An Equilibrium Analysis of Exhaustible Resource Investments

Murray Carlson   Zeigham Khokher

Sheridan Titman*

Finance Department

Graduate School of Business

The University of Texas at Austin

Austin, Texas 78712-1179

(Work in Progress)

June 27, 2000

Abstract. We develop a general equilibrium model of an extractable resource market where both the prices and extraction choices are determined endogenously. The model generates price dynamics that are roughly consistent with observed oil and gas forward and option prices as well as with the two-factor price processes that were calibrated in Schwartz (1997). However, the subtle differences between the endogenous price process determined within our general equilibrium model and the exogenous processes considered in earlier papers can generate significant differences in both financial and real option values.

1. Introduction

Contingent claims analysis is currently being used extensively in the energy industry. For example, energy traders often use models suggested by Black [1], Brennan and Schwartz [2], Schwartz [13] and others for risk management as well as for valuing financial contracts

*Phone: (512) 471-5682 (Carlson), (512) 471-1676 (Khokher) and (512) 232-2782 (Titman). Email: carlsonm@mail.utexas.edu, zkhokher@mail.utexas.edu and titman@mail.utexas.edu.

and real investments. These applications, which typically calibrate the models' parameters using some combination of historical prices and observed forward and option prices, have proven to be successful in valuing and hedging relatively short-term financial contracts.

There is, however, an inherent inconsistency in the application of these models that is likely to create a problem when the models are applied to value and hedge longer horizon investments. Specifically, although the models assume the parameters in the price process are constant, the calibration procedures that are used in practice typically provide for a more flexible specification by allowing the parameters to change with time. Although these procedures generally provide reasonably good approximations when the models are used to interpolate among prices in liquid markets, as we will show, they can generate biases when the methodology is used to extrapolate from observed derivative prices to value long term real investments like the pipelines and other infrastructure needed to exploit oil and gas reserves.

To explore these issues in more detail we develop a general equilibrium model of an extractable resource market where both the prices and extraction choices are determined endogenously. As we show, with plausible parameters the model generates prices that are roughly consistent with observed forward and option prices as well as with the price processes that were calibrated in Schwartz [13]. However, the subtle differences between the endogenous price process determined within our general equilibrium model and the exogenous processes considered in earlier papers can generate significant differences in both financial and real option values.

The fundamental sources of uncertainty in our model arise because of fluctuations in aggregate demand and changes in technology. Aggregate demand, or equivalently the growth rate in GNP, is assumed to follow a mean reverting process while changes in technology, which affect the prices of a potential future substitute for the commodity, fluctuates randomly.[1] As our analysis illustrates, price responses to both sources of uncertainty are determined in part by endogenously determined supply responses. For example, temporary demand shocks have little effect on prices when producers can costlessly increase or decrease supply. Conversely, current prices will fail to respond to shocks that affect the cost of the future substitute when the costs of altering current production are sufficiently high. Hence, for the equilibrium price process to demonstrate the long-term and short-term effects observed in the historical data, it is necessary to consider a setting where producers can alter production at a cost that is significant but not prohibitive.

Our model extends existing general equilibrium models that have appeared in both the finance and economics literature. The model is particularly close in spirit to the Pindyck [9] model, which adds uncertainty to the seminal Hotelling [5] model that describes how the prices of exhaustible resources evolve through time. It is also related to the more recent work of Litzenberger and Rabinowitz [8], who argue that because the option to wait has value in an uncertain environment, resources will be extracted more slowly and prices will appreciate less rapidly than they would in the Hotelling certainty model. In contrast to the Pindyck [9] and Litzenberger and Rabinowitz [8] models, the endogenous price process that arises in our model exhibits mean reversion, which is consistent with the empirical data discussed by Schwartz [13] and others[2]. Moreover, our model is consistent with the

[1]For example, in an application to oil prices one might consider the substitute as tar sands, which cannot be profitably extracted at current prices but are likely to be exploited at future dates when the supply of conventional reserves are exhausted.

[2]In our model, as in the Schwartz model, the volatility of futures prices decreases with the term to maturity, indicating the presence of short-run and long-run components in the price process. This phenomenon is sometimes referred to as the Samuelson effect [12]. As we point out below, our model will differ from the Schwartz model in the exact specification of the price dynamics. Other papers have examined how inventory effects the level of mean reversion in exogenous supply shocks (Deaton and Laroque[3], Routledge, Seppi and Spatt[11]). As will be clear later, adding inventory to our model is not conceptually difficult but would be computationally intensive. We suspect that if we were to add storage to our model it would reduce the effects of mean reversion currently generated by our model.

observation that discounted futures prices may be both above and below the current spot price (i.e. futures curves can be in weak contango or backwardation)[3]. As will become clear, these results are not simply due to the stochastic nature of the exogenous state variables but arise endogenously from the assumed frictions associated with the supply responses.

Our model generates insights about the evolution of natural resource prices that can potentially have important implications on the valuation and hedging of long dated financial or real options. In particular, although the endogenous price process generated by our model is qualitatively similar to the price process assumed by Schwartz [13], the functional form of the drift is, in general, non-linear and generates equilibrium price paths with less extreme realizations than would be generated by Schwartz's model. As a result, options, whose payoffs are especially sensitive to these extreme realizations, are generally less valuable in our general equilibrium setting where the extreme realizations are observed less frequently.

The format of the paper is as follows. In the next section, we specify the assumptions of the model and define the equilibrium. In Section 3 we present an example which is useful for developing intuition regarding the economics underlying our full model. Implications of the equilibrium model for futures prices, future price volatilities and production decision are presented in Section 4. Finally Section 5, compares option prices from our model to those of Schwartz [13].

[3]If futures prices are below the current spot price, the futures curve is said to be backwardated. Litzenberger and Rabinowitz [8] make the distinction between weak and strong backwardation. If discounted futures prices are below the spot price, they say the futures curve is weakly backwardated. Contango is the opposite of backwardation.

2. Assumptions and Equilibrium

In this section we present the assumptions underlying the model and the definition of the equilibrium. A brief summary of the overall setup will motivate the rationale behind the detailed assumptions which follow.

The model examines a risk-neutral economy with a finite reserve of a commodity that is owned by each of a continuum of small, potentially heterogeneous producers. Producers optimally extract the commodity in the face of uncertainty regarding the economy wide demand. In addition, there is an alternative source of supply whose marginal extraction cost is known and stochastic. If producers increase production rates beyond what they have been producing in the recent past, they incur a cost that is proportional to the difference between their new production rate and the lagged production rate. This cost is meant to capture the costs associated with developing reserves.

2.1. Reserves. The economy is defined in continuous time with an infinite horizon. Instantaneous borrowing and lending is possible at a constant interest rate r. There is a finite reserve of a commodity, R_0, owned by a continuum of price-taking producers and an inexhaustible supply of a substitute good. The cost of extraction is assumed to be constant across time, but may differ by producer. In equilibrium low cost producers extract their reserves first, so the unit cost of extraction may be of an arbitrary form, $C(R_t)$, but will increase monotonically as reserves are depleted.[4]

The dynamics of the reserve process, which defines how the reserves are depleted over time, can be expressed as:

$$dR_t = -q_t dt \tag{1}$$

[4]Pindyck[9] uses this specification of reserves in his model.

where q_t is the production process and $R(0) = R_0$. Note that there is no exogenous uncertainty in this process. However, the reserves process is random whenever the production process is stochastic.

