in the economy, i.e., the *near-investment* and the *far-from-investment* zones. We consider these two regimes in the approximated model. Figures 4 and 7 shows that the

price behaves differently depending on the active region in the economy. The linear approximation of our structural model is

$$dS_t = \mu_S(S_t, \varepsilon_t)S_t dt + \sigma_S(S_t, \varepsilon_t)S_t dw_{S,t} \quad (93)$$

where

$$\mu_S(S, \varepsilon) = \alpha + \kappa_\varepsilon(\log[S_{Max}] - \log[S]) \quad (94)$$

$$\sigma_S(S, \varepsilon) = \sigma_\varepsilon \sqrt{\log[S_{Max}] - \log[S]} \quad (95)$$

and ε_t is a two-state Markov chain with transition (Poisson) probabilities

$$P_t = \begin{bmatrix} 1 - \lambda_1 dt & \lambda_1 dt \\ \lambda_2 dt & 1 - \lambda_2 dt \end{bmatrix} \quad (96)$$

The process in equations (93)-(95) is exponentially affine conditional on being in a regime, i.e., the process for the logarithm of the price has a linear drift term and volatility. The linearization of these terms is a first order approximation of the “exact” process for the oil price in equations (88) to (90). Equation (96) is the transition matrix for the regime variable ε_t . Here, λ_i can be interpreted as the intensity of a jump process for moving out of state $\varepsilon_t = i$. A second, less important approximation is that these λ 's are constant, something that is not true in the exact model since they depend on the price S_t (or in the state variable z_t in a similar way than the probability of investment presented in figure 5). We set $\varepsilon_t = 1$ in the *far-from-investment* region and $\varepsilon_t = 2$ in the *near-investment* region.

Data Description and Estimation Method Our data set consists of weekly Brent crude oil prices between Apr-1983 and Apr-2005 deflated by the US Consumer Price Index. The average price is 16.29 dollars per barrel in 1983 prices (or 31.53 dollars per barrel in 2005 prices). The annualized standard deviation of weekly returns is 38%. The skewness in crude oil prices for this period is 1.09 and the excess kurtosis is 0.42.

The parameter space for the approximated model in equations (93)-(95) is given by $\Theta = \{\alpha, \kappa_1, \kappa_2, \sigma_1, \sigma_2, S_{Max}, \lambda_1, \lambda_2\}$. We use the maximum likelihood estimator

for regime-switching models proposed by Hamilton (1989). We do a quasi-maximum likelihood estimation by considering only the first two moment of the distribution. This should not have a significant impact on the estimates because we are working with weekly data. The Hamilton’s estimators accounts for the non-linearities due to the regime-shift characteristic of our model. A by-product of the estimation technique are the smoothed inferences for each regime. We follow Kim’s (1993) algorithm, which is a backward iterative process that starts from the smoothed probability of the last observation. The smoothed probabilities are important because they give information about the true regime that was active any given day.

Results The parameter estimates and standard errors of our model are given in Table 3. In general, most parameters are significant implying that there are clearly two regimes in the data for the period studied. The parameters vary across regimes implying that these regimes are significantly different. The economy stays on average one year in the first regime, $\lambda_1 = 1.023$, before switching to the second regime. Moreover, the first regime is the most frequent one, since the economy stays approximately 79.5% of the time in it ($\lambda_2/(\lambda_1 + \lambda_2) = 0.795$). The economy stays in the second regime on average a couple of months before jumping back to regime 1 ($\lambda_2 = 3.967$). The parameter $\alpha = -0.184$ is negative and significant implying that the process for the price has an upper bound at S_{Max} . Also, the estimate for S_{Max} is a reasonable upper bound given the historical path of crude oil prices ($Ln[S_{Max}] = 4.469$). Under the most frequent regime, the crude oil price follows a strong mean-reverting process ($\kappa_1 = 0.319$), i.e., the drift is positive for low spot prices and negative for high prices. The infrequent regime is different, since the mean reversion parameter κ_2 is insignificant. Also, the second regime is characterized to be more volatile than the first regime ($\sigma_2 > \sigma_1$).

Figure 14 shows the crude oil price and the inferred probability of being in the *near-investment* state (regime 2). We can see that most of the time this probability is low (thin line), implying that the economy stays mainly in the *far-from-investment* regime. Also, when the probability is high, most of the times the price decreases very sharply, which is a characteristic of the *near-investment* regime. In the *far-from-investment* periods, the price seems to have a mean reverting behavior. Many of these results are reflected also in the estimates of table 3. Figure 14 shows that the *near-investment* regime is generally for high prices (like in figure 4), but sometimes

it can be for low spot prices as well. This implies that in the exact model the fixed cost components of the irreversible investment are high enough such that the average price is above the marginal price. This allows to generate both, high and low prices in the *near-investment* state.

The smoothed probabilities from the maximum likelihood estimation are also important to validate the predictions about the futures prices. For this we do a simple exercise. First, we use the smoothed probabilities to detect the periods of time where the economy was under one regime or the other. Second, we group the futures curve in different regimes according to the backed out dates.³³ Third, we sort the curves for both regimes by the price of the shortest maturity contract (typically the one-month futures contract with price F_1) and group them according to this price.³⁴ Finally, we compare the behavior of the futures curves under both regimes with the predictions from our model. We follow a very simple approach for this comparison by calculating the sample mean of the shortest maturity contract ($\overline{F_1}$) and the average short-term curvature of the futures curve ($\overline{F_1} - 2\overline{F_6} + \overline{F_{12}}$).³⁵

Table 4 shows the results. There are three important results that validate our model. First, for each regime the column “Nobs” shows the number of observations in every bin (range of F_1 prices). Just by comparing these columns for both regimes we see that the median in the *near-investment* regime is higher than the one in the *far-from-investment* regime. This confirms that on average the prices are higher in the *near-investment* regime. Second, we can see that in both regimes the curvature is positive for high prices and negative for low prices, implying mean reversion under the equivalent martingale measure. This is one of the main predictions for the futures prices in our model. Finally, we see that for high spot prices (i.e., the first three bins {“30-”, “25-30”, “20-25”}), the curvature of the futures curve in the short-term is higher in the *near-investment* investment region.³⁶ This occurs in our model because the convenience yield is higher in the *near-investment* region, which implies higher degrees of backwardation.

³³We have the futures curve for (Nymex) crude oil prices from Jan-90 to Aug-03.

³⁴We use the notation F_i for the futures price of a contract with the nearest maturity to i months.

³⁵The measure of curvature that we choose is the price of a portfolio of futures contracts, where we have a long position in the one-month and one-year maturity contracts and a short position in two six-month contracts. It is easy to see that this can be a measure of the second derivative of the curve for short maturities ($\omega = F_1 - 2F_6 + F_{12}$).

³⁶The results are similar when we use contracts with other maturities for the measure of curvature.

