Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technologies*

Jaime Casassus
Pontificia Universidad Catolica de Chile and
Carnegie Mellon University

Pierre Collin-Dufresne
UC Berkeley

Bryan R. Routledge
Carnegie Mellon University

November 2003
Revised: March 2004
VERY PRELIMINARY

*Thanks to seminar participants at Carnegie Mellon University, the University of Madison-Wisconsin and at the 2003 LACEA Meetings. Correspondence: Jaime Casassus, Graduate School of Industrial Administration, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, 15213.

Equilibrium Commodity Prices with Irreversible Investment and Non-Linear Technologies

Abstract

We model the properties of equilibrium spot and futures oil prices in a general equilibrium production economy with two goods. In our model production of the consumption good requires two inputs: the consumption good and a 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. Equilibrium spot price behavior is determined as the shadow value of oil. The resulting equilibrium oil price exhibits mean-reversion and heteroscedasticity. Further, 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. Further, our model captures many of the stylized facts of oil futures prices. The futures curve exhibits backwardation as a result of a convenience yield, which arises endogenously due to the productive value of oil as an input for production. This convenience yield is decreasing in the amount of oil available in the economy. We test out model using crude oil data from 1982 to 2003. We estimate a linear approximation of the equilibrium regime-shifting dynamics implied by our model. Our empirical specification successfully captures spot and futures data. Finally, the specific empirical implementation we use is designed to easily facilitate commodity derivative pricing that is common in two-factor reduced form pricing models.

Keywords: Commodity prices, Futures prices, Convenience yield, Investment, Irreversibility, General equilibrium

JEL Classification: C0, G12, G13, D51, D81, E2.

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 commodity prices. 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. The model captures many stylized facts of the data.

Robust features exhibited by time series of commodity spot and futures prices are 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 2002)). Interestingly, the empirical evidence also suggests that there are marked differences across different types of commodity prices (e.g., Fama and French (1987)). Casassus and Collin-Dufresne (2002) use panel data (cross-section and time series) 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 2 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 costs[1] 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 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

[1] The 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 (2002) 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 time[2] in oil futures markets, whereas stock-outs are the exception rather than the rule.

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 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

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

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 in the investment region. This convenience yield reflects the benefit to smoothing the flow of oil used in production. It is decreasing in the outstanding stock of oil wells.

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 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 (for reasonable parameters) 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.

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 (2002)). Second, equilibrium models deliver economically consistent long-term predictions.

[3] Real 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)).

This may be a great advantage compared to reduced from models, which, due to the non-availability of data, may be hard to calibrate for long-term investment horizons.

With this in mind we estimate our model. The price follows a highly non-linear dynamics whose moments need to be calculated numerically. For this reason, we consider a linear approximation for the price process described above. The approximated model is desirable as well because, once estimated, it can be used straightforwardly for financial applications, like valuation or risk management. Our model has regime switching between the *near-investment* and the *far-from-investment* regions. The linearization implies that the price process is exponentially affine conditional on the regime. Under this representation is it straightforward to calculate a good approximation of the likelihood. Therefore, we use the quasi-maximum likelihood technique of Hamilton (1989) to estimate our model with crude oil data from 1990 to 2003. We find that most parameters are significant for both regimes, which validates our model. There is an infrequent state that is characterized by high prices and negative return and a more frequent that has lower average prices 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. 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.

The model presented here is related to many non-commodity areas. Our model is based upon a Cox, Ingersoll and Ross (1985) economy.[4] Dumas(1992) and Uppal (1993) follow CIR and set up the grounds for analyzing dynamic GE models in two-sector economies, particularly, they study the real-exchange rate in an international economy with two countries and shipping cost for transfers of capital. 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

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

goods 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.[5] Fixed adjustment 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. In the investments literature, 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.

Section 2 presents the model. Section 3 characterizes equilibrium commodity prices in our benchmark model with irreversibility and costly oil production. Section 4 considers the special case, where the oil flow rate of each well is flexible with adjustment costs for this type of flexibility. Section 5 presents the empirical estimation of the model and discusses its economic implications. Finally, Section 6 concludes.

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

[5] Surprisingly, there are not many models that use these type of production technologies in continuous time. Recently, Hartley and Rogers (2003) has extended the Arrow and Kurz (1970) two-sector model to an stochastic framework and use this type of production technology with private and government capital as inputs.

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

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

Their 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 equation (3) below:

$$dQ_t = -(i_t + \delta)Q_t dt + X_t dI_t \tag{2}$$

$$dK_t = (f(K_t, i_t Q_t) - C_t)\,dt + \sigma K_t dw_t - \beta(X_t; Q_t, K_t) dI_t. \tag{3}$$

The oil 'industry' produces a flow of oil at rate i_t and depreciates at rate δ. 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:

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

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).[6] 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 that

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

The case where $\beta_K = \beta_Q = 0$ can be recovered by taking the appropriate limit as shown in Jeanblanc-Picque and Shiryaev (1995) and we discuss it in the appendix. Further, to insure that investment is feasible we assume that:

$$\beta_K < 1.$$

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.

For simplicity we assume in this section that the extraction rate per unit time of each oil well is fixed at $i_t = \bar{i}$. This is meant to capture the fact that it is very costly to increase or decrease the production flow of oil wells. In practice this is true within certain limits. We thus reconsider the model with an optimally chosen extraction rate in the presence of adjustment costs in Section 4.

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.[7] The creation of X_t new oil costs $\beta(X_t; Q_t, K_t)$ of the numeraire good. This cost is borne only when investment occurs

[6] The 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.

[7] There 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.

($dI_t = 1$). Uncertainty in our economy is captured by the Brownian motion w_t which drives the diffusion term of the return of the production technology in equation (3). We assume that there exists an underlying probability space $(\Omega, \mathbf{F}, P)$ satisfying the usual conditions, and where $\mathbf{F} = \{\mathcal{F}\}_{t\geq 0}$ is the natural filtration generated by the Brownian Motion w_t.

Given our previous discussion it is natural to seek an investment policy of the form $\{(X_{T_i}, T_i)\}_{i=0,1,\ldots}$ where $\{T_i\}_{i=0,\ldots}$ 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:

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

Let us denote by $J(t, K, Q) = \sup_{\mathcal{C};\mathcal{A}} E_t[\int_t^\infty e^{-\rho s} U(C_s) ds]$ 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 exists. We note that this is slightly different than in traditional models with 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:

$$dW_t = (\alpha(\bar{i}Q_t)^\eta K_t^{1-\eta} - C_t - \beta_X(i+\delta)Q_t)dt + \sigma K_t dw_t \tag{6}$$

Since along each path, the agent can freely choose the ratio of oil to capital stock $Z_t = \frac{Q_t}{K_t}$, the above suggests that she should optimally do so to maximize point-wise the expected return of total wealth, i.e., such as to $\max_Q \left[\alpha(\bar{i}Q)^\eta K^{1-\eta} - \beta_X(i+\delta)Q\right]$, which gives:

$$\frac{Q_t}{K_t} = \left(\frac{\alpha i^\eta \eta}{\beta_X(\bar{i}+\delta)}\right)^{\frac{1}{1-\eta}} \equiv Z^* \tag{7}$$

This suggests that it is optimal to maintain a constant ratio of oil wells to consumption good stock point-wise. It also gives the optimal investment policy, which should satisfy:

$$dK_t + \beta_X dQ_t = 0 \tag{8}$$

Using equations (7) and (8) we may rewrite the dynamics of W_t as

$$\frac{dW_t}{W_t} = ((1-\eta)(\bar{i}Z^*)^\eta - c_t^u)\,dt + \sigma dw_t \tag{9}$$

where we define

$$C_t = c_t^u W_t. \tag{10}$$

The proposition below verifies that if

$$a^u := \frac{1}{\gamma}\left\{\rho - (1-\gamma)\left(\alpha(1-\eta)(iZ^*)^\eta - \gamma\frac{\sigma^2}{2}\right)\right\} > 0 \tag{11}$$

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 condition (11) holds then the optimal value function is given by*

$$J_u(t,K,Q) = e^{-\rho t}\frac{(a^u)^{-\gamma}(K+\beta_X Q)^{1-\gamma}}{1-\gamma} \tag{12}$$

The optimal consumption policy is

$$C_t^* = a^u(K_t^* + \beta_X Q_t^*) \tag{13}$$

and the investment policy is characterized by:

$$\frac{Q_t^*}{K_t^*} = Z^* \tag{14}$$

where Z^* is the constant defined in equation (7).

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

$$\begin{aligned}\frac{dJ_u(t,K_t,Q_t)+U(t,C_t)dt}{J_u(t,K_t,Q_t)} &= \left\{(a^u)^\gamma (c_t^u)^{1-\gamma} - (1-\gamma)c_t^u\right\}dt && (15)\\ &+\left\{(1-\gamma)\left(\alpha(\bar{i}Z_t)^\eta - \beta_X(\bar{i}+\delta)Z_t - \frac{\gamma\sigma^2}{2}\right) - \rho\right\}dt + (1-\gamma)\sigma dw_t && (16)\end{aligned}$$

where we have defined $C_t = c_t^u(K_t + \beta_X Q_t)$. Using the definition of Z^* and a^u in respectively (7) and (11) we have:

$$J_u(T,K_T,Q_T) + \int_0^T U(t,C_t)dt \le J_u(0,K_0,Q_0) + \int_0^T J_u(t,K_t,Q_t)(1-\gamma)\sigma dw_t \tag{17}$$

Taking expectations on both sides (and assuming that the stochastic integral is a martingale) we obtain:

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

with equality when we choose the controls $c_t^u = a^u$ and $Z_t = Z^*$. Further we note that for this choice of controls, we have:

$$\frac{dJ_u}{J_u} = -a^u dt + \sigma(1-\gamma)dw_t \tag{19}$$

which implies that

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

under the assumption (11). It also shows that the stochastic integral above is a square integrable martingale for this choice of control. Letting $T \to \infty$ in (18) shows that our candidate value function indeed is the optimal value function and confirms that the chosen controls are optimal. $\Box$

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 (1976) or Cox-Ingersoll-Ross (1985a) economy. In that case the condition on the coefficient a^u becomes:

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

which we assume below for simplicity.

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

$$\frac{dZ_t}{Z_t} = (i-\delta)dt + \sigma dw_t \tag{21}$$

$$\frac{dK_t}{K_t} = \sigma dw_t \tag{22}$$

It follows that if the following condition holds:

$$a^l := \rho + (1-\gamma)\left(\eta(\bar{\imath}+\delta) - (1-\eta)(\eta+\gamma(1-\eta))\frac{\sigma^2}{2}\right) > 0 \tag{23}$$

then, we have

$$J_l(0, K_0, Q_0) := E[\int_0^\infty e^{-\rho t}\frac{(C_t^l)^{1-\gamma}}{1-\gamma}dt] = \frac{(K_0\alpha\bar{\imath}^\eta)^{1-\gamma}}{(1-\gamma)a^l} \tag{24}$$

We collect the two previous results and a few simple 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. $J_l(t,K,Q) \le J(t,K,Q) \le J_u(t,K,Q)$.
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 (11) and (23) are satisfied, i.e., that $a^l, a^u > 0$.

2.3 Optimal consumption 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:

$$\sup_{\{C \geq 0\}} \{-\rho J + U(C) + \mathcal{D}J\} = 0 \tag{25}$$

where $\mathcal{D}$ is the Itô operator

$$\mathcal{D}J(K, Q) \equiv (f(K, \bar{i}Q) - C)\, J_K + \frac{1}{2}\sigma^2 K^2 J_{KK} - (\bar{i} + \delta) Q J_Q \tag{26}$$

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 (25) 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

$$C_t^* = J_K^{-\frac{1}{\gamma}}. \tag{27}$$

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 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

$$S_t = \frac{J_Q}{J_K} \tag{28}$$

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 $\{K_0, Q_0\}$.[8] 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.

$$J_1 = J_2 \tag{29}$$

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 (smooth pasting) conditions.[9]

$$J_{1K} = (1 - \beta_K) J_{2K} \tag{30}$$

$$J_{1Q} = -\beta_Q J_{2K} + J_{2Q} \tag{31}$$

$$0 = -\beta_X J_{2K} + J_{2Q} \tag{32}$$

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

[9] For 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 (30) to (33) result directly from equation (29) 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 $-J_{1QQ} + \beta_X J_{1KQ} = 0$.