The production process is defined only so long as reserves exist. Given a production policy, the time to exhaustion of the reserves is defined by the following stopping time:

$$R_0 = \int_0^\tau q_t dt. \tag{2}$$

The planning horizon defined by this stopping time may or may not be finite.

2.2. Uncertainty. The (inverse) demand function for the commodity is assumed to be of the form, $p_t = f(q_t; y_t)$. The parameter y_t characterizes inter-temporal demand shocks that arrive according to the process:

$$\frac{dy}{y} = \alpha_y(y)dt + \sigma_y(y)dz_y \tag{3}$$

We will be focusing on the case where this process is mean-reverting with a constant diffusion, so that $\alpha(y) = \kappa_y(\mu_y - \ln(y))$ and $\sigma(y) = \sigma_y$.

We assume that a substitute for the commodity exists with effectively infinite reserves. The substitute is not currently produced because of excessive marginal extraction costs, S_t. However, technological innovations arrive stochastically and affect this cost:

$$\frac{ds}{s} = \alpha_s(s)dt + \sigma_s(s)dz_s. \tag{4}$$

We focus on the case where this process is a geometric brownian motion with constant drift, $\alpha_s(s) = \mu_s$.

The substitute commodity essentially caps demand at its marginal cost. Thus, the effective market demand function is of the form:

$$p(q_t; y_t, S_t) = \min(s_t, \frac{y_t}{q_t}) \tag{5}$$

where q_t is the current amount produced from conventional reserves.

2.3. Definition of Equilibrium. Producers, who are assumed to be price-takers, make production decisions that maximize the market value of their reserves, net of the expected costs of extraction. Note that, since the market value of reserves is a function of the equilibrium price, optimal production decisions and market clearing prices must be determined simultaneously. In equilibrium, at each point in time and in each state, producers correctly conjecture the future evolution of prices and incorporate this information into their production decision.

In addition to marginal extraction costs that depend on the level of reserves, $C(R)$, we assume that producers incur a cost whenever production rates increase. Although the study of more general setup costs is possible, we assume that this cost is proportional to the magnitude of the increase of the optimal production over the existing production rate:

$$f(q_t; q_{t-}) = \gamma(q_t - q_{t-})^+ \tag{6}$$

where γ is a constant, q_t and q_{t-} are the chosen and the existing production rates respectively. As mentioned at the beginning of this section, the form of this cost function is meant to capture the cost of developing new reserves in a reduced form.

To solve for equilibrium prices we consider the dual problem of a Social Planner who maximizes discounted expected consumer surplus in excess of producer surplus. More specifically, at a given point in time this social surplus, SS, is defined as:

$$SS(q_t; y_t, s_t, R_t, q_{t-}) = \int_0^{q_t} p(q; y, s)dq - C(R)q_t - f(q_t; q_{t-}) \tag{7}$$

The social planner chooses production rates to maximize the discounted expected value of the following expression:

$$V(R_t, y_t, s_t, q_{t-}) = \max_{q_t} E_t \int_t^{\infty} e^{-r(s-t)} SS(q_s; y_s, s_s, R_s, q_{s-})ds \tag{8}$$

subject to

$$\int_t^{\tau} q_s ds = R \quad \text{a.s} \tag{9}$$

where τ is a stopping time indicating the date at which reserves are fully depleted. Under conditions outlined in Dixit and Pindyck [4], the solution to this problem will coincide with production policies generated within a competitive equilibrium. The advantage of casting the problem in terms of maximizing social welfare is that traditional dynamic programming techniques can be applied to solve the problem numerically. Once optimal production policies are determined, equilibrium prices are determined, state-by-state, by the market clearing condition implied by the demand curve.

3. A Simple Example

This section considers a simplified version of our model that can be solved analytically. The intuition developed from this example is helpful for understanding the more general model, which must be solved numerically.

We consider the following simplified demand process, which makes closed form solutions possible:

$$p_t(q_t) = \frac{y + \epsilon_t}{q_t} \tag{10}$$

where q_t is the amount produced, y is a constant and the ϵ_t are positive IID shocks. Clearly, demand shocks are temporary in this setting and we can interpret this sequence of demands as being the limiting case for the class of mean-reverting shocks.

The timing of the information and decisions is as follows. At the beginning of each decision epoch, t, the current level of reserves is known to be R_t. Producers observe a shock to the demand curve ϵ_t and make their optimal production decisions. The resulting market clearing price is given by $p_t = p_t(q_t)$. Immediately after the production decisions have been made, the level of reserves drops to $R_{t+1} = R_t - q_t$.

We first provide, a closed-form solution for the case in which there are no extraction costs, no costs associated with altering production rates and no substitute commodity. We then characterize the solution in a more general setting that includes setup costs. Although we cannot provide a closed form solution for the equilibrium in this latter case, we describe the form of the optionality introduced by the setup costs and show how it modifies the optimal response to demand and supply shocks. Our results illustrate how, in the absence of setup costs, a mean-reverting state variable generates prices that are random walks. Hence, our results suggest that setup costs are a necessary feature of a model where prices have temporary as well as permanent components.

3.1. The Equilibrium without Setup costs. In this section we solve for the equilibrium in a simple case and analyze its properties. Recall that this simple case does not consider extraction costs, costs for altering production rates or a substitute commodity. We solve for the equilibrium by reformulating the dynamic optimization problem as the following static variational one:

$$\max_{q_t} E \sum_{t=0}^{\infty} SS\left(q_t\right) \tag{11}$$

subject to

$$\sum_{t=0}^{\infty} q_t = R_0 \quad \text{a.s.} \tag{12}$$

Proposition 1 characterizes the equilibrium price dynamics in this simplified case.

Proposition 1. *Discounted prices in a competitive equilibrium are martingales. Thus, for $s > t$*

$$e^{-rt} p_t = E_t(e^{-r(s-t)} p_s). \tag{13}$$

Moreover, the price of the commodity at an arbitrary time is a function of two random state variables, ϵ_t and R_t:

$$p_t = \frac{ay + \epsilon_t}{R_t} \tag{14}$$

where $a = \frac{1+r}{r}$.

Notice that the above implies that the discounted expected value of the future spot price is the current spot price. Thus, at every point in time prices are expected to rise at the riskless interest rate (i.e. $p_t = E_t\left[e^{-r(s-t)} p_s\right]$ for $s > t$), which suggests that a natural extension of the standard Hotelling [5] result holds in this example. This result is also noted in Pindyck [9] as would be expected given that our example is a special case of his model. (The simplifications allows us to examine the behavior of the volatility of the price process.)

Note also that as long as we are in a setting where reserves are never exhausted, which is the case for this example, the forward curve will be defined by[5]:

$$f_{t,s} \equiv E_t(p_s) = e^{r(s-t)} p_t. \tag{15}$$

This is consistent with a result in on the instantaneous drift of the resource price. Uncertainty alone does not necessarily create the backwardation result in Litzenberger and Rabinowitz [8].