4.7 Commodity Risk Premium

As shown in Figure 15 the model predicts that the commodity risk premium is time-varying: it is positive in the *far-from-investment* regime but negative in the *near-investment* regime. To understand why, consider a positive productivity shock which results in increased oil consumption. This has two opposing effects on oil prices. On the one hand, it reduces the available supply which tends to raise prices. On the other hand, it brings the economy closer to the investment trigger, which tends to lower prices via expected future higher supply. The latter effect dominates in the *near-investment* region, whereas the former dominates in the *far-from-investment* regime. This explains the switch in the sign of the covariance of oil price with productivity shocks and therefore the time variation in commodity risk-premia.

To investigate empirically if we can find some support for the time variation in risk-premia predicted by the model we run simple time series regression of returns on investment in oil on the market return (proxied by the S&P 500). We allow for the beta in the regression to be time-varying and use as conditioning variables the smoothed inferred probability of being in the *near-investment* regime as well as the slope of the futures curve. More specifically, we run the following regression:

$$r_{j,t+1}^e = a_t + b_t r_{M,t+1}^e + \epsilon_{t+1}$$

where $a_t = a_0 + a_1 z_t$, $b_t = b_0 + b_1 z_t$ and z_t is the conditioning variable (i.e., *near-investment* regime probability or slope of futures curve). For the return on the investment in oil we use two proxies. One is the return on a spot investment in oil (analogous to a buy and hold transaction in a stock). The other is the return to a fully collateralized long futures position in oil (probably the more common approach for investors to take positions in energy markets, e.g., Gorton and Rouwenhorst (2005), Erb and Harvey (2005)). Figure 16 shows the cumulative returns associated with both strategies. The return to the futures strategy clearly dominates the spot return strategy reflecting the existence of a convenience yield, the non-monetary dividend which accrues to the holder of the spot but not of the futures contract.

The result of the regression are presented in table 5. They provide support for the fact that risk-premia on oil are time varying and related to the investment regime. In

particular, consistent with the prediction of our theoretical model, the risk-premium is significantly negatively related to the smoothed probability of being in the *near-investment* regime. Note that the coefficient on the market return interacted with the *near-investment* regime probability is very significant for both commodity spot and futures return series (cf. row 2 of panels A and B in table 1). Conditioning on the regime probability increases the R^2 from 1% to around 5%. On the other hand, the slope of the futures curve predicts significantly only the return on the spot price and not on the collateralized futures position. This is consistent with the fact that, to a first order, the slope at the short end of the futures curve moves with the convenience yield. When the slope increases, the convenience yield decreases and the return on the spot position, which is effectively an ex-dividend return, increases (as can be seen from the very significant positive coefficient on slope in row 3 of panel A). On the other hand, the return to the collateralized futures position is not significantly affected by a change in the convenience yield (the coefficient in row 3 of panel B on slope is not significant). This suggests that commodity risk premium, while time-varying, is not driven by the slope of the futures curve.

We note that most other coefficients are not statistically significant, which indicates that there is little evidence for an unconditional risk-premium for investing in oil (at least in our data set). The biggest gain in predictability appears to come from conditioning on the investment regime.³⁷

5 Conclusion

We develop an equilibrium model for spot and futures oil prices. Our model considers the commodity as an input for a production technology in an explicit way. This feature endogenizes one of the main assumptions in standard competitive models of storage, i.e., the demand function. Our model generates positive convenience yields and long period of backwardation in futures curves without the necessity of running out of oil, like in the standard “stock-out” literature. Convenience yields arise

³⁷Our predictability results are robust to using discrete or log returns for market and/or commodity returns. They are also unchanged if we use excess returns instead of gross returns. One caveat applies with respect to our use of the regime probability: The coefficients of the model used to infer the smoothed probability of the regime have been estimated with the whole data, and therefore there is somewhat of a forward looking bias in that statistic. This is typical of this type of study, e.g., Lettau and Ludvigson (2001).

endogenously due to the productive value of the oil, which is consistent with the predictions of the “Theory of Storage”. This convenience yield is high when the stocks of commodity are low, and viceversa. By modeling explicitly risk-averse agents, we can investigate risk-premia associated with holding of stocks of commodities versus futures contracts.

Equilibrium spot price behavior is endogenously determined as the shadow value of oil. Our model makes predictions about the dynamics of oil spot prices and futures curves. The equilibrium price follows an heteroscedastic mean-reverting process. The spot price is non-Markov, because there are two regimes in our economy that depend on the distance to the investment region. For reasonable parameters, the futures curves are most of the time backwardated. Also, the two regimes imply that two futures curve with similar spot prices can have very different degrees of backwardation. Further, the model predicts time varying risk premium on oil: positive in the *far-from-investment* regime and negative in the *near-investment* regime.

We estimate the model using the Simulated Method of Moments for futures prices and macroeconomic data. We find that the model captures many of the stylized facts of our data set. In particular, our model can reproduce the mean and volatilities of futures prices for maturities up to 36 months and also the average consumption-output ratios, consumption of oil-output ratio and real interest rates. We estimate a linear approximation of our model with crude oil prices from 1983 to 2004 and find evidence for regime switching behavior consistent with the predictions of the model. Further, we find that, consistent with the model predictions, excess returns on oil are predictable and related to the inferred probability of investment regime.

Appendix

A Numerical Techniques

In this appendix we delineate the numerical algorithm used to solve the Hamilton-Jacobi-Bellman equation in (24) with boundary conditions represented by equations (28) to (31).

The first step is to use the homogeneity of the solution to reduce the state space (see Subsection 2.3.3). After this is done, the solution of the problem is represented by a nonlinear second order ODE in the state variable $z_t = \log(Q_t/K_t)$. The boundary conditions are also expressed in terms of z_t . Now, we need to determine the value function $j(z)$ in equation (33). The nonlinear HJB equation for $j(z)$ depends on (i) the optimal control c_t^* , and on (ii) the optimal investment strategy $\{z_1, z_2\}$ determined by the boundary conditions. Unfortunately, the optimal control itself depends on the value function $j(z)$. This implies that $j(z)$, c_t^* , z_1 and z_2 need to be simultaneously determined.

We use an iterative method to solve for $j(z)$. The main idea is to build a conditionally linear ODE for $j(z)$ so it is possible to apply a finite-difference scheme. The selection of the initial guess is extremely important for the convergence of the iteration. We assume that $j^0(z) = 1$ which corresponds to the solution when the oil is not relevant for the production technology ($\eta = 0$). In this case we also know that it is never optimal to invest $z_1^0 \rightarrow \infty$.

For every iteration m (for $m = 0 \dots \infty$) we do the following steps:

- Determine the optimal consumption c^{*m} as a function of $j^m(z)$ using equation (42).
- We recognize that the ODE for $j^{m+1}(z)$ determines the value function when it is optimal not to invest in new stocks of commodity. We name this function as $j_{noinv}^{m+1}(z)$. We calculate the coefficients of the ODE for $j_{noinv}^{m+1}(z)$. It is important to notice that this ODE is linear conditional on c^{*m} .
- Determine the optimal commodity/capital ratio z_2^{m+1} using the super contact condition in equation (47). Conditional that it is optimal to invest in new commodity stocks, the returning point is always z_2^{m+1} independent of what was the value of z_t before investment was made. Using this argument we define the extended value matching condition as

$$j_{inv}^{m+1}(z) = j^m(z_2^{m+1}) \left(\frac{1 - \beta_K + e^z(\beta_X - \beta_Q)}{1 + e^{(z_2^{m+1})} \beta_X} \right)^{1-\gamma}. \quad (A1)$$

This equation represents the value function when the representative agent is forced to invest.