These equations imply that

$$(\beta_X - \beta_Q)J_{1K} - (1-\beta_K)J_{1Q} = 0. \tag{33}$$

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

$$J(K,Q) = a_0^{-\gamma}\frac{K^{1-\gamma}}{1-\gamma}j(z) \tag{34}$$

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

$$z = \log\left(\frac{Q}{K}\right) \tag{35}$$

and a_0 is a constant.[10] The dynamic process for z_t is obtained using a generalized version of Itô's Lemma.

$$dz_t \;=\; \mu_{zt}dt - \sigma dw_t + \Lambda_z dI_t^* \tag{36}$$

where

$$\begin{aligned} \mu_{zt} &= -\left(f(1,\bar{i}e^{z_t}) - c_t^* + \bar{i} + \delta - \frac{1}{2}\sigma^2\right) & (37)\\ \Lambda_z &= z_2 - z_1 & (38) \end{aligned}$$

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 (29), 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

[10] Simply for convenience, we set introduce the coefficient $a_0 > 0$ so that as noted in the previous section in the special case $\eta = 0$ (e.g., oil is not used for production), $j(z) = 1$.

$z_t > z_1$ 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.

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

The jump in oil wells is

$$\frac{Q_2}{Q_1} = 1 + x^* \tag{40}$$

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

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

Finally, the optimal consumption from (27) can be rewritten in terms of j as:

$$c_t^* = \frac{C_t^*}{K_t^*} = a_0 \left(j(z_t) - \frac{j'(z_t)}{(1-\gamma)} \right)^{-\frac{1}{\gamma}} \tag{42}$$

Plugging this into the Hamilton-Bellman-Jacobi in equation (25) we obtain one-dimensional ODE for the function j.

$$\begin{aligned} \theta_0 j(z) + \theta_1 j'(z) + \theta_2 j''(z) + a_0 \gamma \left(j(z) - \frac{j'(z)}{1-\gamma} \right)^{1-\frac{1}{\gamma}} & \\ + \alpha(\bar{i}\, e^z)^\eta \big((1-\gamma) j(z) - j'(z) \big) &= 0 \end{aligned} \tag{43}$$

where

$$\theta_0 = -\rho - \frac{1}{2}\gamma(1-\gamma)\sigma^2, \quad \theta_1 = -\bar{i} - \delta - \frac{1}{2}(1-2\gamma)\sigma^2, \quad \theta_2 = \frac{1}{2}\sigma^2 \tag{44}$$

To determine the investment policy, $\{z_1, z_2\}$, the value-matching condition of equation (29) becomes:

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

Lastly, using the homogeneity there are only two super-contact conditions to determine that capture equations (30), (31), and (32).[11] They are

$$
\begin{aligned}
(1-\gamma)e^{z_1}(\beta_X - \beta_Q)j(z_1) - \big(1 - \beta_K + e^{z_1}(\beta_X - \beta_Q)\big)\, j'(z_1) &= 0 \qquad (46)\\
(1-\gamma)e^{z_2}\beta_X j(z_2) - (1 + e^{z_2}\beta_X)j'(z_2) &= 0 \qquad (47)
\end{aligned}
$$

The following proposition summarizes the above discussion and offers a verification argument.

Proposition 3 *Suppose that we can find two constants* z_1, z_2 $(0 \le z_1 \le z_2)$ *and a* $C^2(z_1, \infty)$ *function* $j(\cdot)$*, which solve the ODE given in equation (43) with boundary conditions (45), (46), and (47), then the value function is given by*

$$J(t, K, Q) = e^{-\rho t} a_0^{-\gamma} \frac{K^{1-\gamma}}{1-\gamma} j(z) \tag{48}$$

where $z = \log \frac{Q}{K}$*. Further the optimal consumption policy is given in equation (42). The optimal investment policy consists of a sequence of stopping times and investment amounts,* $\{(T_i, X_{T_i})\}_{i=0,2\ldots}$ *given by* $T_0 = 0$ *and:*

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

$$X_0^* = Q_0 \frac{e^{-z_0}(1 - \beta_K) - e^{-z_2} - \beta_Q}{e^{-z_2} + \beta_X} \tag{49}$$

  Then start anew 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:*

$$T_i = \inf\{t > T_{i-1} : z_{t^-} = z_1\} \quad i = 1, 2, \ldots \tag{50}$$

[11] In a similar way, if $\beta_K = \beta_Q = 0$ the two super-contact conditions presented in footnote (9) become the same condition $(1 + (1-\gamma)e^{z_1}\beta_X)j'(z_1) - (1 + e^{z_1}\beta_X)j''(z_1) = 0$.

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

$$X^*_{T_i} = Q_{T_i} \frac{e^{-z_1}(1-\beta_K) - e^{-z_2} - \beta_Q}{e^{-z_2} + \beta_X} \tag{51}$$

Proof Applying the generalized Itô's lemma to our candidate value function we find:

$$\begin{aligned} dJ(t, K_t, Q_t) + U(t, C_t)dt &= \frac{e^{-\rho t} a_0^{-\gamma} K_{t^-}^{1-\gamma}}{1-\gamma} \Big\{ \Big[\hat{\theta}_0(z_t) j(z_t) + \hat{\theta}_1(z_t) j'(z_t) + \theta_2 j''(z_t) \\ &\quad + a_0^{\gamma}(c_t)^{1-\gamma} - c_t \big((1-\gamma)j(z_t) - j'(z_t)\big)\Big] dt + \big((1-\gamma)j(z_t) - j'(z_t)\big)\sigma dw_t \\ &\quad + \sum_{T_i \le t} \left(\frac{1-\beta_K + e^{z_{T_i^-}}(\beta_X - \beta_Q)}{1+\beta_X e^{z_{T_i}}}\right)^{1-\gamma} j(z_{T_i}) - j(z_{T_i^-}) \Big\} \end{aligned} \tag{52}$$

where for simplicity we have defined $\hat{\theta}_0(z) = \theta_0 + (1-\gamma)\alpha(\bar{i}e^z)^\eta$ and $\hat{\theta}_1(x) = \theta_2 - \alpha(\bar{i}e^z)^\eta$ and $C_t = c_t K_t$. Suppose we can find a function $j(\cdot)$ defined on some closed domain D, such that for any $y > x$ (with $y, x \in D$) we have

$$\left(\frac{1-\beta_K + e^x(\beta_X - \beta_Q)}{1+\beta_X e^y}\right)^{1-\gamma} j(y) - j(x) \le 0$$

and

$$\hat{\theta}_0(z)j(z) + \hat{\theta}_1(z)j'(z) + \theta_2 j''(z) + \sup_c \left[a_0^\gamma (c)^{1-\gamma} - c\big((1-\gamma)j(z) - j'(z)\big)\right] \le 0$$

then we have

$$J(T, K_T, Q_T) + \int_0^T U(t, C_t)dt \le J(0, K_0, Q_0) + \int_0^T \frac{e^{-\rho t} a_0^{-\gamma} K_{t^-}^{1-\gamma}}{1-\gamma} \big((1-\gamma)j(z_t) - j'(z_t)\big)\sigma dw_t \tag{53}$$

Under the assumption of the proposition j is such a function. Furthermore the Bellman equation (43) guarantees that for the candidate choice of control for consumption (given in (42)) the drift is zero, and the value matching condition (45) insures that at the optimum the jump is zero. Thus taking expectation (and assuming that the stochastic integral is a martingale) we get

$$E\left[J(T,K_T,Q_T)+\int_0^T U(t,C_t)dt\right] \le J(0,K_0,Q_0) \tag{54}$$

with equality for our choices of optimal controls. It remains to show that $\lim_{T\to\infty} E\left[J(T,K_T,Q_T)\right] = 0$ and that the stochastic integral is indeed a true martingale. To be completed... $\square$

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. We first describe the pricing kernel and financial asset prices. Next, we use the marginal value of a unit of oil, as in equation (28), 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 the markets are dynamically complete, the pricing kernel is characterized by the representative agent's optimal solution (see Duffie (1996)). First, define the risk-free money-market account whose price is B_t. The process for the money market price is

$$\frac{dB_t}{B_t} = r_t dt + \Lambda_B dI_t \tag{55}$$

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. In equilibrium, the Λ_B is a constant.

The pricing kernel for our economy satisfies

$$\frac{d\xi_t}{\xi_t} = -\frac{dB_t}{B_t} - \lambda_t dw_t \tag{56}$$

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.[12]

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

$$\xi_t = e^{-\rho t}\frac{J_K(K_t, Q_t)}{J_K(K_0, Q_0)} \tag{57}$$

$$r_t = f_K(K_t, \bar{i}Q_t) - \sigma\lambda_t \tag{58}$$

$$\lambda_t = -\sigma\frac{K_t J_{KK}}{J_K} \tag{59}$$

$$\Lambda_B = -\frac{\beta_K}{1-\beta_K} \tag{60}$$

where $f_K(.,.)$ is the first derivative of the production function with respect its first argument. Moreover, the equilibrium interest rate and market price of risk are only functions of the state variable z_t, i.e. $r_t = r(z_t)$ and $\lambda_t = \lambda(z_t)$. [13]

Proof See the Appendix. □

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 equation (59) 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 (30) we can calculate the size of the jump in the stochastic

[12] The 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.

[13] We 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.

discount factor and note that it depends on the oil investment cost structure. In particular, recall from (4) that creating X_t new oil wells costs $\beta(X_t, Q_t, K_t)$ units of the consumption good. Equilibrium financial prices will jump if $\beta_K > 0$ where β_K determines how the cost function is related to the size of the numeraire industry.

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 (28), is a function of the ratio of oil to numeraire good state variable, z_t.

$$S_t = \frac{J_Q}{J_K} = \frac{e^{-z_t} j'(z_t)}{(1-\gamma)j(z_t) - j'(z_t)} \tag{61}$$

To characterize the oil spot price behavior, consider the spot price at the investment boundary, z_1. From the smooth-pasting condition in equation (32), the oil price immediately after new investment is

$$S_{2,t} = \beta_X \tag{62}$$

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 (33) implies that

$$S_{1,t} = \frac{\beta_X - \beta_Q}{1 - \beta_K} \tag{63}$$

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

$$\Lambda_S = \frac{\beta_Q - \beta_K \beta_X}{1 - \beta_K} \tag{64}$$

Since oil is not a traded financial asset, the jump in the price of oil can be different that the Λ_B jump in financial prices. The only situation that produces both asset and oil prices that have no jumps is when there is no fixed cost to investing in oil ($\beta_K = \beta_Q = 0$), hence investment is not lumpy. However, it is also possible to generate continuous asset prices and discontinuous oil prices ($\beta_K = 0$, $\beta_Q > 0$).

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. The simulations that follow illustrate this case.

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 parameters for the examples shown in this section are in Table 1. 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 implies that it is not optimal to make a new investment as soon as the spot price (marginal value 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; that is $S(z_1) = S(z_2)$.

We use the maximum price $S_{\max}$ to partition the state space into two regimes. On the right in Figure 4 with $z_t \geq z_{Smax}$ is the *far-from-investment* zone. In this region, investment is new oil in the short term is sufficiently unlikely, and the oil price is decreasing in z_t. On the left in Figure 4 with $z_1 < z_t \leq z_{Smax}$ is the *near-investment* zone. In this region, the likelihood of investment in new oil dominates and a decrease in the stock of oil, z_t declines, reduces the price in anticipation of the increased future oil stocks. 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 (at least once) 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 3 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. Note that there is a one-to-one mapping between $\{S_t, \varepsilon_t\}$ and z_t.