Another interesting consequence of the supply responses is that they turn temporary demand shocks into permanent price shocks. We can see this using the fact that the spot price and forward prices are related by Equation (15) and thus shocks to next period's spot price are attenuated and transmitted to all forward prices. More precisely, since one step ahead forward prices are directly proportional to next period's spot price, shocks to the spot price are transmitted throughout the entire forward curve.

With this solution at hand, it is easy to characterize the variance of both spot and forward prices. These results are recorded in Proposition 2.

Proposition 2. *At any point in time the conditional variance of next period's spot price is given by:*

$$Var_t(p_{t+1}) = \frac{Var_t(\epsilon_{t+1})}{R_{t+1}^2}. \tag{16}$$

and we can calculate the variance of the logarithm of the future spot price as:

$$Var_t\left(\log p_{t+s}\right) = A + (s-t)\sigma_\eta^2 \tag{17}$$

[5]In a setting where reserves may be depleted in finite time the instantaneous change in discounted prices will be a martingale until the reserves are depleted Routledge, Seppi and Spatt [11] carefully analyze forward curves in a related setting where inventories may be exhausted. In this case there is no closed form solution for the forward curve.

where A and σ_η are constants.

Remember that $R_{t+1} = R_t - q_t$ is in the information set at time t. Thus, the first part of the proposition makes it clear that the effect of a demand shock is greatly attenuated by supply responses. To see this, consider what would happen in the following period were producers not to alter their production from the current level. In this case, the variance of the next period price would be $Var(\epsilon_{t+1})/q_t^2$ which is clearly higher since current production is much lower than the total remaining reserves.

In addition to their dampening effect, Proposition 2 illustrates a second implication of the producer's supply responses. The variance of the log of the future spot price is linear in the holding period. This fact has two empirical implications. First, variance ratio tests should indicate that log prices follow a random walk. Second, the implied volatilities of options on this commodity's forward prices should be constant. The latter implication illustrates that the endogenous supply responses transform a mean reverting state variable into a random walk. It is also interesting to note that at long horizons the Black [1] model of option pricing should become exact, since at long horizons the spot price is lognormally distributed.

3.2. The Behavior of Prices with Infinite Setup Costs. When setup costs are infinite, it can never be optimal to increase production.[6] It is instructive to consider the behavior of prices under a deterministic policy where the quantity produced decreases at the rate of interest.

Proposition 3. *Consider the following deterministic production policy:*

$$q_t = e^{-rt} q_0$$

[6]In order to focus on a non-trivial case, we assume that initial production can be chosen costlessly.

where q_0 is chosen so that $R_0 = q_0 \sum_{t=0}^{\infty} e^{-rt}$. Under this policy, the forward curve slopes upward at the rate of interest:

$$E(p_t) = e^{rt} \frac{y}{q_0}$$

for all $t > 0$. Moreover, spot prices may be temporarily above or below $\frac{Y}{q_t}$ implying that forward curves may be in contango or backwardation. Finally, the term structure of volatility is declining and constant.

The proof follows immediately from the definitions of forward prices and the term structure of volatility. We will see in the next section that this case, which is the polar opposite of the case considered in the previous subsection, provides an approximate description of prices when setup costs are high.

3.3. The Equilibrium with Setup Costs. It is clear from the simple case discussed above, that if equilibrium prices are to have any temporary components, frictions must be introduced into the model. To induce temporary components, we add a cost associated with increasing production rates. The modified objective function, which includes a cost proportional to the change in the production rate that is incurred is given below:

$$\max_{q_t} E\left[\sum_{t=0}^{\infty} SS(q_t) - \sum_{t=0}^{\infty} \gamma(q_t - q_{t-1})^+\right] \tag{18}$$

subject to

$$\sum_{t=0}^{\infty} q_t = R_0 \quad \text{a.s}$$

The additional cost greatly complicates the analysis of the problem making a closed form solution impossible. We will focus here only on the form of the optimal production policy and indicate why closed form solutions in this case are not possible.

Proposition 4. *At any time t.the optimal production policy, q_t,must satisfy one of two first order conditions*

$$e^{-rt}\frac{y+\varepsilon_t}{q_t}+e^{-rt}B_t=k_t \tag{19}$$

$$e^{-rt}\frac{y+\varepsilon_t}{q_t}-\gamma+e^{-rt}B_t=k_t \tag{20}$$

here B_t is a binary option with a payoff γe^{-r}if $q_{t+1}>q_t$ and k_t is a constant. Moreover, there is a range of demand shocks where there will be no supply response.

The first order conditions to this problem depend on whether or not producers increase their production rates in the current period. First, if the current demand shock is low, so that producers will want to decrease production, Equation (19) must be satisfied. This first-order condition is in effect if the current decision is to decrease the rate of production in the current period. The value of this option is a function of the current production choice and if the current production rate is decreased its value increases. Therefore a decrease in production will cause an increase in the left hand side of the Equation.

On the other hand, if the current demand shock is high then producers will want to increase production which implies that Equation (20) must be satisfied. In this case we need to consider the effect of an increase in the production rate on the binary option whose price is B_t. As current production increases, its value falls. In addition, production increases cause the current spot price to fall. Thus, increasing production causes the left hand side of the equation to decrease.

As shown in the Appendix, there is an intermediate range of demand shocks for which there will be no supply response. This implication is illustrated in Figure 1 which plots the above first-order conditions. In order to make the figure easier to interpret, the dependence of k_t on the level of the demand shock is ignored.[7] The current production rate is determined by the intersection of the equations for the first-order condition and the current level of k_t. As illustrated, in very high demand states the production rate increases and in very low demand states the production falls.[8] However, there is a range of intermediate shocks for which production will remain unchanged. The size of this region is proportional to γ, the proportional cost of increasing the production rate.

The form of the optimal production policy has important implications for the commodity price process. Notably, if the cost of increasing production is suitably high, commodity prices will inherit some of the temporary nature of the demand shocks. In addition, over long horizons there will be some impact from the endogenous supply responses. Therefore, in this setting we would expect to see both permanent and temporary components in the commodity's price.

There is little more we can say about the optimal solution to the Social Planner's problem in this setting without characterizing the solution to this problem numerically. No analytic solution exists for either the Lagrange multiplier process or for the value of the binary option that appears in the first-order conditions. In the next section we generalize this example, solve the problem numerically and characterize the interesting aspects of the price process.

4. The Numerical Solution to the General Model

We now move to the solution of the more general model introduced in Section 2. As indicated above, in order to proceed with the analysis we must apply computational techniques to solve the model numerically. This section begins with a brief discussion of

[7] It is possible to show that λ is less sensitive to the demand shock than is the level of demand.

[8] i.e. λ_t intersects the "high demand" and "low demand" curves to the right and left, respectively, of the lagged production level.

the solution methodology and then moves on to study the behavior of the model under various parameterizations.