- Use a finite-difference scheme to solve for the value function $j_{noinv}^{m+1}(z)$. The finite difference discretization defines a tridiagonal matrix that needs to be inverted to determine the value of $j_{noinv}^{m+1}(z)$. Instead of doing this, we eliminate the upper diagonal of this matrix. At this point the value of $j_{noinv}^{m+1}(z)$ depends only on the value of $j_{noinv}^{m+1}(z - \Delta z)$. We choose a z_{min} negative enough to ensure that at that level it is optimal to invest, and then we solve the value function for higher z_t . At every point we choose the maximum of the value from investing ($j_{inv}^{m+1}(z)$) and the value of no investing which comes from the finite-difference scheme. This maximum determines the value of $j^{m+1}(z)$. The optimal trigger z_1^{m+1} is endogenously determined when the representative agent is indifferent between investing and postponing the investment. The algorithm described above is a more efficient way than solving independently for $j_{inv}^{m+1}(z)$ and $j_{noinv}^{m+1}(z)$ and then choosing $j^{m+1}(z) = \max(j_{inv}^{m+1}(z), j_{noinv}^{m+1}(z))$.
- Check for the convergence condition. If it not satisfied we start a new iteration with the updated value of $j^{m+1}(z)$.

Once $j(z)$ has converged it is straight forward to calculate spot commodity prices from equation (71). For the futures prices we use an implicit finite-difference technique. This is simpler than the solution for $j(z)$ since the coefficients of the PDE and boundary conditions and boundaries $\{z_1, z_2\}$ are known at the beginning of the scheme.

References

- Abel, Andrew B., and Janice C. Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review* 84, 1369–1384.
- , 1996, Optimal investment with costly reversibility, *The Review of Economic Studies* 63, 581–593.
- , 1997, An exact solution for the investment and value of a firm facing uncertainty, adjustment costs, and irreversibility, *Journal of Economic Dynamics and Control* 21, 831–852.
- Arrow, Kenneth J., and Mordecai Kurz, 1970, *Public Investment, the Rate of Return and Optimal Fiscal Policy* (Johns Hopkins Press: Baltimore).
- Baldursson, Fridrik M., and Ioannis Karatzas, 1997, Irreversible investment and industry equilibrium, *Finance and Stochastics* 1, 69–90.
- Bertola, Giuseppe, and Ricardo J. Caballero, 1994, Irreversibility and aggregate investment, *The Review of Economic Studies* 61, 223–246.
- Bessembinder, Hendrick, Jay F. Coughenour, Paul J. Seguin, and Margaret M. Smoller, 1995, Mean-reversion in equilibrium asset prices: Evidence from the futures term structure, *Journal of Finance* v50 n1, 361–375.
- Bobenrieth, Eugenio S. A., Juan R. A. Bobenrieth, and Brian D. Wright, 2002, A commodity price process with a unique continuous invariant distribution having infinity mean, *Econometrica* 70(3), 1213–1219.
- Brennan, Michael J., 1958, The supply of storage, *American Economic Review* 48, 50–72.
- , 1991, The price of convenience and the valuation of commodity contingent claims, in D. Lund and B. Oksendal, Eds.: *Stochastic Models and Option Values*, North Holland.
- , and Eduardo S. Schwartz, 1985, Evaluating natural resource investments, *Journal of Business* v58 n2, 135–57.
- Caballero, Ricardo J., and Eduardo M. R. A. Engel, 1999, Explaining investment dynamics in u.s. manufacturing: A generalized (s,s) approach, *Econometrica* 67(4), 783–826.
- Carlson, Murray, Zeigham Khokher, and Sheridan Titman, 2002, An equilibrium analysis of exhaustable resources investments, University of Texas Working Paper.

- Casassus, Jaime, and Pierre Collin-Dufresne, 2005, Stochastic convenience yield implied from commodity futures and interest rates, *Journal of Finance* 60, 2283–2331.
- Chambers, Marcus, and Roy Bailey, 1996, A theory of commodity price fluctuations, *Journal of Political Economy* 104(5), 924–957.
- Cochrane, John H., 2001, *Asset Pricing* (Princeton: University Press).
- Cox, John C., Jonathan E. Ingersoll Jr., and Steve A. Ross, 1985, An intertemporal general equilibrium model of asset prices, *Econometrica* 53, 363–384.
- Deaton, Angus, and Guy Laroque, 1992, Competitive storage and commodity price dynamics, *Journal of Political Economy* 104, 896–923.
- Dixit, Avinash, 1991, A simplified treatment of the theory of optimal regulation of brownian motion, *Journal of Economic Dynamics and Control* v15 n4, 657–73.
- , 1993, *The Art of Smooth Pasting* (Vol. 55 in Fundamentals of Pure and Applied Economics, Harwood Academic Publishers).
- Dixit, Avinash K., and Robert S. Pindyck, 1994, *Investment under Uncertainty* (Princeton University Press: Princeton).
- Duffie, Darrell, 1996, *Dynamic Asset Pricing Theory* (Princeton: University Press).
- , and Kenneth Singleton, 1993, Simulated moments estimation of markov models of asset prices, *Econometrica* 61, 929–952.
- Dumas, Bernard, 1991, Super contact and related optimality conditions, *Journal of Economic Dynamics and Control* v15 n4, 675–85.
- , 1992, Dynamic equilibrium and the real exchange rate in a spatially separated world, *Review of Financial Studies* 5, 153–180.
- Erb, Claude B., and Campbell R. Harvey, 2005, The tactical and strategic value of commodity futures, National Bureau of Economic Research, Working Paper.
- Fama, Eugene F., and Kenneth R. French, 1987, Commodity futures prices: some evidence on forecast power, premiums, and the theory of storage, *Journal of Business* v60 n1, 55–73.
- , 1988, Business cycles and the behavior of metals prices, *Journal of Finance* v43, 1075–1093.
- Finn, Mary G., 1995, Variance properties of solow’s productivity residual and their cyclical implications, *Journal of Economic Dynamics and Control* v19, 1249–1281.