Proposition 5 *The oil price in equation (61) is governed by the following two-regime stochastic process*

$$\begin{aligned} dS_t &= \mu_S(S_t,\varepsilon_t)S_t dt + \sigma_S(S_t,\varepsilon_t)S_t dw_t + \Lambda_S dI_t && (65)\\ \mu_S(S_t,\varepsilon_t) &= r(S_t,\varepsilon_t) - y(S_t,\varepsilon_t) + \sigma_S(S_t,\varepsilon_t)\lambda(S_t,\varepsilon_t) && (66)\\ \sigma_S(S_t,\varepsilon_t) &= \frac{(S_t + e^{-z(S_t,\varepsilon_t)})\Lambda(S_t,\varepsilon_t) - e^{-z(S_t,\varepsilon_t)}\gamma\sigma}{S_t} && (67)\\ \Lambda_S &= \frac{\beta_Q - \beta_K \beta_X}{1-\beta_K} && (68) \end{aligned}$$

where

$$\varepsilon = \begin{cases} 1 & if \quad z > z_{Smax} \\ 2 & if \quad z_1 < z \leq z_{Smax} \end{cases} \qquad (69)$$

and where $r(S_t,\varepsilon_t) = r(z_t)$ *and* $\lambda(S_t,\varepsilon_t) = \lambda(z_t)$ *as in Proposition 1,* $z(S_t,\varepsilon_t) = z_t$ *and* $y(S_t,\varepsilon_t) = y_t$ *is the convenience yield defined later in equation (75).*

Proof See the Appendix. □

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 shows some degrees of 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 a lot of oil to produce capital, i.e., Q decreases quickly 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 one above the dashed-line and two 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_{Smax} 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. Note that variable z_t remains most of the time between -8 and -6 which is right near the boundary. For our example, 53% of the time the oil price is above the marginal cost (that is $z_1 < z_t < z_2$) and 10% of the time the economy is in the *near-investment* region ($z_t < Z_{Smax}$).[14]

From the previous discussion, the non-monotonicity in the relationship between the state variable, z_t and the spot price, S_t, is crucial for the regime shifting behavior

[14]Recall 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 $z_1 < z_t < z_2$.

of the spot price. The size of the hump in Figure 4 is determined by the optimal investment policy z_1 and z_2. In order for the hump to be large, investment in new oil wells needs to be large; that is the size of $z_2 - z_1$. To understand how investment policy is affected by our model parameters, Figure 9 shows the investment strategy defined z_1 and z_2 under various parameters. The graph on the upper-left corner shows the effect of economies of scale in the strategy $\{z_1, z_2\}$. The bigger is the fixed cost component β_K, the more is the investment delay (z_1 is decreasing in β_K). In these charts the difference $z_2 - z_1$ gives an idea of the optimal number of oil wells to be built from equation (39). When the fixed cost component is small, the number of new oil wells is low (in the limiting case, investment is infinitesimal). For higher fixed cost, the investment increases because of higher levels of economies of scale. The graph to the right shows that a higher marginal cost delays investments. The lower-left graph of figure 9 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 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.

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.[15] The stochastic process for the futures price is

$$\frac{dF_t}{F_t} = \mu_{F,t}dt + \sigma_{F,t}dw_t + \Lambda_F dI_t \tag{70}$$

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

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

Proposition 6 *The equilibrium futures price $F(z,t,T)$ in equation (70) satisfies $\mu_{F,t} = \sigma_{F,t}\lambda_t$ and $F(z_1,t,T) = F(z_2,t,T)$, implying $\Lambda_F = 0$ and the following partial differential equation*

$$\frac{1}{2}\sigma^2 F_{zz} + (\mu_z - \sigma\lambda_t)F_z + F_t = 0 \tag{71}$$

with boundary condition

$$F(z,T,T) = S(z) \tag{72}$$

Proof See the Appendix. □

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 slopped forward curve) is implied by the convenience yield. For example, Casassus and Collin-Dufresne (2002) 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. In our model, we can determine the convenience yield implicitly from equilibrium prices using the no-arbitrage condition for tradable assets

$$E_t^*\left[\frac{dS_t}{S_t}\right] = \frac{dB_t}{B_t} - \frac{dY_t}{S_t} \tag{73}$$

where E_t^* is the expectation under the equivalent martingale measure. 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 financial asset the convenience yield would be the dividend flow that implies no 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 implicit cumulative net convenience yield Y_t has the following dynamics:

$$dY_t = y_t S_t dt + \Lambda_Y S_t dI_t \tag{74}$$

If we compare the risk-adjusted drift of the price in equation (66) with the one from (73) we conclude that [16]

$$
\begin{aligned}
y_t &= \frac{\bar{i}}{S_t}(f_q(K_t, \bar{i}Q_t) - S_t) - \delta && (75)\\
\Lambda_Y &= \Lambda_B - \Lambda_S && (76)
\end{aligned}
$$

In our setting, there are two components to 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_q 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 when the economy is in the *near-investment* region than when it is in the *far-from-investment* region (see figure 10). Interestingly, in Section 4, where we allow for optimal oil extraction, i_t^*, this term vanishes in the case that the adjustment costs for substituting inputs are zero, i.e. $y_t = -\delta$.

The second component of convenience yield is the predictable jump that occurs in prices at the time of oil investment. $\Lambda_Y < 0$ represents the singular component, which represents arbitrage profits that agents could make were they able to buy the commodity in the investment region. Note that if one could short-sell the money-market fund and buy the commodity, one would lock a risk-free profit of $-\Lambda_Y > 0$. Of course, the commodity is not a financial assets, and its 'price' is the shadow value to the consumers of using it as an input to production, which is very high just prior to investing.

Figure 11 shows the futures prices for different spot prices and maturities. As with the process for spot prices in Proposition 2, 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 while the thin lines are for spot prices in the *near-investment* region. 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 can only decrease,

[16] The 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_q would be expressed in terms of f_z which has a less clear economic meaning.

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 $S_t = 15$ in figure 11). When the price is near the maximum price the futures curves are downward-sloping, i.e., backwardation (see the curves when $S_t = 30$). The expected price is below the current price, because of a high probability of an increase in oil supply (or a high net convenience yield). 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. In general, the futures curve are steeper when the spot price is in the *near-investment* region. This is a direct implication of a higher convenience yield in this region. This can also be interpreted as a likely sooner investment to create new oil. Our model also generates non-monotonic curves (see the humped curve when $S_t = 25$ 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 $S_t = 25$ 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.

Recall from equations (60) and (64) that both asset prices and Oil spot prices may jump at the (predictable) investment in oil. However, as shown in Proposition 3, 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 12. 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 heteroscedasticity in equilibrium futures prices. First, the volatility curves depend on the spot price. In most affine reduce-form models for commodity prices (see for example Schwartz (1997)), the futures price is proportional to the spot price, thus relative volatility ratio is assumed to be constant for any given spot price and maturity date T which does not occur in our model . 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 $\sigma_S(S_{max}, \varepsilon, t) = 0$). As in figure 7, a negative volatility implies a negative correlation between the spot price and the futures price. For example, if the spot price is very high and is in the *far-from-investment* region it has a very low volatility and is positively correlated with shocks in capital. In the near future, the price is expected to be in the *near-investment* region and to be negatively correlated with shocks in capital. This implies that the spot and futures price have negative correlation, which is shown with negative relative volatility values in the figure. This is the case for a 3-year maturity contract when the spot price is $S_t = 30$ (in the *far-from-investment* region).[17]

4 Extensions - Flexible production with adjustment cost

In solving the representative agent model in equation (5) we made the simplifying assumption that the production of oil in equation (2) was fixed at $i_t = \bar{i}$. In this section we explore the effect of relaxing this assumption by extending our model to include an optimal demand rate of oil. We consider a variant of the model proposed in the previous sections, where the production technology $f(K_t, i_t Q_t)$ is flexible in the sense that the sector chooses optimally the fraction of oil to use as an input i_t^*. There are adjustment costs for this type of flexibility when the optimal demand rate deviates from some target rate $\bar{i}$. Changes in the stocks of capital K_t and oil wells Q_t produces adjustment in the input rate, i_t.

[17] The relative volatility or hedge-ratio features that occur in our model are very similar to those found in the storage-based equilibrium model of Routledge, Seppi, and Spatt (2000).

The problem of the representative agent is similar to the one before, but the dynamics of the stocks K and Q account for the flexible demand/production of oil and the adjustment costs, ψ. The stochastic differential equations for these dynamics are

$$\begin{aligned} dQ_t &= -(i_t + \delta)Q_t dt + X_t dI_t && (77) \\ dK_t &= \left(f(K_t, i_t Q_t) - \frac{\psi}{2}(i_t - \bar{i})^2 K_t - C_t \right) dt + \sigma K_t dw_t - \beta(X_t) dI_t && (78) \end{aligned}$$

There is an extra first order constraint that determines the optimal demand for oil i_t^*. It is easy to check that this FOC is

$$f_q(K_t, i_t^* Q_t) - \psi(i_t^* - \bar{i})\frac{K_t}{Q_t} = S_t \tag{79}$$

where f_q is the first derivative of the production function with respect its second argument, i.e., the amount of oil used as an input. In equilibrium, marginal benefit of an extra unit of oil (f_q) minus adjustment costs equals its marginal cost (S_t). This model nests the previous one, because as the adjustment cost ψ tends to infinity, the optimal demand rate moves to the target rate

$$\lim_{\psi \to \infty} i_t^* = \bar{i} \tag{80}$$

Overall, the main results of the paper remain unchanged when we try different adjustment costs. The spot price follows a mean-reverting process and the same two regimes are present in the economy.

Figure 13 shows the dynamics of the demand rate and the convenience yield for different adjustment costs. The thin lines shows the values of these variables when the economy is in the *near-investment* region and the thick lines for the *far-from-investment* region. The plot on the left shows the logarithm of the input ratio, $\log(i_t^* Q_t / K_t)$, determined by equation (79). In the limiting case, when adjustment costs are very high, this relation is linear ($i_t^* \to \bar{i}$). The more oil in the economy, the more it is used as an input for the production technology f. With reasonable adjustment costs ($\psi = 1, \psi = 0.1$) this increasing relation is also true, but there is a tendency to use more oil when oil stocks are low. The reason is that the price of oil is lower for lower stocks near the investment region, so this benefit (from less costs of inputs) justifies adjustments in the demand rate. When the production technology can be adjusted costlessly ($\psi = 0$), what really matters is the price of oil S_t instead of the stocks of oil in the economy.[18] When the oil price is high, the input rate is low, and viceversa. In this case, the (log) input ratio inherits the two regimes present in the oil price.

The plot in the right of figure 13 shows the convenience yield for different adjustment costs. Interestingly, the convenience yield is increasing in the adjustment costs in the *near-investment* region (where the stocks of oil are small compared to the amount of capital in the economy). With small adjustment costs this benefit is less important because the oil can be replaced by capital at a small cost. In the case with no adjustment costs ($\psi = 0$), the convenience yield is zero and the net convenience yield equal the negative depreciation cost. Here there are no benefits from holding oil in the *no-investment* region, because it can be substituted (locally) at no cost. This is clear from equations (75) and (79) with $\psi = 0$. There remains a *singular* component to oil prices, which could be interpreted as a 'convenience yield' at the investment boundary. This result shows that a two-input production technology is necessary but not sufficient to generate positive convenience yields. It is necessary to have some degree of rigidity in the technology or adjustment costs (this is similar to the findings of Carlsson et al. (2002)).

5 Implementation and Estimation

We want to understand the empirical properties and implications of the model in Section 2. To do this we match the conditional moments of the oil prices in the data with the ones predicted by the structural model. We estimate a linear approximation version of the commodity pricing model in Proposition 5. 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

[18]Of course, the oil/capital ratio affects the price, but when $\psi = 0$ the input ratio is only a function of S_t. In this case, we can get the input ratio in closed-form from equation (79), i.e. $i_t^* Q_t / K_t = (\alpha\, \eta / S_t)^{\frac{1}{1-\eta}}$.

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 the structural model in Section 2 is

$$dS_t = \mu_S(S_t, \varepsilon_t) S_t dt + \sigma_S(S_t, \varepsilon_t) S_t dw_t \tag{81}$$

where

$$\mu_S(S, \varepsilon) = \alpha + \kappa_\varepsilon(\log[S_{Max}] - \log[S]) \tag{82}$$

$$\sigma_S(S, \varepsilon) = \sigma_\varepsilon \sqrt{\log[S_{Max}] - \log[S]} \tag{83}$$

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

$$P_t = \begin{bmatrix} 1 - \lambda_1 dt & \lambda_1 dt \\ \lambda_2 dt & 1 - \lambda_2 dt \end{bmatrix} \tag{84}$$

The process in equations (81)-(83) 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 (65) to (67). Equation (84) 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 in 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 Jan-1982 and Aug-2003 deflated by the US Consumer Price Index. The average price is 15.41 dollars per barrel in Jan-1982 prices (or 30.16 dollars per barrel in Aug-2003 prices). The annualized standard deviation of weekly returns is 38.1%. The skewness in crude oil prices for this period is 1.26 and the excess kurtosis is 0.62.

The parameter space for the approximated model in equations (81)-(83) 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 are doing 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 2. 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, $\lambda_1 = 0.984$, before switching to the second regime. Moreover, the first regime is the most frequent one, since the economy stays approximately 83.7% of the time in it ($\lambda_2/(\lambda_1 + \lambda_2) = 0.837$). The economy stays in the second regime on average a couple of months before jumping back to regime 1 ($\lambda_2 = 5.059$). The parameter α 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 ($S_{Max} = 39.8$). Figure 14 displays a graphical representation of the estimates of the drift and volatility of returns of oil. The figure shows that under the most frequent regime (thick line), 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 (thin line), since for reasonable prices the drift is significantly negative and almost constant ($\alpha = -0.248$ and $\kappa_2 = 0.055$). Also, the second regime is characterized to be more volatile than the first regime ($\sigma_2 > \sigma_1$).

Figure 15 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 2. Figure 15 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.[19] 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.[20] 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 ($\overline{F_1} - 2\overline{F_6} + \overline{F_{12}}$).[21]

Table 3 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

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

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

[21] The 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 ($\omega = F_1 - 2F_6 + F_{12}$).