4.1. The Computational Technique. The equilibrium is characterized by the solution to the constrained social planner's problem defined by equation (8). This problem is conceptually straightforward to solve using the standard recursive techniques of dynamic programming. For example, given an initial estimate for the value function in a given state, $V_0(R, y, s, q_-)$, one can apply value iteration techniques in order to converge to the fixed point that describes the solution as well as the production policy associated with the optimum (see, for example, Puterman [10]). Given the optimal production policy, it is then possible to determine equilibrium prices as a function of the state variables, as well as to describe the equilibrium price dynamics, by working with the transition density of the resulting Markov chain.

The problem with solving for the equilibrium arises for practical reasons. Typically, the first step in solving these types of dynamic problems numerically is to form a discrete approximation to the continuous state space (see, for exapmle, Kushner and Dupuis [7]). This gives rise to a problem known in the numerical methods literature as the "Curse of Dimensionality": as the dimensionality of the state space increases, the number of points in the discrete approximation to the state space increases geometrically. The problem we are studying here has four state variables, (R, y, s, q_-), and one continuous choice variable, the production rate. Thus, the computational and storage requirements of the problem are considerable. We deal with this issue by applying numerical algorithms that can efficiently exploit the structure of the problem.

4.2. Two Benchmark Examples. In this subsection we examine the behavior of the numeric solution to the general model under two extreme assumptions about the startup costs. First, we set these costs to zero and analyze the model's output in light of the results from Section 3. Second, we analyze the model when setup costs are very high, in which case results should be similar to those described in Proposition 3. A detailed analysis of these particular parameterizations will illustrate the basic forces underlying the fully specified general equilibrium.

Relevant characteristics of the equilibrium without startup costs are illustrated in Figure 2. In the leftmost column the conditional behavior of future prices is examined. The forward curve is the solid line in the top panel. Consistent with the analytic results in Section 3, forward prices grow from the spot price at the rate of interest. This is true for all levels of the state variables; thus, temporary demand shocks cause parallel shifts in the entire forward curve. Two measures of the volatility of future prices are examined in the bottom panels. The standard deviation of future log prices is proportional to the square root of time and, therefore, the term-structure of volatility is flat.[9] Notice that this occurs in the model despite the fact that demand shocks are temporary and is a direct consequence of the costless supply responses. Also note that supply responses considerably dampen demand shocks, resulting in price volatilities that are an order of magnitude smaller than demand volatility.

Two characteristics of the optimal supply policy are apparent from the figure and are illustrated in the rightmost column. First, average production decreases with time, consistent with the fact that prices are expected to increase. Second, quantities are about

[9]We define the term-structure of volatility for a stochastic process, x_t, as the relationship between $\sqrt{\frac{\mathrm{var}(x_t)}{t}}$ and t.

as volatile as demand shocks, indicating that all changes in demand are matched by changes in the quantity supplied.

With high startup costs the behavior of the model is very different as is illustrated in Figure 3. In contrast to the case just described, the forward curve may be in backwardation or in contango (i.e. the forward prices do not increase at the rate of interest). High realizations of demand are associated with steeply backwardated forward curves as a result of producer's (optimal) reluctance to increase production. Evidence of such reluctance can also be seen in panels (g) and (i) where we see that the volatility of the supply response is low and the volatility of the spot price is high. Note also that the resulting equilibrium spot price volatilities are very similar to those of the demand shock.

The dynamics of the forward curves in the two benchmark examples are compared in Figure 4. First, we choose two distinct points in time, each with an associated forward curve. The relationship between the shape of the curves in the two panels is of interest. In panel (a) we see that the two forward curves are parallel. This illustrates the fact that temporary shocks have a equal effect on all future prices when supply responses are costless. On the other hand, when supply responses are costly, temporary shocks have a larger impact on short-term prices than on long term prices; hence, pairs of forward curves are not necessarily parallel (see panel (c)).

We can further clarify the dynamics of the forward curves if we compare the spot price process to the 2-year forward price process. When startup costs are zero, the forward price process looks very much like the spot price process. In contrast, when startup costs are high the spot price process is considerably more volatile than the forward price process, indicating that prices have a mean reverting tendency. Hence, we see that setup costs are necessary to generate mean reversion in the exhaustible resource price process.

4.3. The Base Case Parameterization. In this subsection we present the solution to the general model under a parameterization that we refer to as the base case. Parameter values are itemized in Table 1. The startup costs for this case are between those in the two polar cases discussed above. Other inputs to the model are unchanged.

Under the base case parameterization, equilibrium price dynamics in our economy are qualitatively similar to those described by the partial equilibrium model of Schwartz and Smith [14][10]. A brief description of their specification will facilitate comparison to our endogenous prices. They assume that the log of the spot price, X_t, is composed of a short-run Ornstein-Uhlenbeck deviation process, χ_t, and a long-run Brownian motion process, ξ_t. More precisely:

$$\begin{aligned} X_t &= \xi_t + \chi_t \\ d\chi_t &= -\kappa \chi_t dt + \sigma_\chi dz_\chi \\ d\xi_t &= \mu_\xi dt + \sigma_\xi dz_\xi \end{aligned}$$

With this specification, forward prices are given by the following equation:[11]

$$\ln(F_{t,t+s}) = \xi_t + e^{-\kappa s}\chi_t + \mu_\xi s + \frac{1}{2}\left[(1 - e^{-2\kappa s})\frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2 s + 2(1 - e^{-\kappa s})\frac{\rho_{\xi\chi}\sigma_\chi\sigma_\xi}{\kappa}\right]$$

If χ_t is positive (negative) forward curves will be in weak backwardation (contango). This effect diminishes exponentially with time and the long end of the forward curve slopes upward at the rate μ_ξ.

[10] Schwartz and Smith show how their model is similar to those of Gibson and Schwartz and Schwartz.
[11] We will ignore the market prices of risk for the two factors as they do not affect the shape of the forward curves. See Schwartz and Smith for the fully specified forward prices.

We will also be interested in their term structure of volatility which is given by:

$$\frac{\mathrm{var}(\ln(p_{t+s}))}{s} = \sigma_\xi^2 + \frac{1-e^{-2\kappa s}}{s}\frac{\sigma_\chi^2}{2\kappa} + 2\frac{1-e^{-\kappa s}}{s}\frac{\rho_{\xi\chi}\sigma_\chi\sigma_\xi}{\kappa} \tag{21}$$

Their term structure of volatility is independent of the level of the state variables and declines over time to a constant level σ_ξ^2.

Like the price processes described in Schwartz and Smith, the endogenous price process generated by our model has both a short-run mean reverting component, and a long-run growth component. Our forward curves may be in backwardation or in contango, depending on the level of the demand shock. In addition, the term structure of volatility is downward sloping.

Observation 1. ***[Forward Curves]*** *The forward curves in the economy can be in backwardation or in contango (see Figure 5).*

The forward curves are in weak backwardation (contango) depending on whether the demand shock process is below (above) its long-run mean. This effect is a direct result of the fact that supply responses are costly.

Observation 2. ***[Term Structure of Volatility]*** *The term structure of volatility is downward sloping (see Figure 5).*

The reason for the increased short run volatility is that current supply responses are constrained and hence exogenous shocks cause increased volatility at the short end of the curve. However, the long end of the curve exhibits lower volatility since the effect of exogenous shocks is dampened by producers supply responses.[12]

[12]It is well known that some markets may have an increasing or humped term structure of volatility.