- , 2000, Perfect competition and the effect of energy price increases on economic activity, *Journal of Money, Credit and Banking* v32n3, 400–416.
- Gibson, Rajna, and Eduardo S. Schwartz, 1990, Stochastic convenience yield and the pricing of oil contingent claims, *Journal of Finance* v45n3 (July), 959–76.
- Gorton, Gary, and K. Geert Rouwenhorst, 2005, Facts and fantasies about commodity futures, Yale ICF, Working Paper.
- Grossman, Sanford J., and Guy Laroque, 1990, Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* 58, 2551.
- Guesnerie, Roger, 1975, Pareto optimality in non-convex economies, *Econometrica* 43, 1–29.
- Gustafson, Robert L., 1958, Carryover levels for grains: A method for determining amounts that are optimal under specified conditions., Technical Bulletin No. 1178, U.S. Department of Agriculture, Washington, DC.
- Hamilton, James D., 1989, A new approach to the economic analysis of nonstationary time series and the business cycle, *Econometrica* 57, 357384.
- Harrison, J. Michael, 1990, *Brownian Motion and Stochastic Flow Systems* (Krieger Publishing Company: Florida).
- Harrison, Michael, Thomas Sellke, and Allison Taylor, 1983, Impulse control of brownian motion, *Mathematics of Operations Research* v8 n3, 454–466.
- Hartley, P. M., and L. C. G. Rogers, 2003, Two-sector stochastic growth models, University of Cambridge Preprint Series.
- Jeanblanc-Picque, Monique, and Albert N. Shiryaev, 1995, Optimization of the flow of dividends, *Russian Math. Surveys* v50, 257–277.
- Kaldor, Nicholas, 1939, Speculation and economic stability, *The Review of Economic Studies* 7, 1–27.
- Karatzas, Ioannis, and Steven E. Shreve, 1998, *Methods of Mathematical Finance* (Applications of Mathematics 39, Springer-Verlag: New York).
- Kim, Chang-Jin, 1993, Dynamic linear models with markov-switching, *Journal of Econometrics* 60, 1–22.
- Kogan, Leonid, 2001, An equilibrium model of irreversible investment, *Journal of Financial Economics* v62 n2 November, 201–45.

- , 2004, Asset prices and real investment, *Journal of Financial Economics* 73, 411–431.
- , Dmitry Livdan, and Amir Yaron, 2005, Futures prices in a production economy with investment constraints, National Bureau of Economic Research, Working Paper.
- Lettau, Martin, and Sydney C. Ludvigson, 2001, Consumption, aggregate wealth, and expected stock returns, *Journal of Finance* 56, 815–849.
- Litzenberger, Robert H., and Nir Rabinowitz, 1995, Backwardation in oil futures markets: theory and empirical evidence, *Journal of Finance* 50, 1517–1545.
- Mamaysky, Harry, 2001, The term structure of interest rates and the price of durable goods, Working Paper Yale University.
- McDonald, Robert, and Daniel Siegel, 1986, The value of waiting to invest, *Quarterly Journal of Economics* v101 n4, 707–27.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–888.
- , 1975, An asymptotic theory of growth under uncertainty, *The Review of Economic Studies* 42, 375–393.
- Newbery, David M., and Joseph E. Stiglitz, 1981, *The Theory of Commodity Price Stabilization, A Study of the Economics of Risk* (Clarendon Press: Oxford).
- Paddock, James L., Daniel R. Siegel, and James L. Smith, 1988, Option valuation of claims on real assets: The case of offshore petroleum leases, *Quarterly Journal of Economics* 103, 479–508.
- Pindyck, Robert S., 1988, Irreversible investment, capacity choice, and the value of the firm, *American Economic Review* 78, 969–985.
- Richard, Scott F., and M. Sundaresan, 1981, A continuous time equilibrium model of forward prices and futures prices in a multigood economy, *Journal of Financial Economics* 9, 347–371.
- Ross, Stephen A., 1997, Hedging long run commitments: Exercises in incomplete market pricing, *Economic Notes by Banca Monte* 26, 99–132.
- Routledge, Bryan, Duane Seppi, and Chester Spatt, 2000, Equilibrium forward curves for commodities, *Journal of Finance* LV n 3, 1297–1338.
- Scarf, Herbert, 1960, The optimality of (s; s) policies in the dynamic inventory problem, In *Mathematical Methods in the Social Sciences K. Arrow, S. Karlin and P. Suppes (ed.)*, Stanford, CA: Stanford University Press pp. 196–202.

- Scheinkman, Jose A., and Jack Schechtman, 1983, A simple competitive model with production and storage, *The Review of Economic Studies* 50(3), 427–441.
- Schwartz, Eduardo S., 1997, The stochastic behavior of commodity prices: Implications for valuation and hedging, *Journal of Finance* v52n3 (July), 923–73.
- , and James E. Smith, 2000, Short-term variations and long-term dynamics in commodity prices, *Management Science* v47n2 (July), 893–911.
- Sundaresan, M., 1984, Consumption and equilibrium interest rates in stochastic production economies, *Journal of Finance* 39, 77–92.
- Telser, Lester G., 1958, Futures trading and the storage of cotton and wheat, *Journal of Political Economy* 66, 233–55.
- Uppal, Raman, 1993, A general equilibrium model of international portfolio choice, *Journal of Finance* v48 n2, 529–53.
- Wei, Chao, 2003, Energy, the stock market, and the putty-clay investment model, *American Economic Review* v93, 313–323.
- Williams, Jeffrey C., and Brian D. Wright, 1991, *Storage and Commodity Markets* (Cambridge University Press: England).
- Working, Holbrook, 1948, Theory of the inverse carrying charge in futures markets, *Journal of Farm Economics* 30, 128.
- , 1949, The theory of the price of storage, *American Economic Review* 39, 1254–62.
- Wright, Brian D., and Jeffrey C. Williams, 1982, The economic role of commodity storage, *Economic Journal* 92(367), 596–614.

Production technologies		
Productivity of capital K (*)	α	0.128
Importance of oil	η	0.04
Demand rate for oil (*)	\bar{i}	0.104
Volatility of capital (*)	σ_K	0.364
Volatility of oil stocks	σ_Q	0.013
Correlation of capital and oil shocks	ρ_{KQ}	0
Depreciation of oil	δ	0.20
Irreversible investment		
Fixed cost (K component) (*)	β_K	0.012
Fixed cost (Q component)	β_Q	0.150
Marginal cost of oil	β_x	12.5
Agents preferences		
Patience	ρ	0.05
Risk aversion (*)	γ	0.52

Table 1: Parameters from the calibration exercise.

Moment conditions	Historical data		Model	
	Sample average	Sample SD	Uncond. Mean	Uncond. SD
Futures prices - 01	23.53	7.31	21.74	6.19
Futures prices - 03	23.11	6.93	21.69	5.94
Futures prices - 06	22.52	6.36	21.64	5.56
Futures prices - 09	22.07	5.90	21.59	5.21
Futures prices - 12	21.73	5.54	21.56	4.87
Futures prices - 18	21.29	5.01	21.50	4.24
Futures prices - 24	21.07	4.63	21.46	3.69
Futures prices - 30	21.01	4.40	21.43	3.22
Futures prices - 36	20.92	4.22	21.41	2.86
Volatility - 01	0.251	0.028	0.203	0.153
Volatility - 03	0.222	0.036	0.202	0.136
Volatility - 06	0.183	0.043	0.195	0.115
Volatility - 09	0.156	0.047	0.187	0.099
Volatility - 12	0.138	0.049	0.177	0.088
Volatility - 18	0.117	0.049	0.156	0.073
Volatility - 24	0.102	0.039	0.136	0.065
Volatility - 30	0.095	0.031	0.118	0.061
Volatility - 36	0.098	0.039	0.104	0.057
Consumption/GDP	0.617	0.008	0.616	0.019
(Petroleum consumption)/GDP	0.016	0.004	0.017	0.007
Real interest rates	0.015	0.015	0.015	0.002

Table 2: Historical and implied moments by the model using parameters in Table 1.