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.[22] This occurs in our model because the convenience yield is higher in the *near-investment* region, which implies higher degrees of backwardation.

6 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 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. We estimate a linear approximation version of our model with crude

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

oil prices from 1982 to 2003. Our empirical specification successfully captures spot and futures data. Finally, the specific empirical implementation we use is designed to easily facilitate commodity derivative pricing that is common in two-factor reduced form pricing models.

References

Arrow, Kenneth J., and Mordecai Kurz, 1970, *Public Investment, the Rate of Return and Optimal Fiscal Policy* (Johns Hopkins Press: Baltimore).

Bailey, Roy, and Marcus Chambers, 1996, A theory of commodity price fluctuations, *Journal of Political Economy* 104(5), 924–957.

Bessembinder, Hendrick, Jay F. Coughenour, Seguin Paul J, and Margaret Monroe Smoller, 1995, Mean-reversion in equilibrium asset prices: Evidence from the futures term structure, *Journal of Finance* VOL. L NO.1 (March), 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, 2002, Maximal affine model of convenience yields implied from interest rates and commodity futures, Carnegie Mellon University Working Paper.

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).

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.

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.

Gibson, Rajna, and Eduardo S. Schwartz, 1990, Stochastic convenience yield and the pricing of oil contingent claims, *Journal of Finance* v45n3 (July), 959–76.

Grossman, Sanford J., and Guy Laroque, 1990, Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* 58, 2551.

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.

James L. Paddock, 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.

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 lenear 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.

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., 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* (Claredon Press: Oxford).

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.

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	0.05
Importance of oil,	η	0.02
Demand rate for oil,	$\bar{i}$	0.05
Volatility of return on capital,	σ	0.25
Depreciation of oil,	δ	0.01
Irrevertsible investment		
Fixed cost (K component),	β_0	0.005
Fixed cost (Q component),	β_1	0.01
Marginal cost of oil,	β_2	20
Agents preferences		
Patience,	ρ	0.05
Risk aversion,	γ	1.5

Table 1: Parameters used for examples and simulations in Section 2.

Parameter	Estimate	Std.Error	t-ratio
λ_1	0.984	0.312	3.2
λ_2	5.059	1.589	3.2
α	-0.248	0.103	-2.4
κ_1	0.319	0.116	2.7
κ_2	0.055	0.148	0.4
σ_1	0.251	0.010	25.8
σ_2	0.808	0.063	12.9
S_{Max}	39.8	0.399	99.7

Table 2: 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	***far-from-investment* state**			***near-investment* state**		
(\$/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 3: 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.

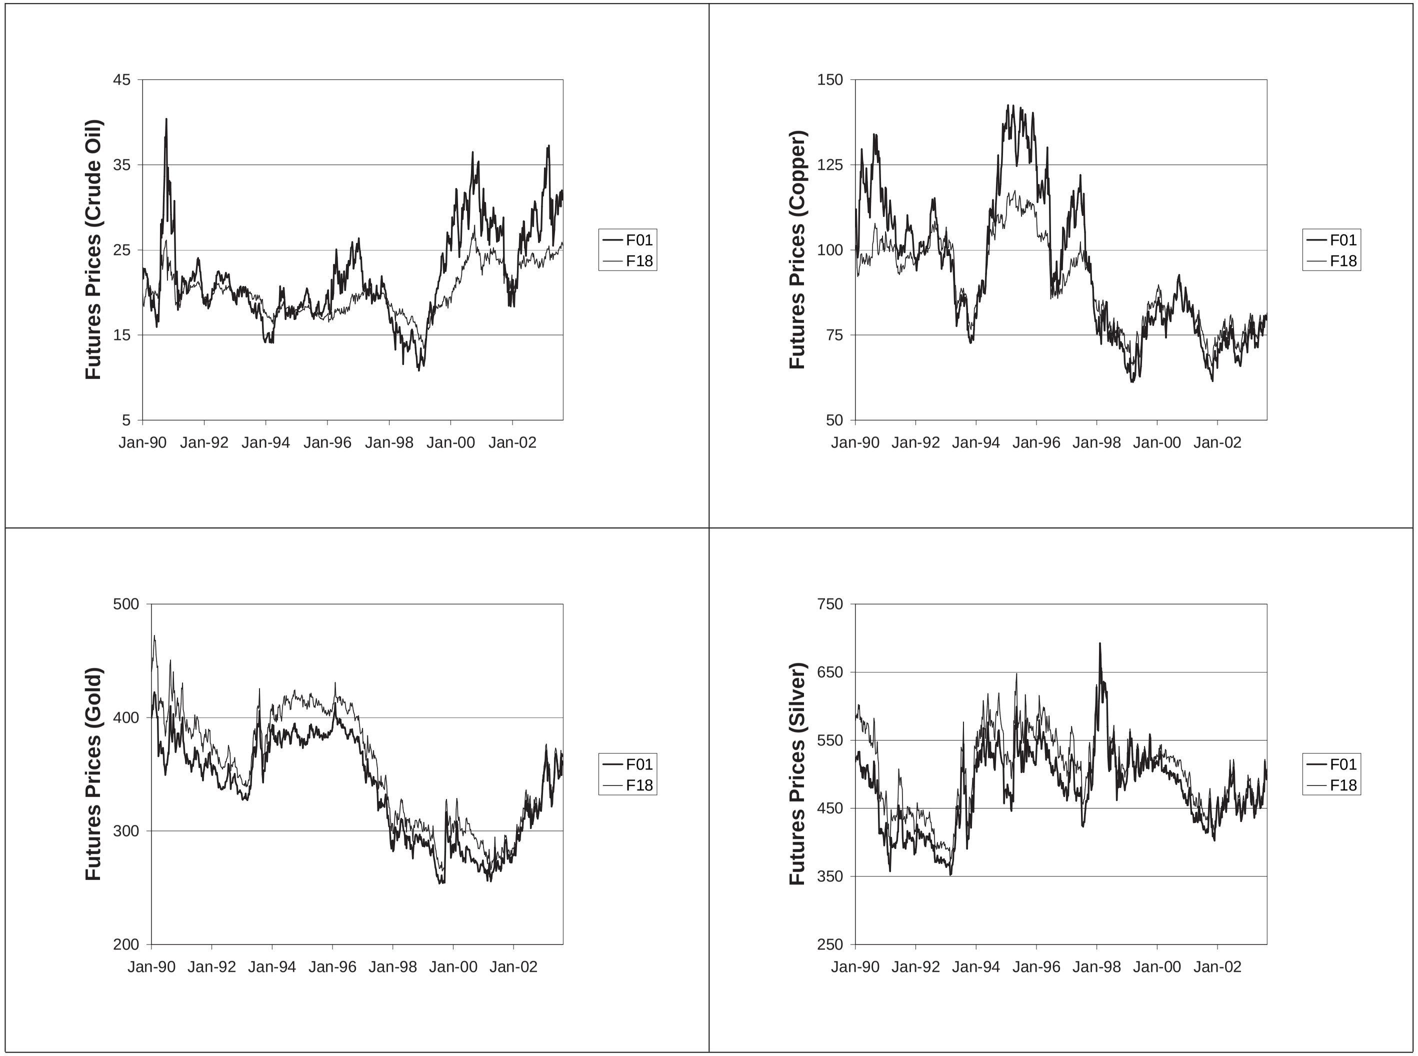

Figure 1: F01 and F18 futures contracts on crude oil, copper, gold and silver between 1/2/1990 to 8/25/2003. Oil prices are in dollars per barrel, copper prices are in cents per pound, gold prices are in dollars per troy ounce and silver prices are in cents per troy ounce.

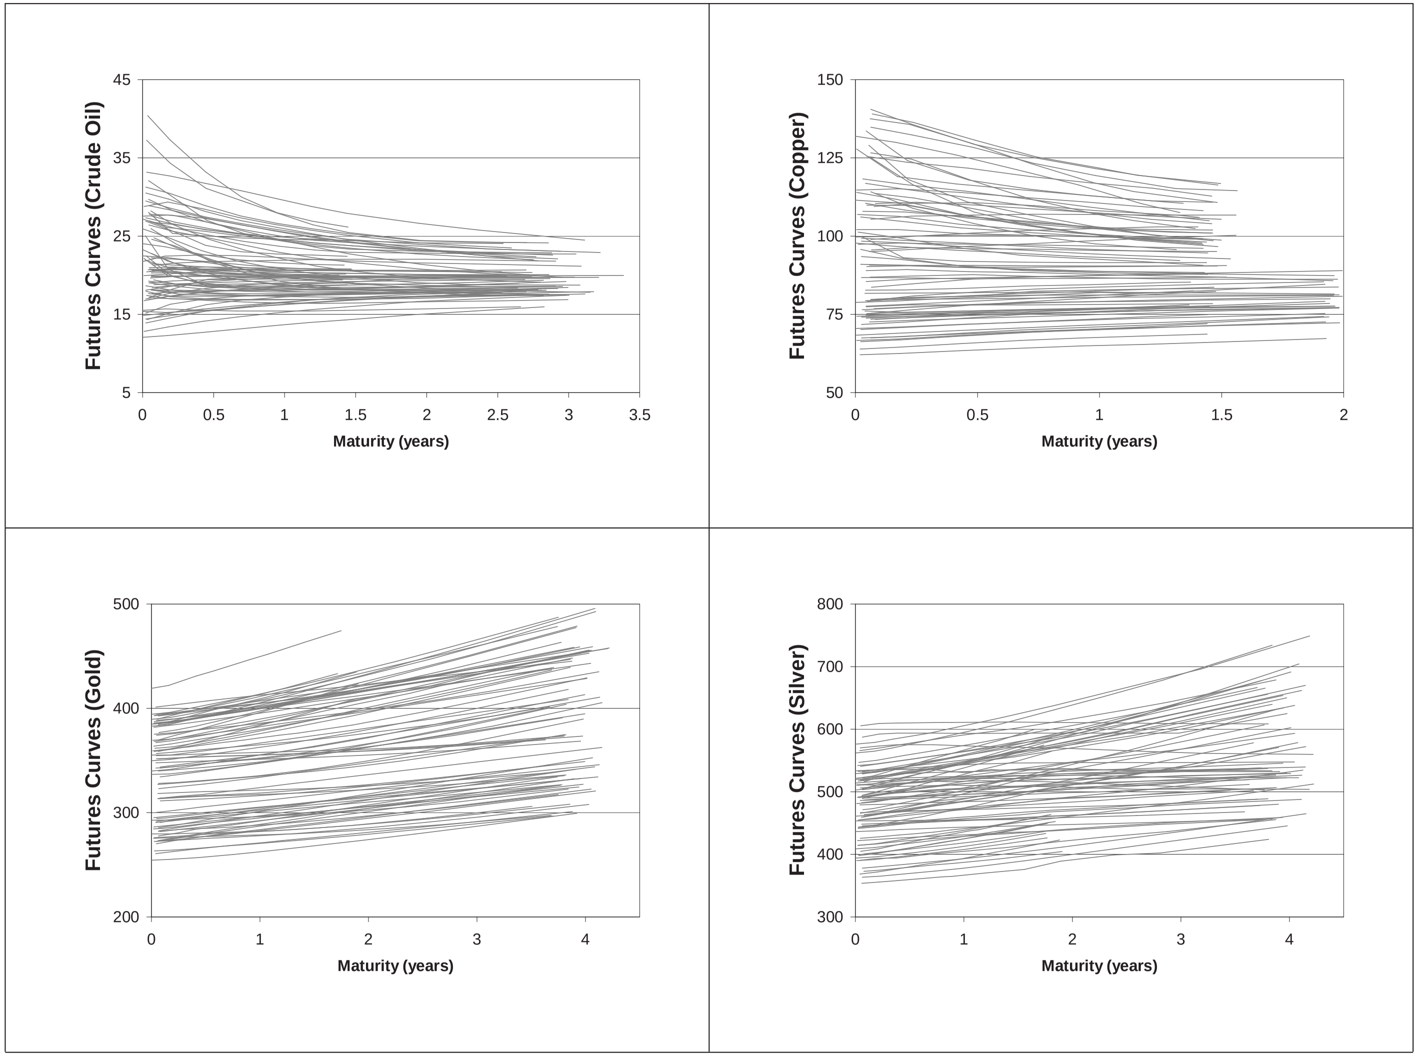

Figure 2: Monthly term structures of futures prices on crude oil, copper, gold and silver between 1/2/1990 to 8/25/2003. Oil prices are in dollars per barrel, copper prices are in cents per pound, gold prices are in dollars per troy ounce and silver prices are in cents per troy ounce.