4.4. Comparative Static Analysis. In this subsection the sensitivity of the model's output to variation in the input parameters is examined. In particular, changes in the level of reserves, the interest rate and volatility of the two sources of uncertainty will be discussed.

We begin by studying the effect of reserve levels.

Observation 3. ***[Reserve Levels]*** *All forward prices rise as reserves are depleted but the effect on the term structure of volatility is small (see Figure 6).*

Intuitively, as reserves are depleted we would expect to see the level of prices increase. This is indeed the case as shown in Figure 6 where panels (a) and (b) show forward curves at high and low reserve levels. Notice that prices at both the short and long end of the forward curve are higher when reserves are low. It is interesting to note that the term structure of volatility remains virtually unchanged as reserves decrease. This is again consistent with the behavior of volatilities in the Schwartz and Smith model as reflected in Equation 21. As reserves approach exhaustion, however, the price process will be more dependent on the cost process for the alternative technology and we would expect to see the term structure of volatility change.

Observation 4. ***[Interest Rates]*** *A decrease in the level of the interest rate increases prices and decreases the slope of the forward curves in the long run (see Figure 7).*

This observation extends the standard Hotelling result on the slope of the forward curve. The reason for the increase in prices is clear if one considers a two period model.

These aberrations have traditionally been attributed to seasonality in the arrival of information. Alternatively, Hong [6] has suggested that these violations of the Samuelson effect are a result of heterogeneity in the informational endowments of agents in the economy. We conjecture that it would be straightforward to allow for such violations in our model by specifying start up costs as a function of current production levels. However, we retain our current specification, as such violations are rare in the applications we have in mind.

In the last period, all reserves will be produced. Due to the fact that reserves are limited, this will result in a "scarcity rent" for the resource owners. The present value of this scarcity rent governs the first period production choice. Obviously, if interest rates fall, the benefit of holding reserves for another period rises. Thus, fewer producers extract the resource in the first period, increasing the current price.

Observation 5. ***[Demand Shock Volatility]*** *An increase in demand volatility has no effect on forward prices and causes price volatilities to rise. (See Figure 8.)*

In the simple example developed in Section 3 equation (14) shows that the spot price of the resource does not depend on the volatility of the demand shock. This is also the case in the more general setting. Comparison of panels (a) and (b) show that forward prices are insensitive to a change in the demand volatility from 15% to 20% per year. There is, however, a direct and intuitive effect on the term structure of volatility as is illustrated in panels (c) and (d).

5. Implications for Option Pricing

As described in the previous section, the commodity price process generated by the model has two components: a short-run mean reverting component, and a long-run growth component. Schwartz [13] and Schwartz and Smith [14] examine the ability of an empirical model with these characteristics to explain futures prices for several commodities. In a separate paper, Schwartz and Miltersen [15] describe how to use such a two-factor model to price options on commodities. In this section, we examine the ability of the Schwartz and Smith two-factor model to price options on commodities whose prices are generated by our model.

Schwartz [13] describes how to employ the Kalman Filter to estimate the parameters in his model using time-series data on a group of forward prices. Once the model parameters are identified it is then possible to price a broad range of financial instruments, including European options on the underlying commodity. Although the commodity prices from our model exhibit a two-factor behavior similar to that of Schwartz and Smith [14], the functional form describing the behavior of the factors is different. The endogenous supply response imposes a non-linear drift on the short-run component; when large demand shocks arrive, producers optimally increase (or decrease) production. However, because adjusting production is costly, small demand shocks do not result in large supply responses. Intuitively, the drift in the short-run component is "locally" linear but overall non-linear, the result being that large temporary shocks are significantly dampened. The important effect from the point of view of pricing options is that the distribution of prices from our model has truncated tails relative to those predicted by the Schwartz and Smith [14] model that is calibrated to a time-series of data generated by our model. We demonstrate this effect under the base-case parameterization described in Table 1.

We perform a straightforward experiment to analyze the ability of the Schwartz and Smith two factor model to predict option prices in our setting. First, a time series of forward curves are simulated. We assume that the sampling interval is weekly, that the time series observations are available for the last year and that monthly contracts extending out two years are observable. We then use this data to calibrate the two factor model. Option prices implied by the calibrated Schwartz and Smith model are then calculated and compared to those generated by our equilibrium model.

In general, the calibrated two factor model overvalues options with maturities ranging from one to five years. Table 2 summarizes this result. The magnitude of the overpricing is potentially significant for large scale investment projects with a real-options component.

Further research is required to determine a feasible and robust procedure to deal with this bias.

6. Conclusion

In this paper, we have developed a general equilibrium model of exhaustible resource prices and extend the existing literature in a number of directions. Using a simple example we show that uncertainty alone cannot explain the backwardation observed in resource markets. In fact, for resources with flexible production processes forward prices will rise at the rate of interest and temporary demand shocks will be uniformly transmitted throughout the forward curve. In addition, we show that in this context the term structure of volatility will will be low and constant. In light of these results we conclude that, in the absence of frictions, the equilibrium price process will not exhibit the rich behaviour observed for commodities such as oil and gas. Therefore, we incorporate an extra cost associated with developing new reserves. Although introducing this extra cost significantly complicates the analysis and necessitates a numerical solution, we are able to generate endogenous price processes that can exhibit both backwardation of the forward curve and mean reversion in the spot price. We examine the implications of our model for real investment decisions. The equilibrium price process has truncated tails relative to the price distribution implied by the Schwartz and Smith [14] model. As a result we conclude that options whose payoffs are sensitive to extreme realizations are less valuable in our equilibrium setting where the extreme realizations occur less frequently than their model would predict.

A. Appendix: Proofs

Proof of Proposition 1: For the simplified demand process $p_t = \frac{y+\varepsilon_t}{q_t}$ where $\varepsilon_t > -y$, $\varepsilon_t \sim iid$ and $E(\varepsilon_t = 0)$. Consider the objective function equivalent to Equation 11:

$$\max_{q_t} E \sum_{t=0}^{\infty} e^{-rt} p_t q_t$$

subject to

$$\sum_{t=0}^{\infty} q_t = R_0 \quad \text{a.s.}$$

or equivalently,

$$\max_{q_t} \sum_{\omega} \sum_{t} e^{-rt} p_t(\omega)\, q_t(\omega)\, \pi(\omega) - \lambda(\omega) \left[\sum_{t} q_t(\omega) - R_0 \right]$$

where $\lambda(\omega)$ is the Lagrange multiplier process and $\pi(\omega)$ is the probability of a path. This optimization problem implies two first order conditions:

$$e^{-rt} p_t(\omega)\, \pi(\omega) = \lambda(\omega) \tag{A.1a}$$

$$\sum_{t=0}^{\infty} q_t(\omega) = R_0 \quad \forall\, t, \omega \tag{A.1b}$$

Now along any path ω, define $\widehat{\lambda}(\omega) \equiv \frac{\lambda(\omega)}{\pi(\omega)}$ and thus:

$$\widehat{\lambda}(\omega) = e^{-rt} p_t(\omega)$$

Substitute in $q_t = \frac{y+\varepsilon_t}{p_t} = \frac{y+\varepsilon_t}{\widehat{\lambda}(\omega)e^{rt}}$ into Equation A.1b and obtain

$$\widehat{\lambda}(\omega) = \frac{\sum_{t=0}^{\infty}(y+\varepsilon_t)e^{-rt}}{R_0} \quad \forall\ \omega$$