<i>far-from-investment state</i>			<i>near-investment state</i>		
Parameter	Estimate	t-ratio	Parameter	Estimate	t-ratio
λ_1	1.023	3.1	λ_2	3.967	2.5
$1/\lambda_1$	0.978		$1/\lambda_2$	0.252	
$\lambda_2/(\lambda_1 + \lambda_2)$	79.5%		$\lambda_1/(\lambda_1 + \lambda_2)$	20.5%	
κ_1	0.319	2.2	κ_2	0.253	0.5
σ_1	0.287	12.4	σ_2	-0.790	7.2

Common parameters		
Parameter	Estimate	t-ratio
α	-0.184	-1.7
$Ln[S_{Max}]$	4.469	3.9

Table 3: Quasi-maximum likelihood estimates for the regime-switching model for weekly deflated Brent crude oil prices between Jan-1982 and Aug-2003.

F_1 oil prices	<i>far-from-investment state</i>			<i>near-investment state</i>		
(\$/barrel)	Nobs	$\overline{F_1}$	$\overline{F_1} - 2\overline{F_6} + \overline{F_{12}}$	Nobs	$\overline{F_1}$	$\overline{F_1} - 2\overline{F_6} + \overline{F_{12}}$
30-	41	32.4	114.9	32	33.1	181.8
25-30	93	27.3	3.0	35	27.9	92.7
20-25	189	21.8	31.9	17	21.7	41.5
15-20	237	18.1	-7.2	13	18.2	-34.1
10-15	54	13.5	-28.7	2	12.4	-182.0

Table 4: Sample mean of the shortest maturity contract ($\overline{F_1}$) and average short-term curvature of the futures curve ($\overline{F_1} - 2\overline{F_6} + \overline{F_{12}}$) under different regimes and for different groups of crude oil prices between Jan-1990 and Aug-2003. The active regime is inferred by the estimation of the regime-switching model.

A. Coefficient estimates for crude oil returns							
Row	Constant	$r_{M,t+1}^e$	Additional Regressors				R^2
			p_t	$p_t \cdot r_{M,t+1}^e$	sl_t	$sl_t \cdot r_{M,t+1}^e$	
1	-0.07 (0.11)	-0.20 (1.47)					0.01
2	0.17 (0.23)	0.05 (0.30)	-0.88 (0.45)	-1.23 (2.95)			0.05
3	0.48 (0.71)	-0.19 (1.37)			78.86 (2.28)	-2.74 (0.38)	0.03

B. Coefficient estimates for collateralized futures returns							
Row	Constant	$r_{M,t+1}^e$	Additional Regressors				R^2
			p_t	$p_t \cdot r_{M,t+1}^e$	sl_t	$sl_t \cdot r_{M,t+1}^e$	
1	1.05 (1.67)	-0.17 (1.28)					0.01
2	1.12 (1.54)	0.06 (0.41)	-0.07 (0.03)	-1.17 (2.85)			0.04
3	0.87 (1.28)	-0.19 (1.38)			-22.32 (0.65)	-2.72 (0.37)	0.01

Table 5: The table presents the a and b estimates from the following time-series regressions: $r_{j,t+1}^e = a_t + b_t r_{M,t+1}^e + \epsilon_t$ where $a_t = a_0 + a_1 z_t$, $b_t = b_0 + b_1 z_t$ and z_t is the scaling (conditioning) variable. The term $r_{j,t+1}^e$ in the regressions is the (log) real excess return of crude oil spot prices (*Panel A*) or the (log) real excess return of the collateralized futures strategy (*Panel B*). $r_{M,t+1}^e$ is the real excess return of the value-weighted CRSP index. The scaling variables z_t are the smoothed inferred probability of being in the *near-investment* regime, p_t , and the slope on the futures curve, sl_t . The t -statistics is presented in parentheses below each coefficient estimate.