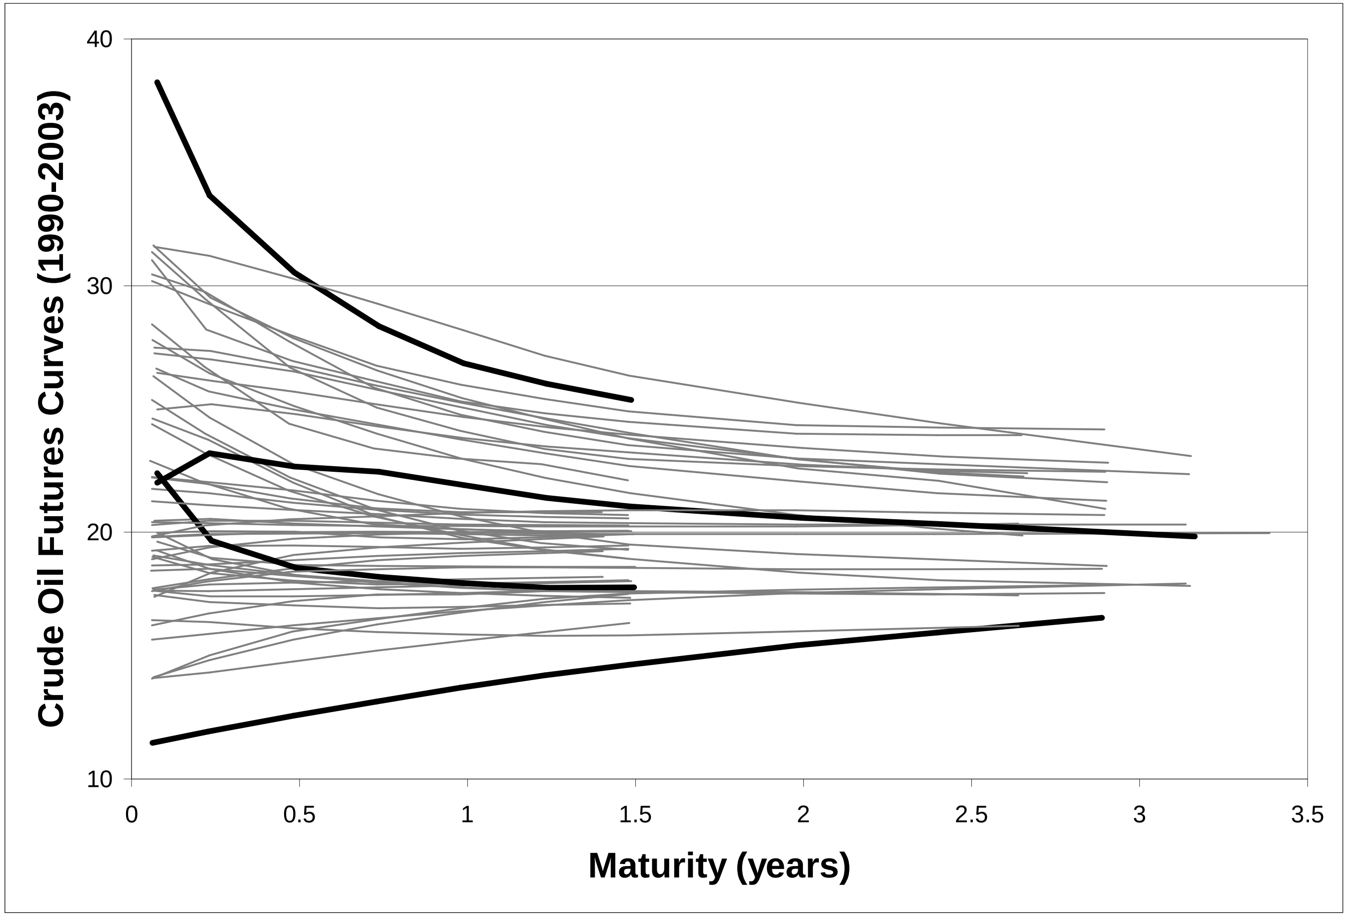

Figure 3: Sample futures curves from 1990-2003 for Brent crude (deflated by U.S. Consumer Price Index).

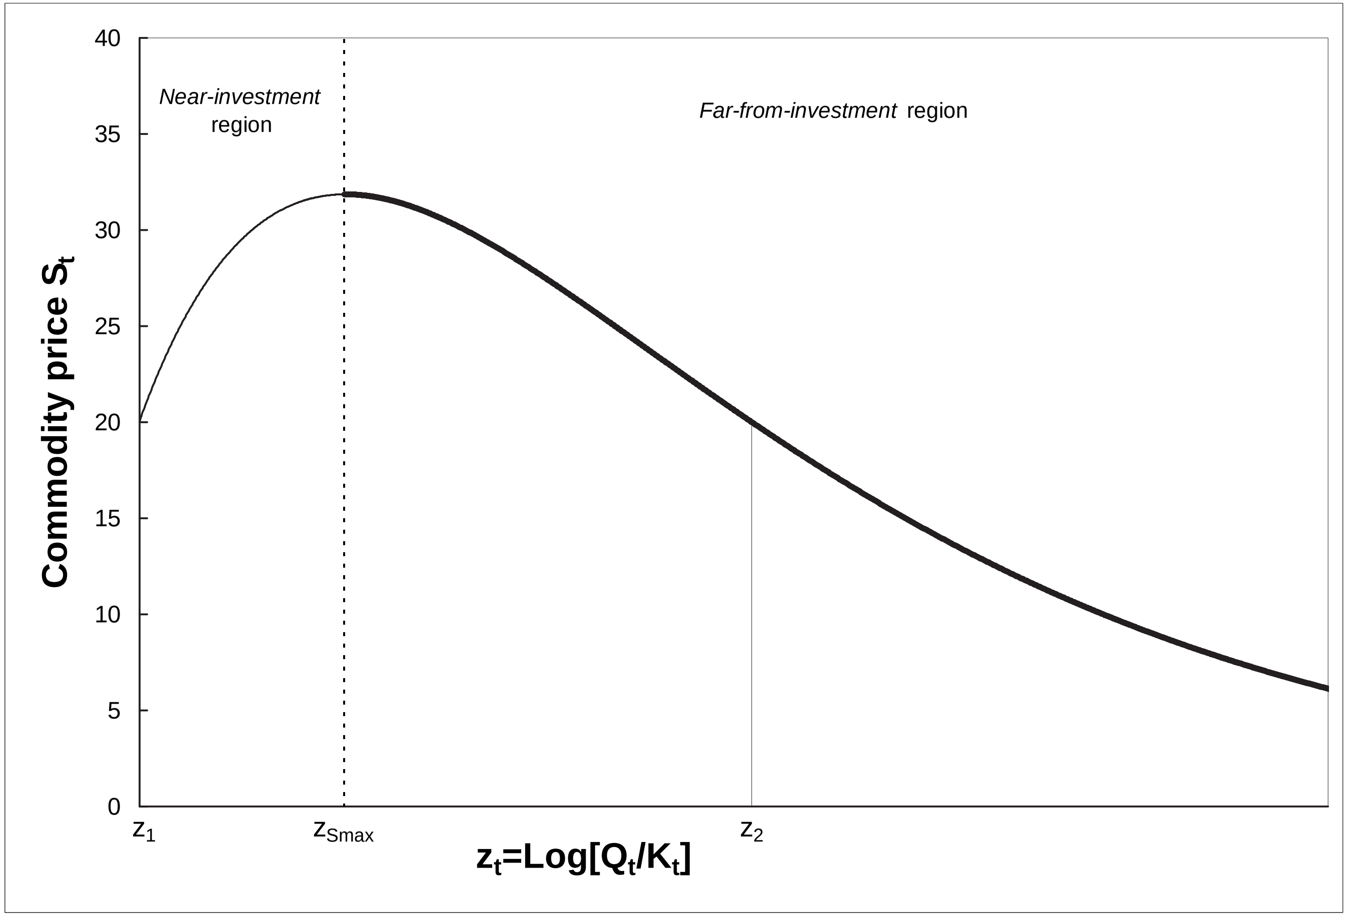

Figure 4: Oil price S_t as a function of the logarithm of the oil wells-capital ratio, z_t. The vertical dashed-line is at z_{Smax} and separates two regions. The thin line shows the oil price in the *near-investment* region ($z_1 < z_t \leq z_{Smax}$) and the thick line is the oil price in the *far-from-investment* region ($z_t \geq z_{Smax}$). We use the parameters in Table 1. In particular, the fixed cost components of the investment are $\beta_0 = 0.005$ and $\beta_1 = 0.01$, and the marginal cost of oil is $\beta_2 = 20$. The equilibrium critical ratios are $z_1 = -8.35$, $z_{Smax} = -7.86$ and $z_2 = -6.88$.

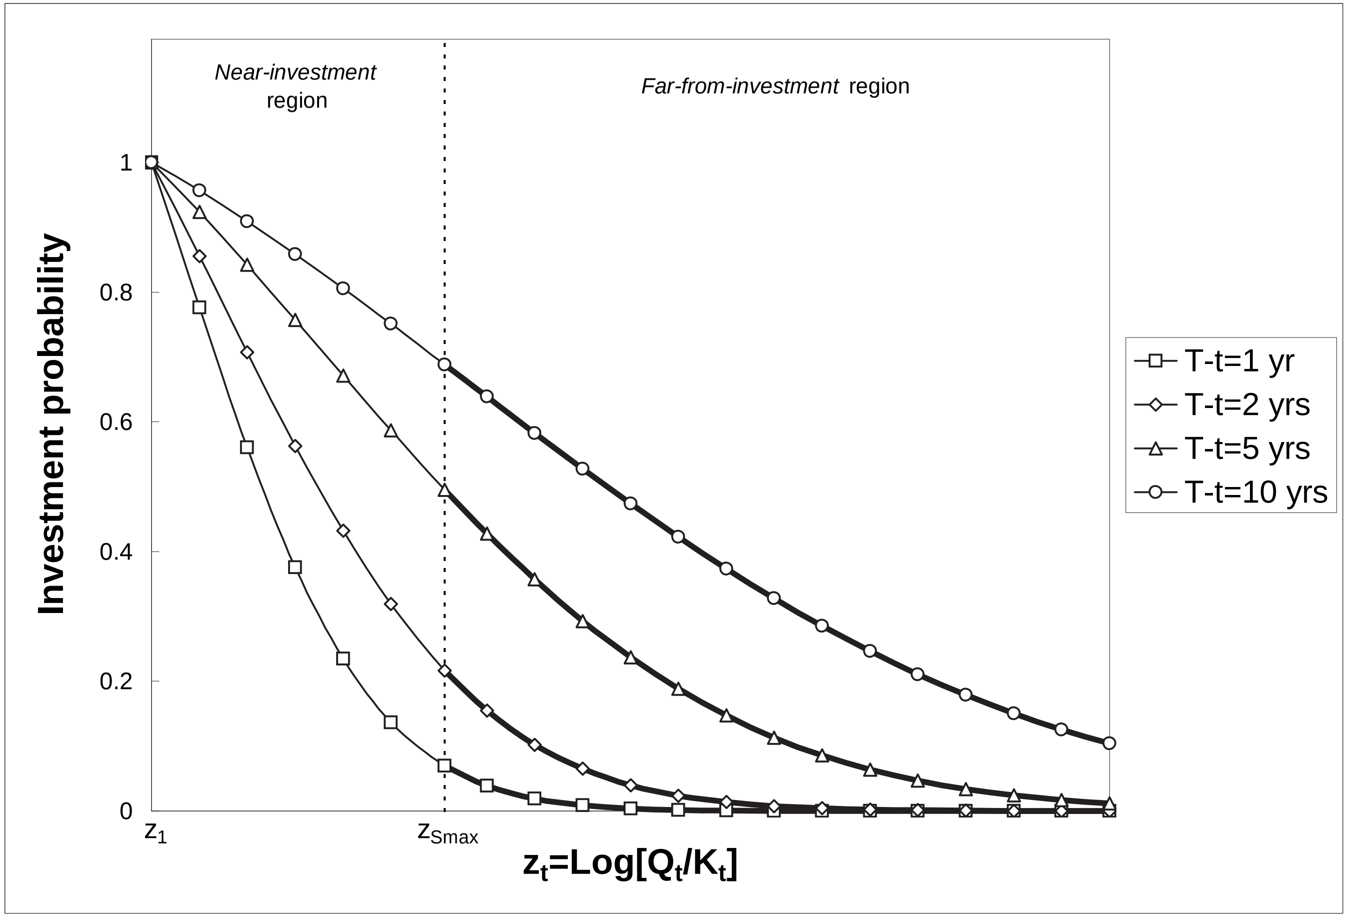

Figure 5: Probability of investing during an interval of time $T - t$ as a function of the logarithm of the oil wells-capital ratio. The thin lines show the probability in the *near-investment* region and the thick line in the *far-from-investment* region. The equilibrium critical ratios are $z_1 = -8.35$ and $z_{Smax} = -7.86$.