Let S be the set of ω such that such that $R_t = R, \varepsilon_t = \varepsilon$ and $p_t = p$ and use Equation A.1a to sum over S:

$$\sum_{\omega \varepsilon S} e^{-rt} p_t(\omega)\,\pi(\omega) = \sum_{\omega \varepsilon S} \lambda(\omega)$$

which implies,

$$\begin{aligned} e^{-rt}p_t &= \frac{\sum_{\omega\varepsilon S}\lambda(\omega)}{\sum_{\omega\varepsilon S}\pi(\omega)} \\ &= \frac{\sum_{\omega\varepsilon S}\pi(\omega)\,\widehat{\lambda}(\omega)}{\sum_{\omega\varepsilon S}\pi(\omega)} \\ &= \sum_{\omega\varepsilon S}\widehat{\lambda}(\omega) * \left(\frac{\pi(\omega)}{\sum_{\omega\ \varepsilon\ S}\pi(\omega)}\right) \\ &= E_t\left[\widehat{\lambda}|\ R,\varepsilon,P\right] \end{aligned}$$

However, recall that $\widehat{\lambda}(\omega) = e^{-rt}p_t(\omega)$. Thus discounted prices are martingales. To obtain the second part of the proposition note that:

$$\begin{aligned} E_0\left(\widehat{\lambda}\right) &= E_0\left(\sum_{t=0}^{\infty} e^{-rt}\frac{y+\varepsilon_t}{R_0}\right) \\ &= \varepsilon_0\left(\sum_{t=0}^{\infty} e^{-rt}\frac{y}{R_0}\right) \\ &= \frac{ay+\varepsilon_0}{R_0} \end{aligned}$$

and similarly for any time t,

$$E_t\left(\widehat{\lambda}\right) = e^{-rt}\left(\frac{ay+\varepsilon_t}{R_t}\right)$$

and hence,

$$p_t = \frac{ay+\varepsilon_t}{R_t}$$

Q.E.D.

Proof of Proposition 2: Given the expression for price $p_t = \frac{ay+\varepsilon_t}{R_t}$,we clarify the form of the reserves process, R_t :

$$\begin{aligned} R_1 &= R_0 - q_0 \\ &= R_0 - \frac{y+\varepsilon_0}{p_0} \\ &= R_0 - \frac{y+\varepsilon_0}{\frac{ay+\varepsilon_0}{R_0}} \\ &= R_0\left(1 - \frac{y+\varepsilon_0}{ay+\varepsilon_0}\right) \\ &= \frac{(a-1)\,yR_0}{ay+\varepsilon_0} \end{aligned}$$

Extending this logic by a simple induction argument, it is apparent that $R_t = R_0 \prod_{i=1}^{t} \eta_t$, where the η_t are IID shocks. Substituting for R_t in the expression for prices we obtain:

$$\begin{aligned} p_t &= \frac{ay+\varepsilon_t}{R_t} \\ &= \frac{ay+\varepsilon_t}{R_0 \prod_{i=1}^{t} \eta_t} \end{aligned}$$

Hence, the term structure of volatility is easily obtained from:

$$\begin{aligned} var\left(\log P_{t+s}\right) &= var\left(\log\left(ay+\varepsilon_t\right)-\sum_{i=t}^{s}\log\eta_i-\log R_0\right) \\ &= A+(s-t)\sigma_\eta^2 \end{aligned}$$

Q.E.D.

Proof of Proposition4 : Consider the objective function equivalent to Equation :

$$\max_{q_t} E\sum_{t=0}^{\infty} e^{-rt}p_t q_t-\gamma(q_t-q_{t-1})^{+}$$

subject to

$$\sum_{t=0}^{\infty} q_t = R_0 \quad \text{a.s.}$$

or equivalently,

$$\max_{q_t}\sum_{\omega}\sum_{t} e^{-rt}\left(p_t\left(\omega\right)q_t\left(\omega\right)-\gamma(q_t-q_{t-1})^{+}\right)\pi\left(\omega\right)-\lambda\left(\omega\right)\left[\sum_{t} q_t\left(\omega\right)-R_0\right]$$

where $\lambda\left(\omega\right)$ is the Lagrange multiplier process and $\pi\left(\omega\right)$ is the probability of a path. However, for some path ω and any time t, the equilibrium solution can be obtained from a consideration of the following simplified equation where we suppress the ω dependence for clarity:

$$e^{-rt}\left(p_t q_t-\gamma(q_t-q_{t-1})^{+}\right)\pi\left(\omega\right)+e^{-r(t+1)}\left(p_{t+1}q_{t+1}-\gamma(q_{t+1}-q_t)^{+}\right)\pi\left(\omega\right)-\lambda\left(\omega\right)q_t$$

The setup costs complicate the problem resulting in the following first order conditions:

$$\left[e^{-rt}p_t - \gamma e^{-rt}1_{(q_t\ >q_{t-1})} + \gamma e^{-r(t+1)}1_{(q_{t+1}>q_t)}\right]\pi(\omega) = \lambda(\omega) \qquad \text{(A.3)}$$

$$\sum_{t=0}^{\infty} q_t(\omega) = R_0 \quad \forall\ t, \omega \qquad (24)$$

Where, $1_{(q_t\ >q_{t-1})}$ is the indicator function on the set defined in the subscript. Now as before define S as the set of ω such that such that $R_t = R, \varepsilon_t = \varepsilon$ and $p_t = p$ and use Equation A.3a to sum over S:

$$\sum_{\omega\varepsilon S}\left[e^{-rt}p_t - \gamma e^{-rt}1_{(q_t\ >q_{t-1})} + \gamma e^{-r(t+1)}1_{(q_{t+1}>q_t)}\right]\pi(\omega) = \sum_{\omega\varepsilon S}\lambda(\omega)$$

Consider first the set of ω, such that $q_t < q_{t-1}$which implies:

$$\begin{aligned}
e^{-rt}p_t\sum_{\omega\varepsilon S}\pi(\omega) + \sum_{\omega\varepsilon S}\pi(\omega)\gamma e^{-r(t+1)}1_{(q_{t+1}>q_t)} &= \sum_{\omega\varepsilon S}\lambda(\omega) \\
e^{-rt}p_t + \frac{e^{-rt}\sum_{\omega\varepsilon S}\pi(\omega)\gamma e^{-r}1_{(q_{t+1}>q_t)}}{\sum_{\omega\varepsilon S}\pi(\omega)} &= \frac{\sum_{\omega\varepsilon S}\pi(\omega)\widehat{\lambda}(\omega)}{\sum_{\omega\varepsilon S}\pi(\omega)} \\
e^{-rt}p_t + e^{-rt}\sum_{\omega\varepsilon S}\gamma e^{-r}1_{(q_{t+1}>q_t)} * \left(\frac{\pi(\omega)}{\sum_{\omega\ \varepsilon\ S}\pi(\omega)}\right) &= \sum_{\omega\varepsilon S}\widehat{\lambda}(\omega) * \left(\frac{\pi(\omega)}{\sum_{\omega\ \varepsilon\ S}\pi(\omega)}\right) \\
e^{-rt}p_t + e^{-rt}B_t &= E_t\left[\widehat{\lambda}\middle|\ R, \varepsilon, P\right] \\
e^{-rt}p_t + e^{-rt}B_t &= k_t
\end{aligned}$$