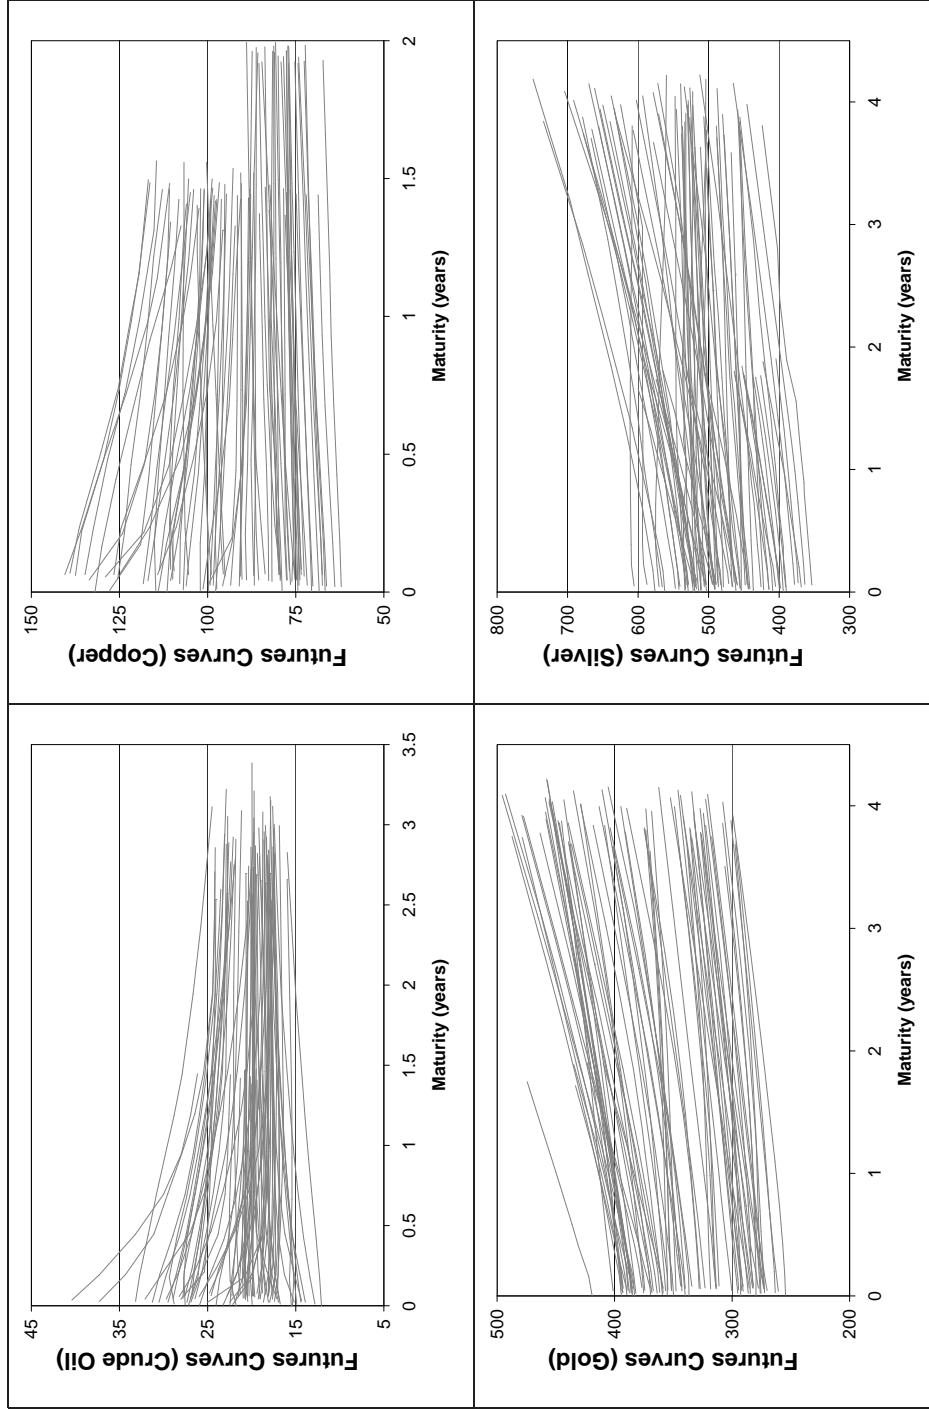


Figure 1: Monthly term structures of futures prices on crude oil, copper, gold and silver between 1/2/1990 to 8/25/2003. Oil prices are in dollars per barrel, copper prices are in dollars per troy ounce and silver prices are in cents per troy ounce.




Figure 2: Sample futures curves from 1990-2003 for Brent crude (deflated by U.S. Consumer Price Index).




Figure 3: Mean and volatility of futures prices for different maturities. The markers show the moments from historical annual crude oil prices from 1983 to 2004. The lines show the moments implied by the model using the parameters in Table 1.




Figure 4: Oil price S_t as a function of the logarithm of the oil wells-capital ratio, z_t . The vertical dashed-line is at z_{Smax} and separates two regions. The thin line shows the oil price in the *near-investment* region ($z_1 < z_t \leq z_{Smax}$) and the thick line is the oil price in the *far-from-investment* region ($z_t \geq z_{Smax}$). We use the parameters in Table 1. In particular, the fixed cost components of the investment are $\beta_K = 0.012$ and $\beta_Q = 0.15$, and the marginal cost of oil is $\beta_x = 12.5$. The equilibrium critical ratios are $z_1 = -8.91$, $z_{Smax} = -8.14$ and $z_2 = -6.29$.




Figure 5: Probability of investing during an interval of time $T - t$ as a function of the logarithm of the oil wells-capital ratio. The thin lines show the probability in the *near-investment* region and the thick line in the *far-from-investment* region.




Figure 6: Simulations for the logarithm of the oil wells-capital ratio z_t (below) and the oil price S_t (above) over time. The thin lines show these variables in the *near-investment* region and the thick lines show them in the *far-from-investment* region. We use the parameters in Table 1. For the path below, the critical ratios are $z_1 = -8.91$, $z_{Smax} = -8.14$ and $z_2 = -6.29$, while for path above the (equilibrium) maximum price is $S_{Max} = 29.58$ and the marginal cost of oil is $\beta_x = 12.5$.

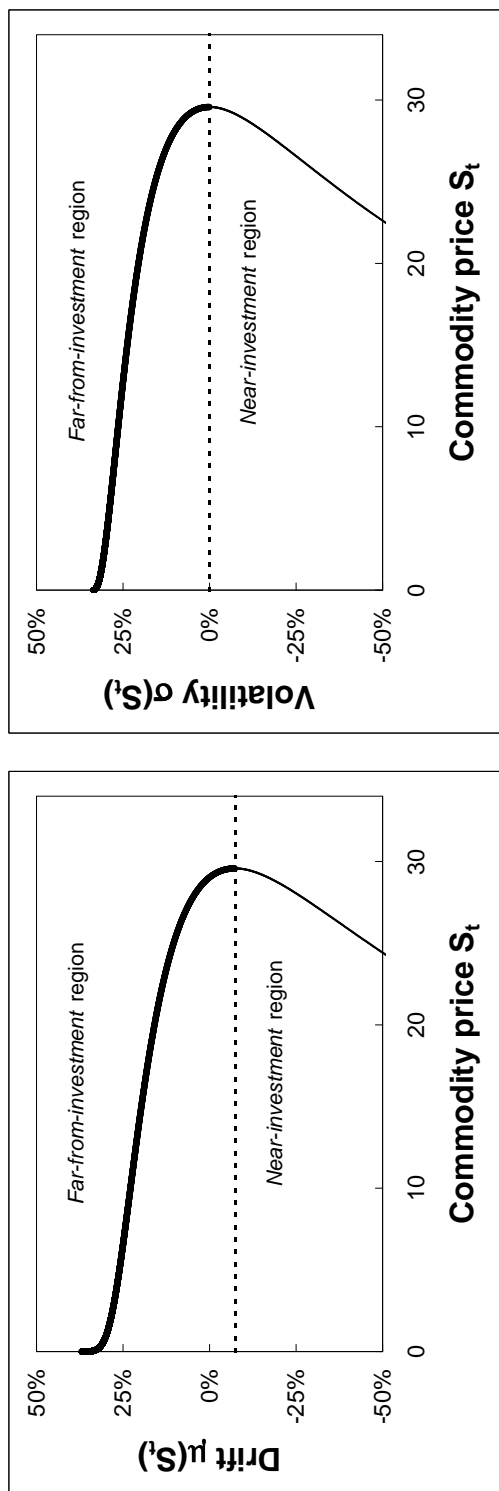


Figure 7: Return and instantaneous volatility of returns in oil price S_t . The horizontal dashed-line separates the two regimes. The thin lines below the dashed-lines show the variables under the *near-investment* regime and the thick lines under the *far-from-investment* regime. We use the parameters in Table 1. In particular, the fixed cost components of the investment are $\beta_K = 0.012$ and $\beta_Q = 0.15$, and the marginal cost of oil is $\beta_x = 12.5$. The endogenous upper bound for the price is $S_{Max} = 29.58$.




Figure 8: Simulated probability density function for the state variable z_t and the commodity price S_t using the parameters in Table 1.