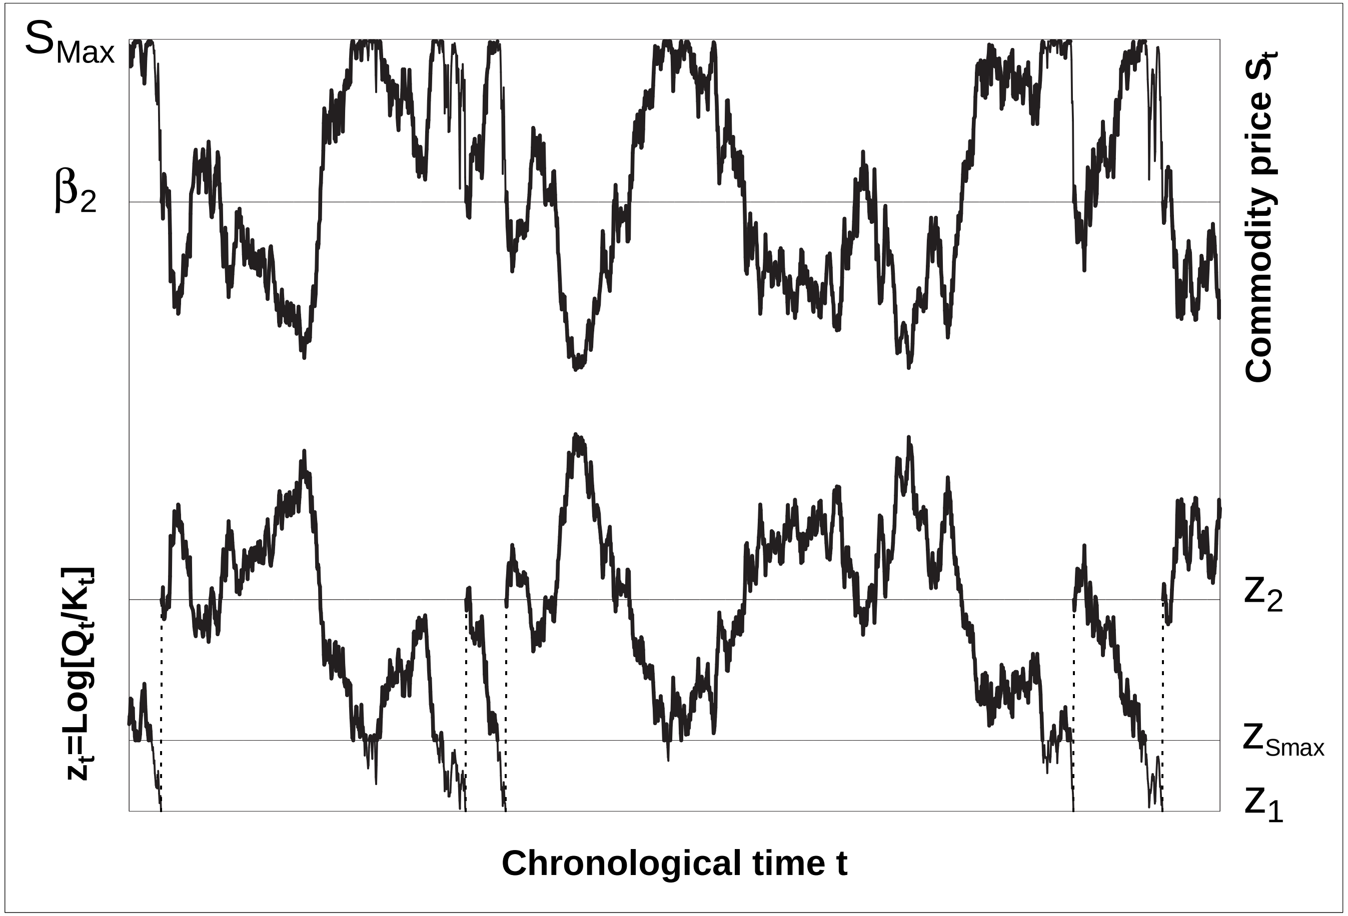

Figure 6: Simulations for the logarithm of the oil wells-capital ratio z_t (below) and the oil price S_t (above) over time. The thin lines show these variables in the *near-investment* region and the thick lines show them in the *far-from-investment* region. We use the parameters in Table 1. For the path below, the critical ratios are $z_1 = -8.35$, $z_{Smax} = -7.86$ and $z_2 = -6.88$, while for path above the (equilibrium) maximum price is $S_{Max} = 31.86$ and the marginal cost of oil is $\beta_2 = 20$.

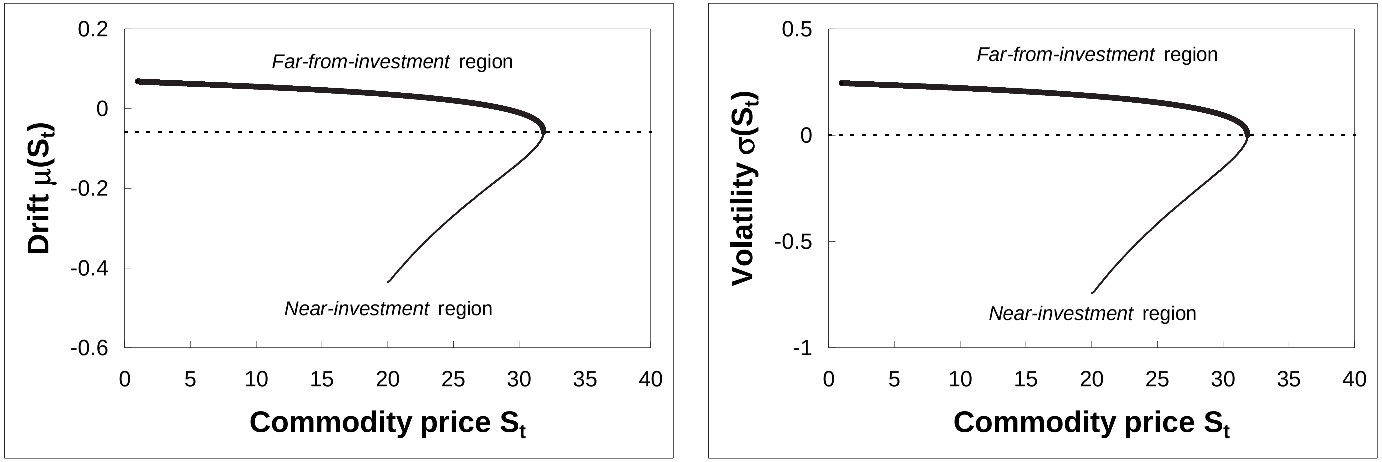

Figure 7: Return and instantaneous volatility of returns in oil price S_t. The horizontal dashed-line separates the two regimes. The thin lines below the dashed-lines show the variables under the *near-investment* regime and the thick lines under the *far-from-investment* regime. We use the parameters in Table 1. In particular, the fixed cost components of the investment are $\beta_0 = 0.005$ and $\beta_1 = 0.01$, and the marginal cost of oil is $\beta_2 = 20$. The endogenous upper bound for the price is $S_{Max} = 31.86$.

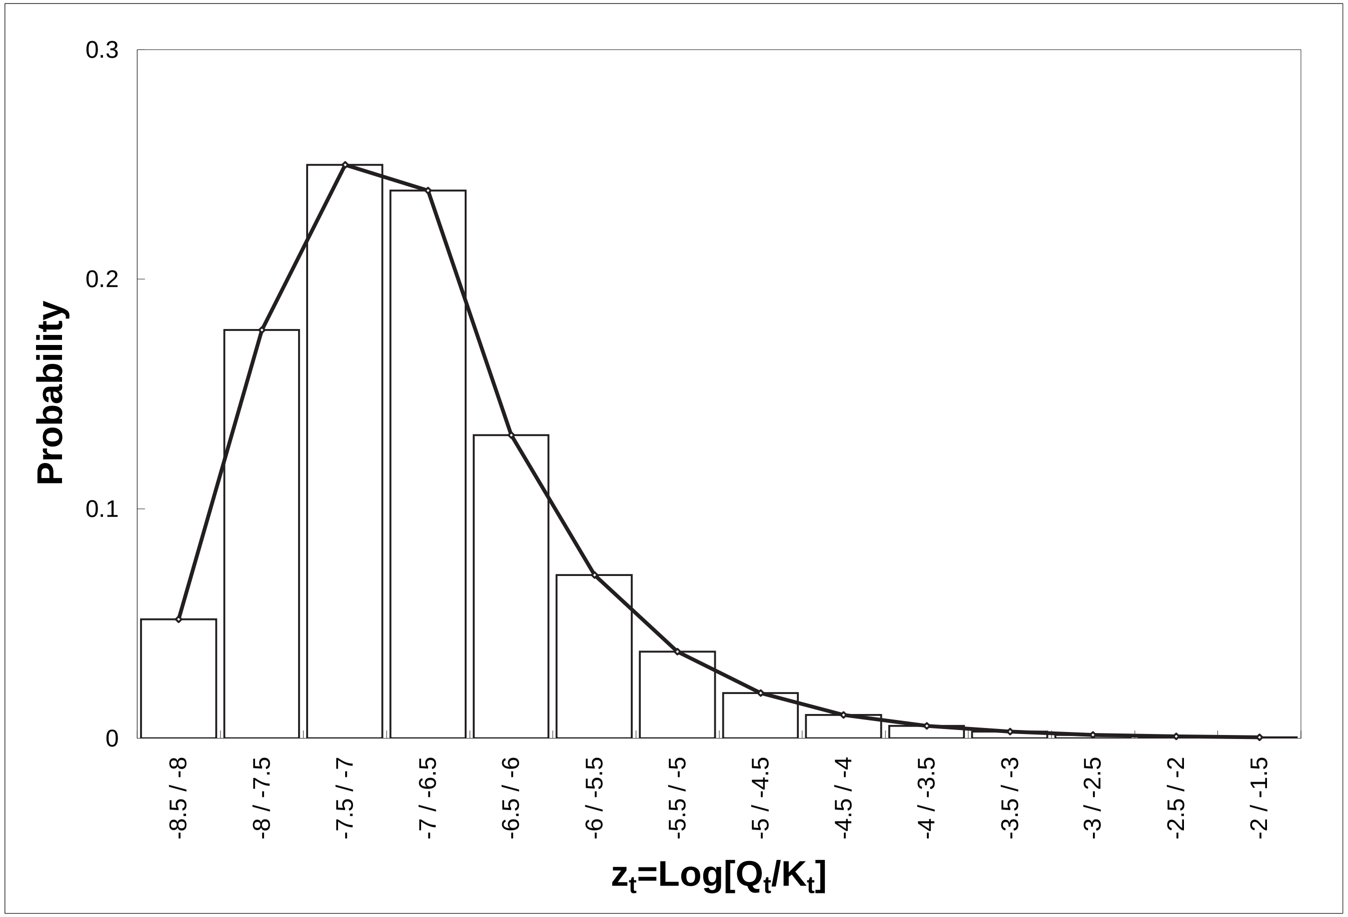

Figure 8: Simulated probability density function (PDF) for the logarithm of the oil wells-capital ratio z_t. We use the parameters in Table 1 and the critical values for z_t are $z_1 = -8.35$, $z_{Smax} = -7.86$ and $z_2 = -6.88$.