Where $B_t = E_t\left[\gamma e^{-r}1_{(q_{t+1}>q_t)}\middle|\ R, \varepsilon, P\right]$, k_t is a random variable.and $1_{(.)}$ is the indicator function. Note that B_t in this case can be interpreted as a binary option which pays off γe^{-r} in those states ω when $q_{t+1}(\omega) > q_t(\omega)$. Similarly for the set of ω, such that

$q_t > q_{t-1}$the following condition holds:

$$e^{-rt}p_t - \gamma + e^{-rt}B_t = k_t$$

Recall that $p_t = \frac{y+\varepsilon_t}{q_t}$ and hence the two conditions can be expresses as:

$$e^{-rt}\frac{y+\varepsilon_t}{q_t} + e^{-rt}B_t = k_t \tag{A.3a}$$

$$e^{-rt}\frac{y+\varepsilon_t}{q_t} - \gamma + e^{-rt}B_t = k_t \tag{A.3b}$$

Intuitively, these first order conditions illustrate that at time t the production, q_t, is chosen to satisfy the R.H.S.which is fixed. Note however that the choice of q_t also influences the value of the binary option B_t. However, increases in q_t serve to decrease the L.H.S of Equation A.3a. Thus, producers will optimally incur the added expense γ and increase production only when ε_t is sufficiently large. On the other hand, when ε_t is small producers will decrease production in accordance with Equation A.3b. However, for intermediate demand shocks, producers may not find it optimal to decrease production and yet the demand shock may not be sufficient to incur the expense of increasing production. Thus neither first order condition is satisfied and production remains unchanged. This proves the proposition

References

[1] Black, F.,1976, The pricing of commodity contracts, Journal of Financial Economics 3,167-179.

[2] Brennan, M. and E. Schwartz, 1985, Evaluating natural resource investments, Journal of Business 20,135-157.

[3] Deaton, A., and G. Laroque, 1996, Competitive Storage and commodity price dynamics, Journal of Political Economy 104, 896-923.

[4] Dixit, A., and R. Pindyck, 1994, Investment under uncertainty (Princeton University Press).

[5] Hotelling, H., 1931, The Economics of Exhaustible Resources, Journal of Political Economy 39,137-175

[6] Hong, H., 2000, A Model Of Returns And Trading In Futures Markets, forthcoming Journal of Finance.

[7] Kushner, H., and P. Dupuis, 1992, Numerical Methods for stochastic control problems in continuous time (Springer Verlag).

[8] Litzenberger, R. and N. Rabinowitz, 1995, Backwardation in oil futures markets: Theory and empirical evidence, Journal of Finance 50, 1517-1545.

[9] Pindyck, R., 1980, Uncertainty and exhaustible resource markets, Journal of Political Economy 88, 1203-1225.

[10] Puterman, M.L., 1994, Markov Decision Processes: Discrete Stochastic Dynamic Programming (John Wiley and Sons, Inc.).

[11] Routledge, B., D. Seppi, and C. Spatt, 2000, Equilibrium forward curves for commodities, Journal of Finance forthcoming.

[12] Samuelson, P., 1965, Proof that properly anticipated prices fluctuate randomly, Industrial Mangement Review 6, 41-49.

[13] Schwartz, E., 1997, The stochastic behaviour of commodity prices: Implications for valuation and hedging, Journal of Finance 52,923-973.

[14] Schwartz, E. and J. Smith, 1998, Short-Term Variations and Long-Term Dynamics in Commodity Prices, working paper.

[15] Schwartz, E. and K. Miltersen, 1998, Pricing of Options on Commodity Futures with Stochastic Term Structures of Convenience Yields and Interest Rates, Journal of Financial and Quantitative Analysis 33,33-59.

Parameter Name	Symbol	Value
Risk-free interest rate	r	0.05
Long-run average demand	μ_y	3.69
Rate of mean reversion of demand	κ_y	1.00
Volatility of demand	σ_y	0.15
Drift of cap	μ_S	0.00
Volatility of cap	σ_S	0.05
Cost of increasing production	γ	0.50
Extraction Cost	C	0.00

Table 1: **Parameter values for the base case.**

Time to Maturity (years)	Option Price (in dollars)		Overpricing
	Equilibrium	Calibration	(in percent)
1	0.0498	0.0522	5
2	0.0476	0.0581	22
3	0.0594	0.0619	4
4	0.0698	0.0653	-6
5	0.0589	0.0685	16

Table 2: **Comparison of option prices.** This table compares the model's actual option prices to the option prices generated by a calibration of the Schwartz and Smith [14] model. The option prices from the calibrated model are, in general, higher than the actual option prices.

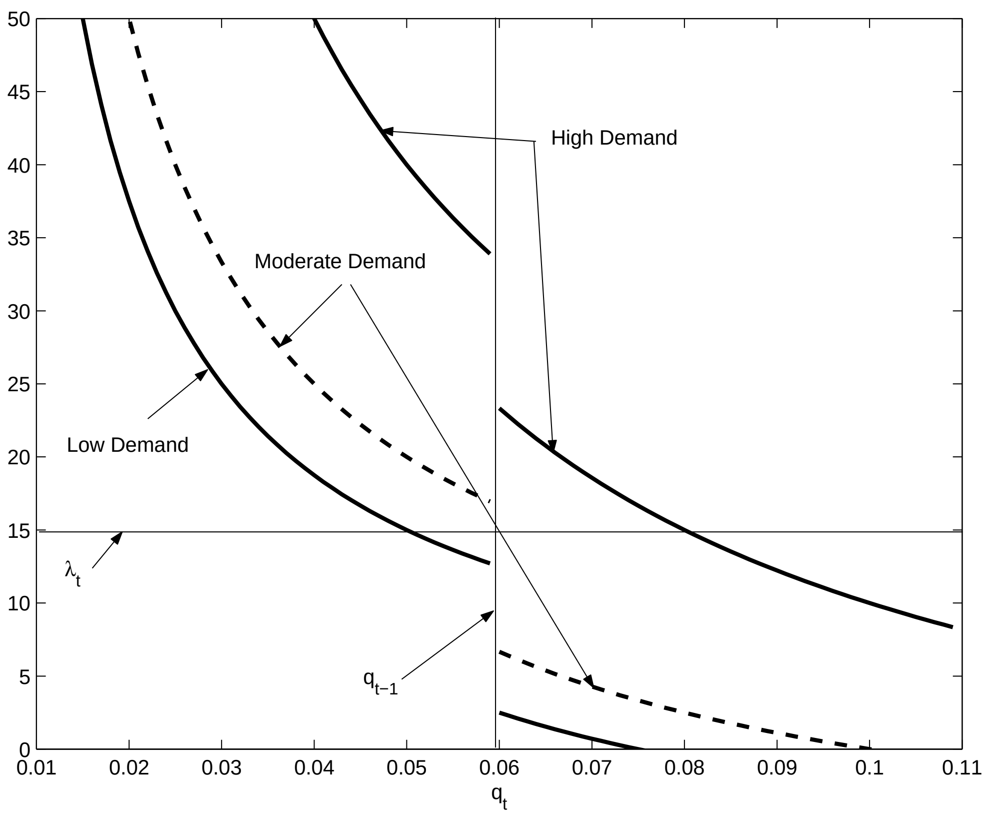

Figure 1: **The first order conditions from the model with startup costs.** The presence of startup costs introduces a region in which production rates will not be changed.