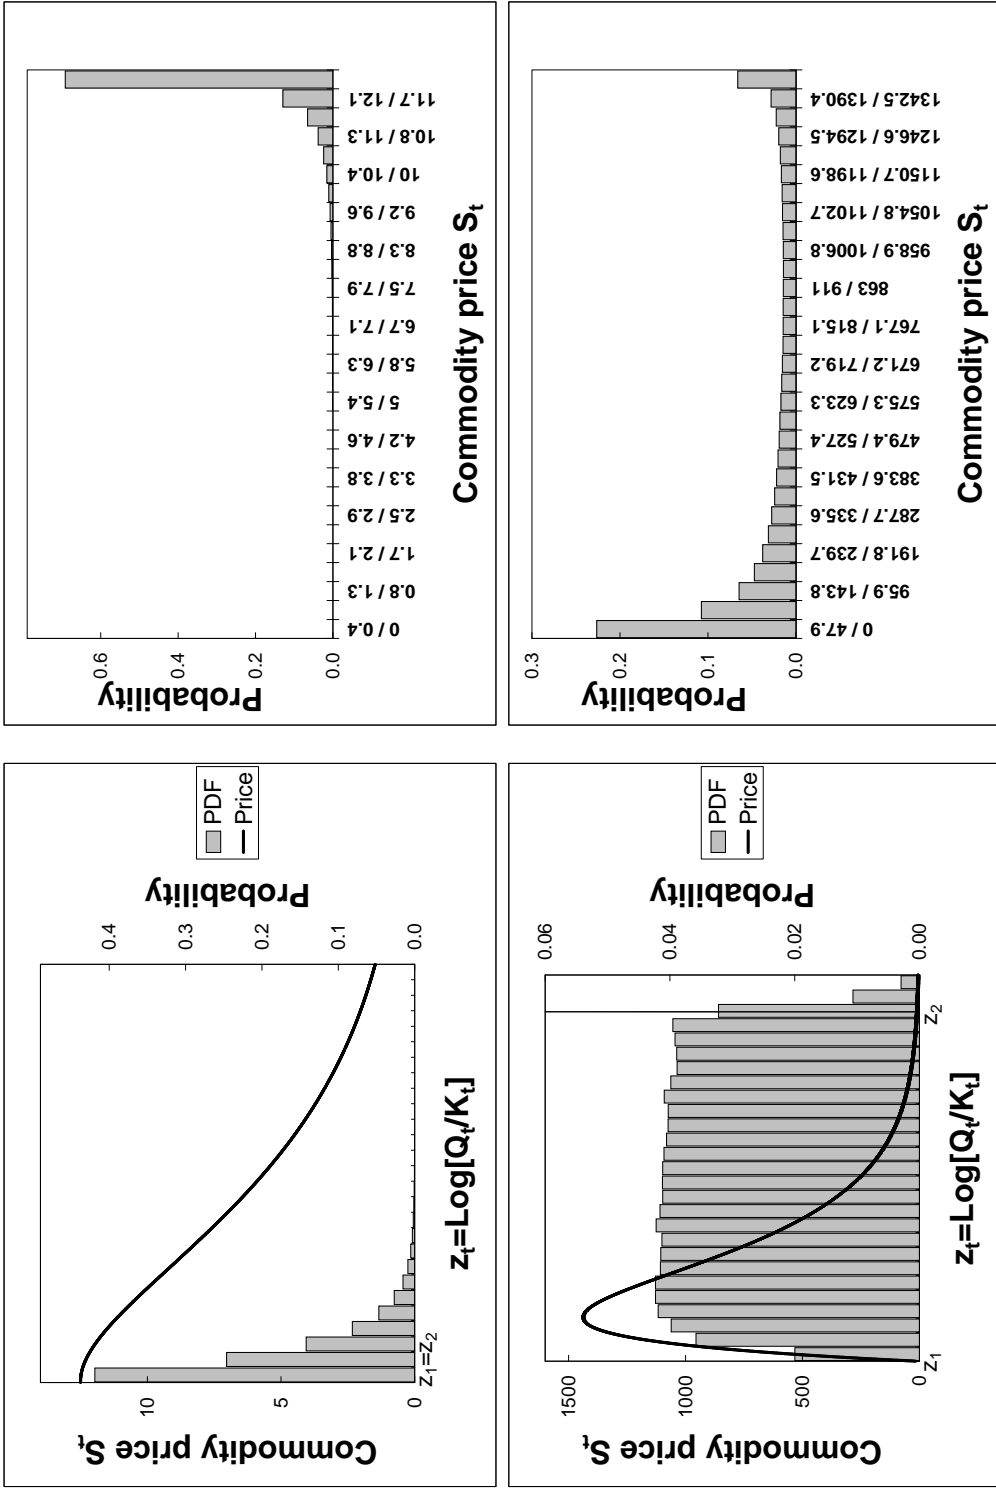


Figure 9: Effect of fixed costs component of the investment in the simulated probability density function of the state variable z_t and of the commodity price S_t . The upper row is for the case without fixed costs and the lower row is when $\beta_K = 0.12$ and $\beta_Q = 1.5$. The rest of the parameters are from Table 1.