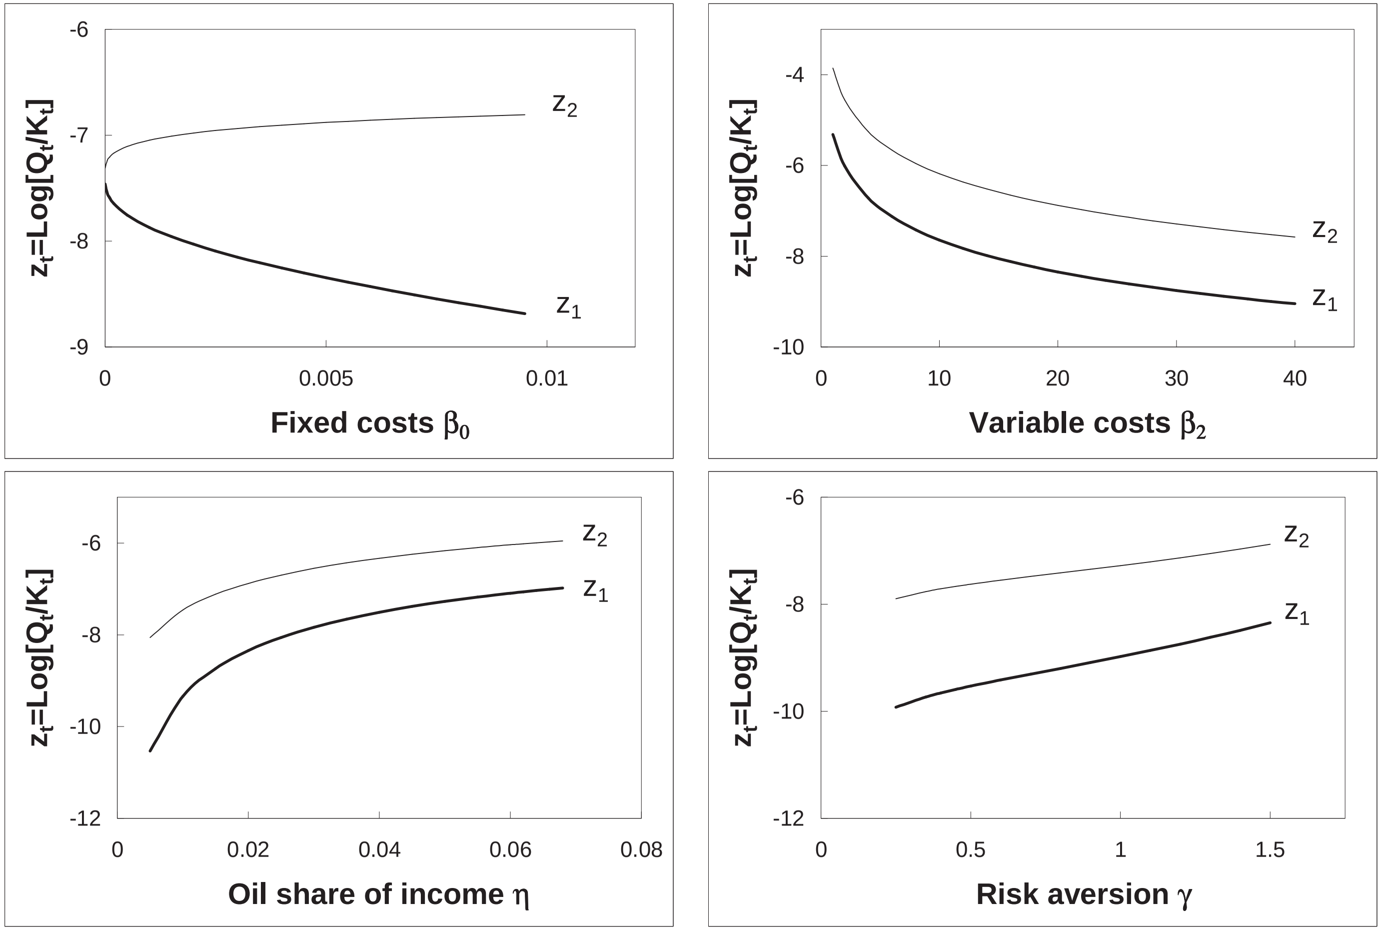

Figure 9: Investment strategy $\{z_1, z_2\}$ for different investment cost structure (β_0, β_2), levels of risk aversion γ and oil share of income η. The thick (below) line corresponds to the investment trigger z_1, while the thin (above) line is the returning point z_2. To summarize both fixed cost components in the parameter β_0, we assume that $\beta_1 = \beta_0\beta_2$ for this plots.

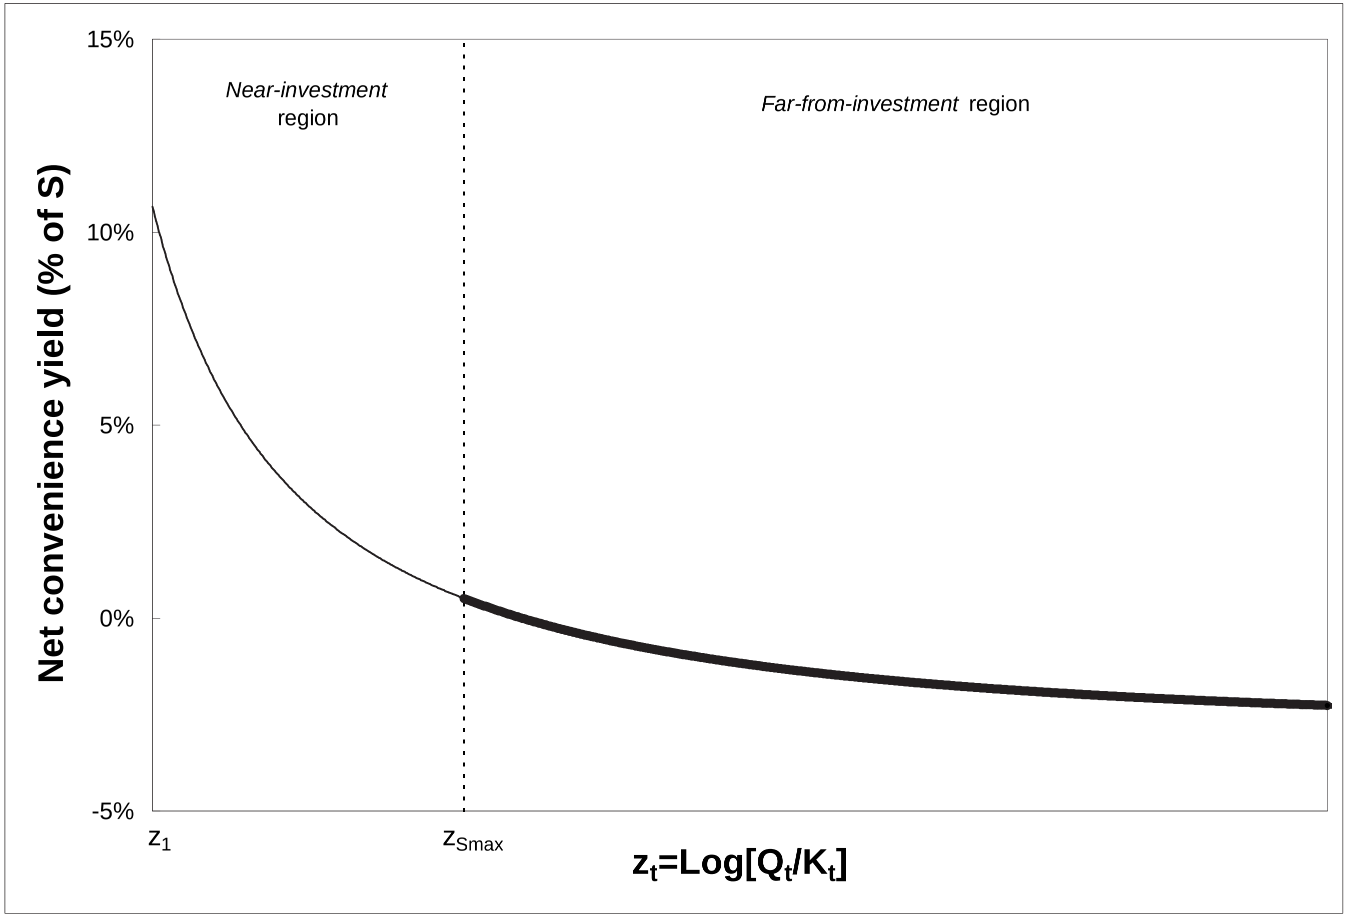

Figure 10: Convenience yield as a function of the state variable z_t when $i_t = \bar{i}$. The thick line is the convenience yield when the economy is in the *far-from-investment* region and the thin for the economy in the *near-investment* region. We use the parameters in Table 1 and the critical values for z_t are $z_1 = -8.35$ and $z_{Smax} = -7.86$.

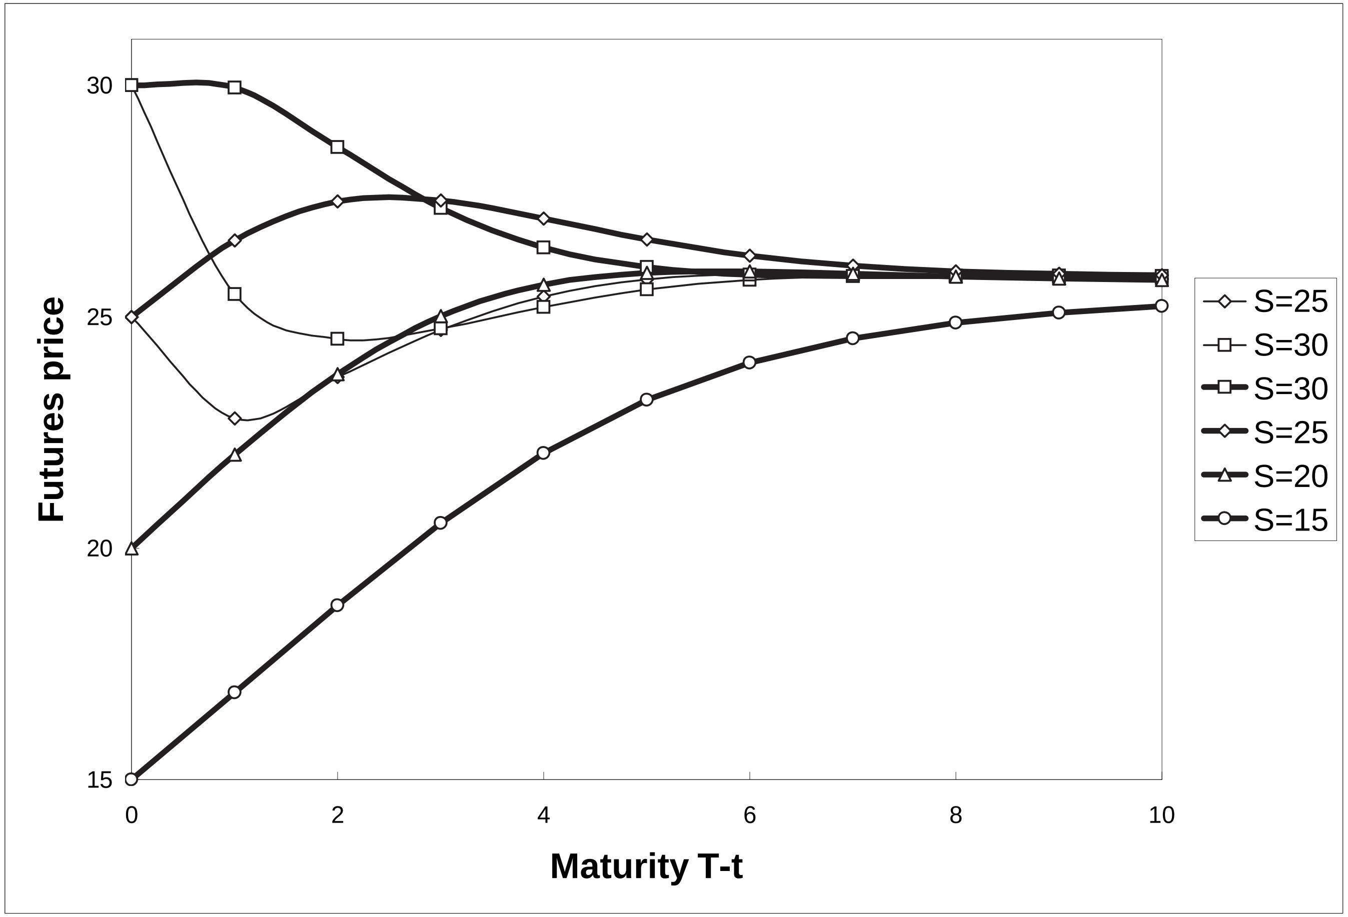

Figure 11: Futures curves for contracts on oil for different spot prices. The thick curves are for spot prices in the *far-from-investment* region and the thin lines when the spot price is in the *near-investment* region. We use the parameters in Table 1 and the endogenous upper bound for the price is $S_{Max} = 31.86$.