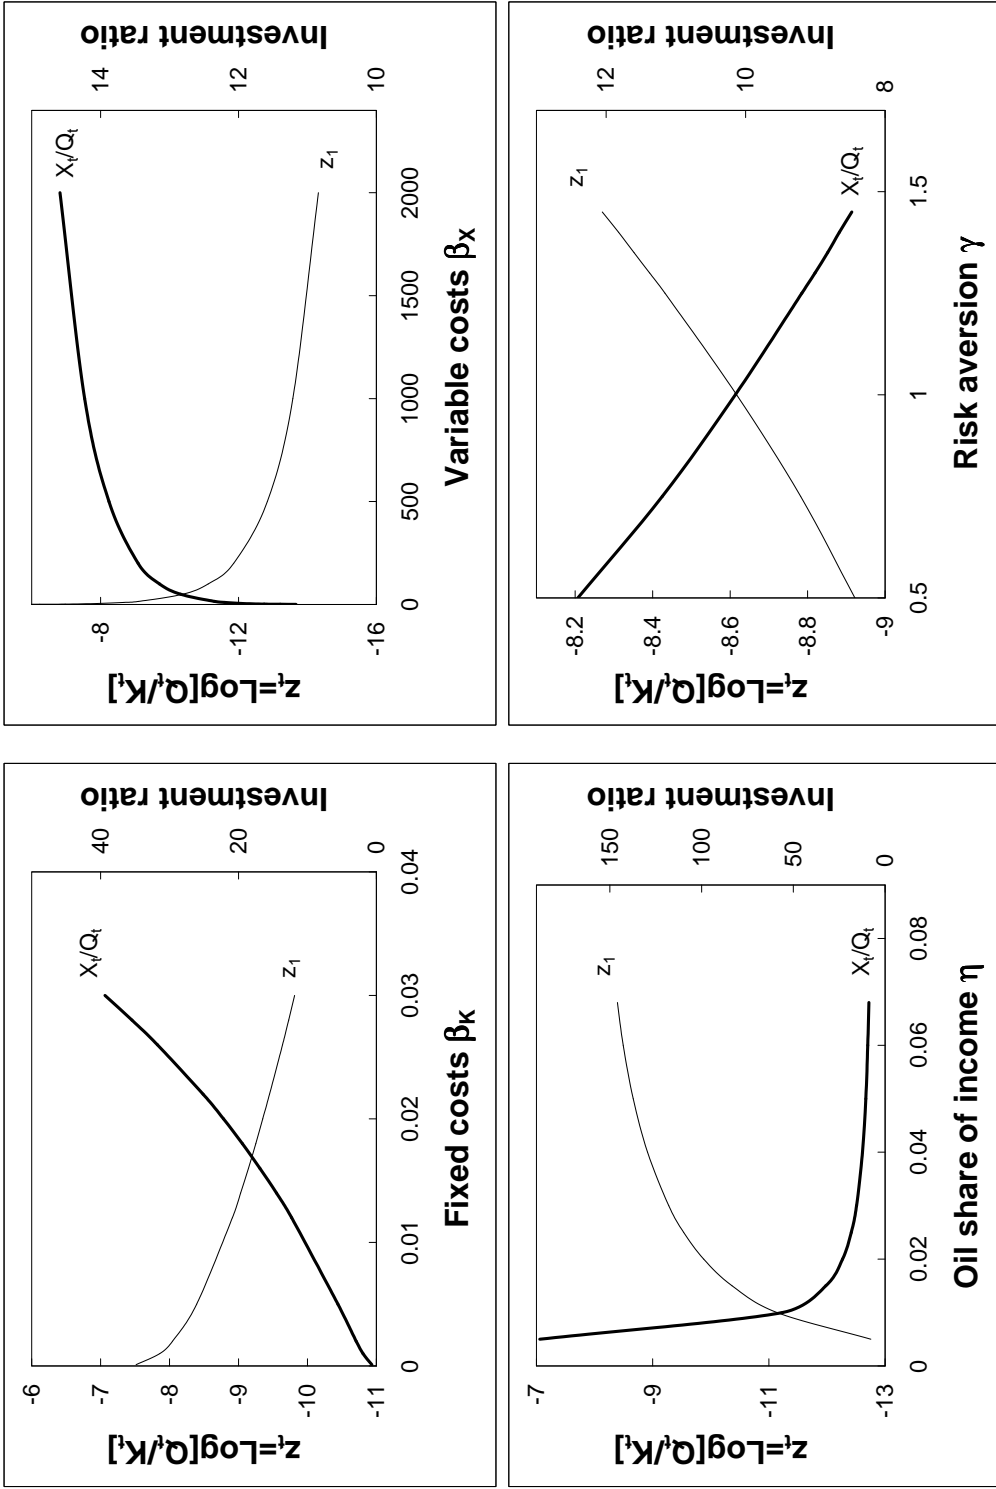


Figure 10: Investment strategy $\{z_1, z_2\}$ for different investment cost structure (β_K, β_X) , levels of risk aversion γ and oil share of income η . The thick (below) line corresponds to the investment trigger z_1 , while the thin (above) line is the returning point z_2 . To summarize both fixed cost components in the parameter β_K , we assume that $\beta_Q = \beta_K \beta_X$ for this plots.




Figure 11: Futures curves for contracts on oil for different spot prices. The thick curves are for spot prices in the *far-from-investment* region and the thin lines when the spot price is in the *near-investment* region. We use the parameters in Table 1 and the endogenous upper bound for the price is $S_{Max} = 29.58$.

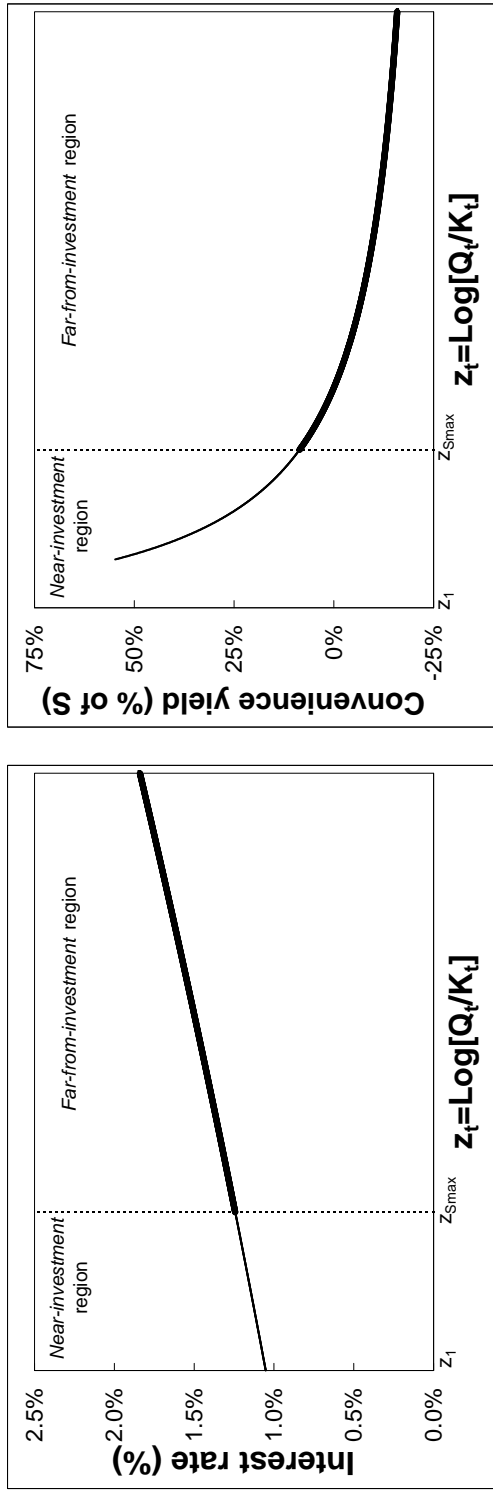


Figure 12: Interest rate and convenience yield as a function of the state variable z_t when $i_t = \bar{i}$. The thick line is the convenience yield when the economy is in the *far-from-investment* region and the thin for the economy in the *near-investment* region. We use the parameters in Table 1.




Figure 13: Relative volatility of futures contracts on oil to spot price volatility for different spot prices and maturities. The thick curves are for spot prices in the *far-from-investment* region and the thin lines when the spot price is in the *near-investment* region. We use the parameters in Table 1 and the endogenous upper bound for the price is $S_{Max} = 29.58$.




Figure 14: Historical Brent crude oil prices between Apr-1983 and Apr-2005 deflated by the US Consumer Price Index (thick line) and inferred probability of being in the *near-investment* state (thin line).




Figure 15: Implied commodity risk premium as a function of the state variable z_t when $i_t = \bar{i}$. The thick line is the risk premium when the economy is in the *far-from-investment* region and the thin for the economy in the *near-investment* region. We use the parameters in Table 1.




Figure 16: Historical performance of alternative investments from Jan-1985 to Jan-2005.