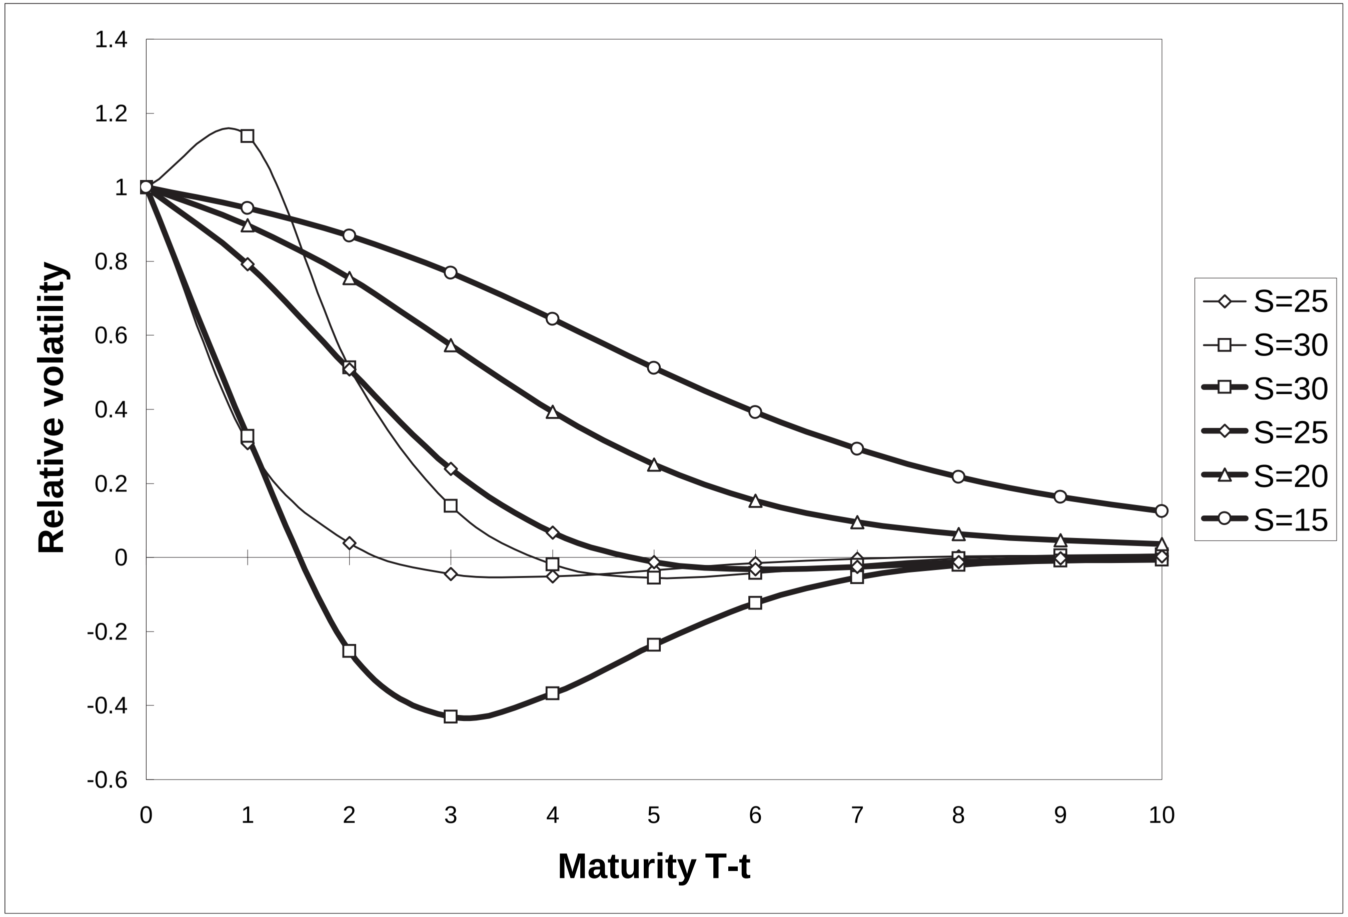

Figure 12: Relative volatility of futures contracts on oil to spot price volatility for different spot prices and maturities. The thick curves are for spot prices in the *far-from-investment* region and the thin lines when the spot price is in the *near-investment* region. We use the parameters in Table 1 and the endogenous upper bound for the price is $S_{Max} = 31.86$.

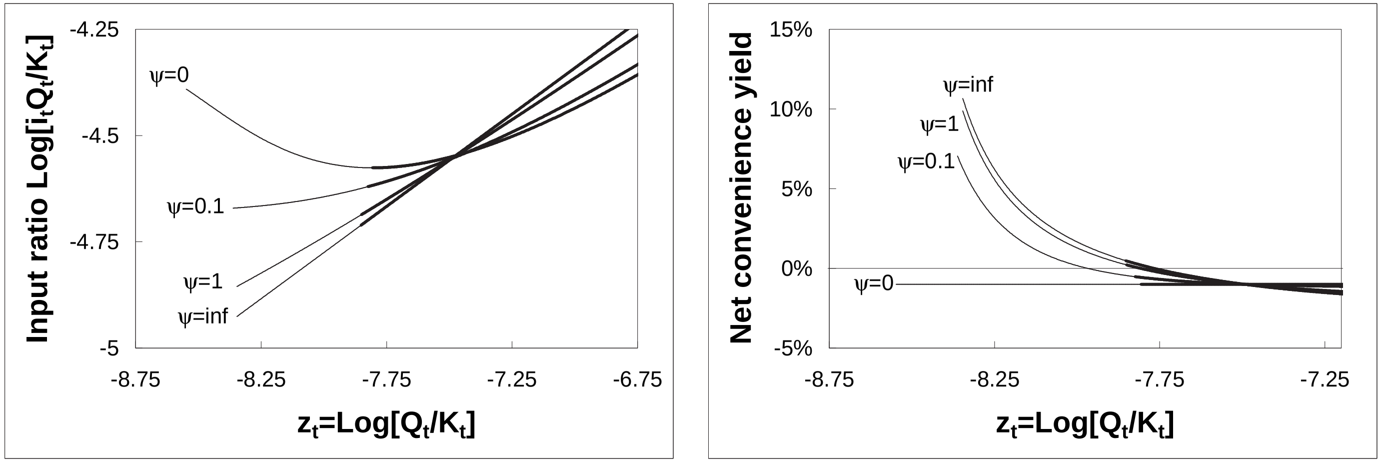

Figure 13: Logarithm of the input ratio for the numeraire production technology and net convenience yield as a function of z_t when the production is flexible and for different adjustment costs. The thin lines show these variables in the *near-investment* region and the thick lines show them in the *far-from-investment* region. We use the parameters in Table 1. In particular, the target input ratio is $\bar{i} = 5\%$ and the depreciation rate is $\delta = 1\%$.

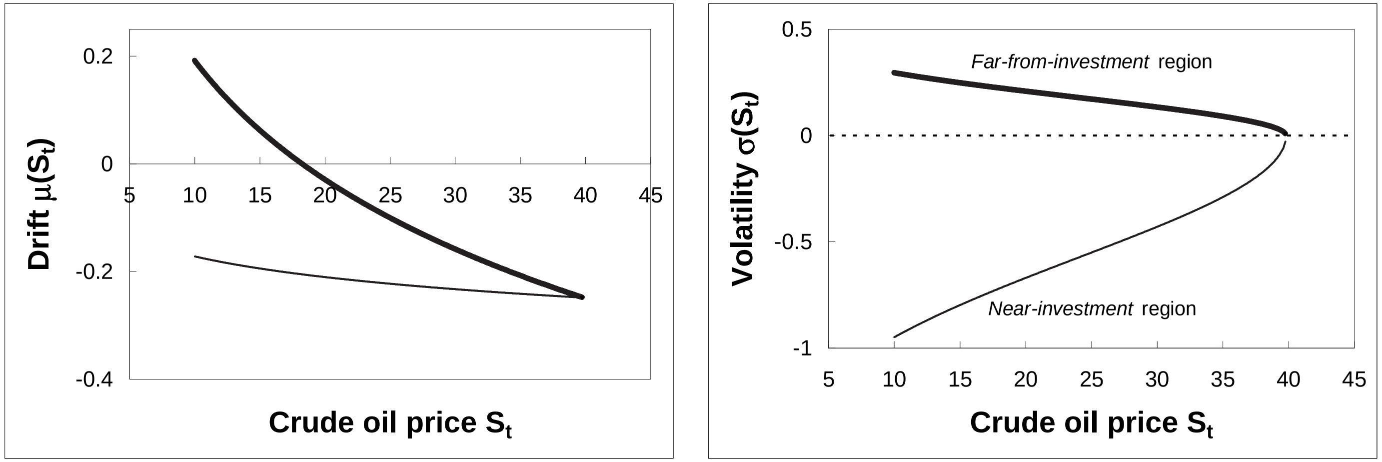

Figure 14: Maximum likelihood estimes for the expected return and instantaneous volatility of returns in oil prices between Jan-1982 and Aug-2003. The thin lines below the dashed-lines show the variables under the *near-investment* regime and the thick lines are the variables under the *far-from-investment* regime. The estimates for the parameters that generate these plots are shown in Table 2.

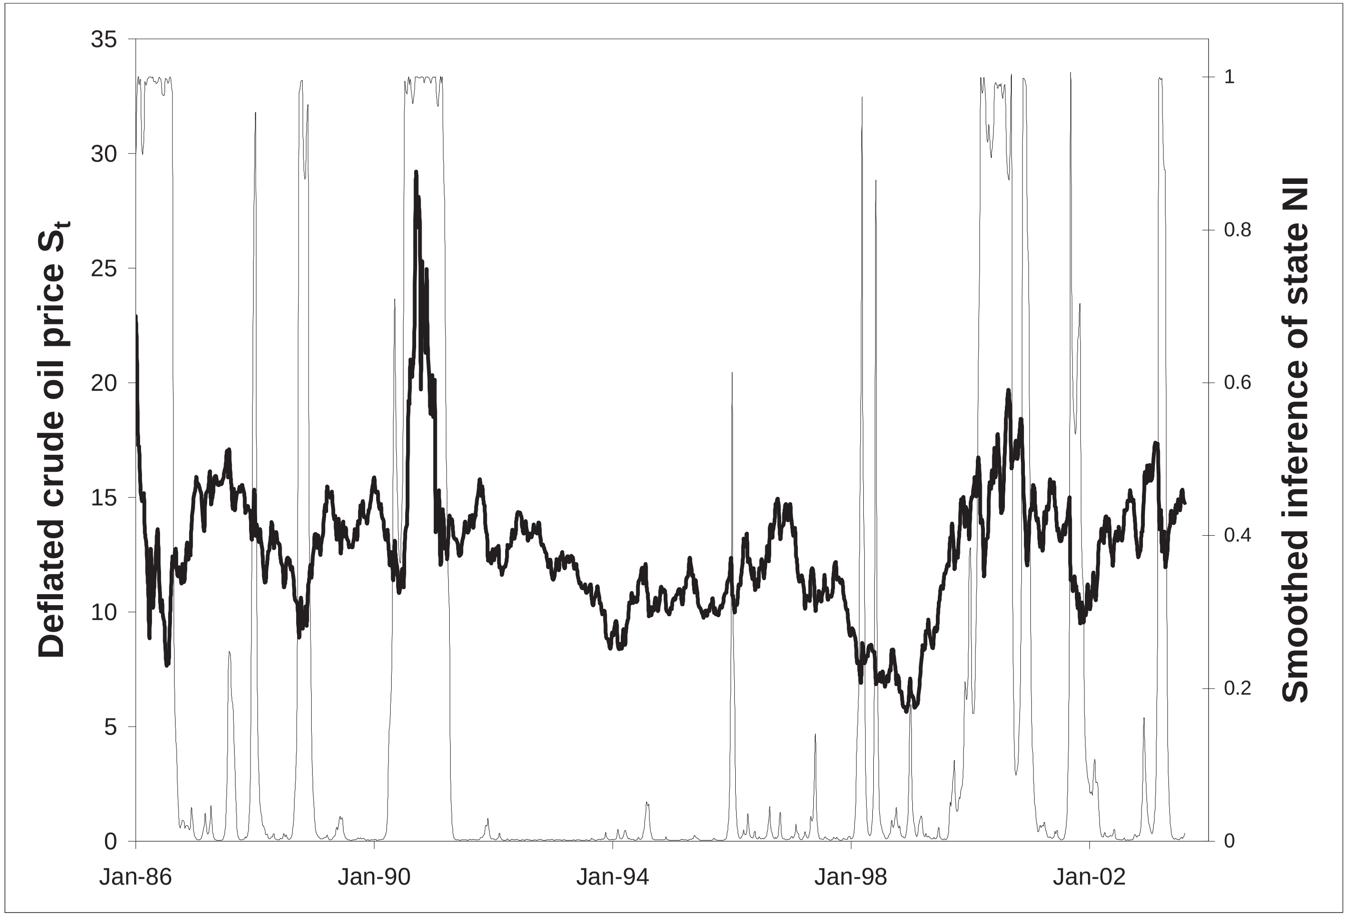

Figure 15: Historical Brent crude oil prices between Jan-1982 and Aug-2003 deflated by the US Consumer Price Index (thick line) and inferred probability of being in the *near-investment* state (thin line).