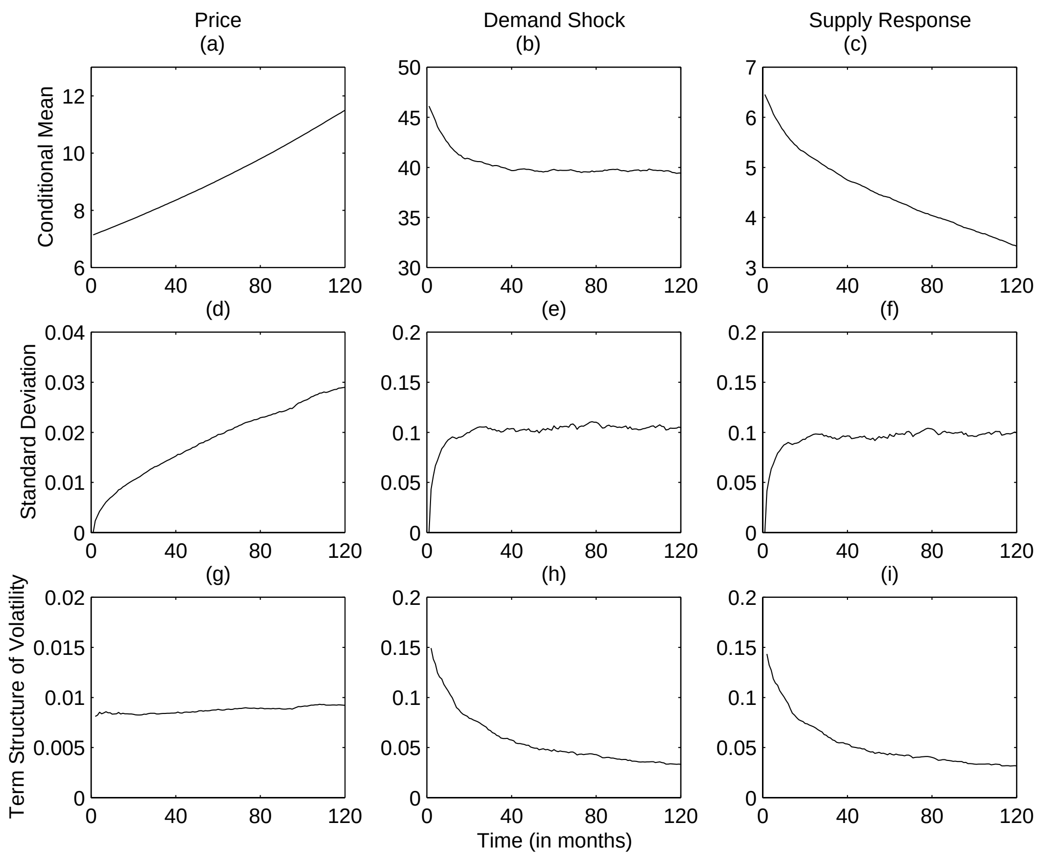

Figure 2: **Summary output from the model with no startup costs.** In the leftmost column, conditional means and variances of the model's equilibrium prices are presented. The middle and right columns give analogous results for the demand shock and optimal supply response.

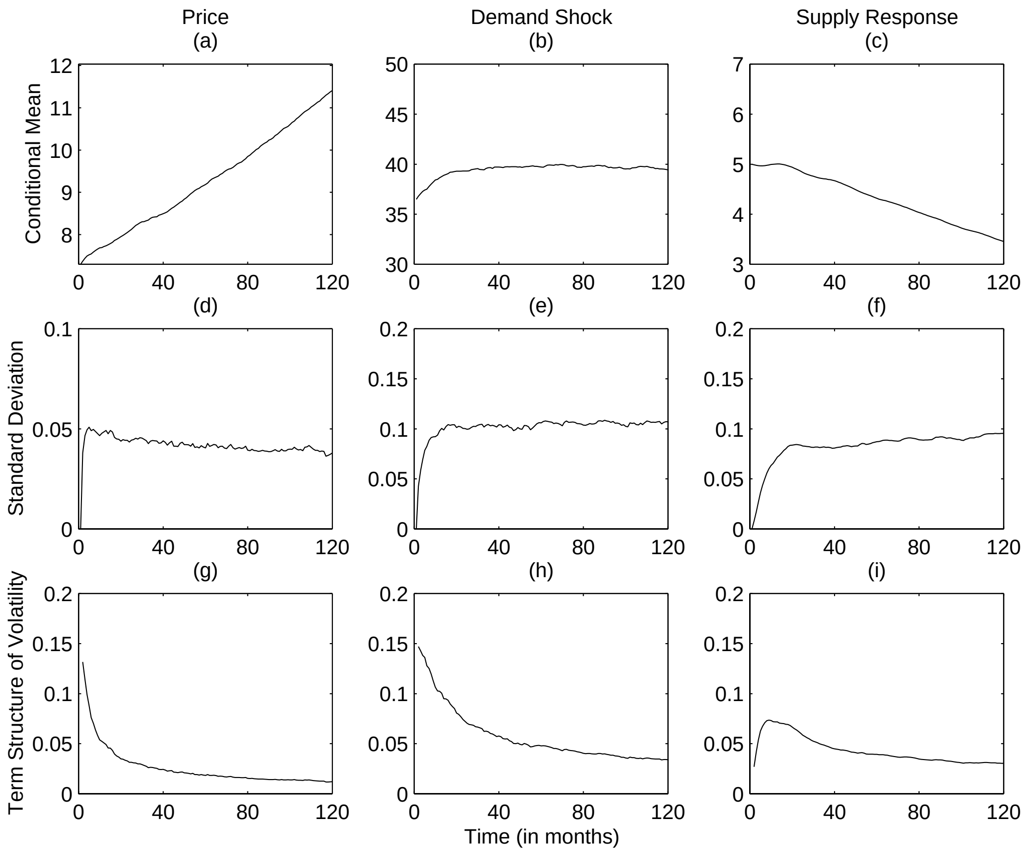

Figure 3: **Summary output from the model with high startup costs.** In the leftmost column, conditional means and variances of the model's equilibrium prices are presented. The middle and right columns give analogous results for the demand shock and optimal supply response.

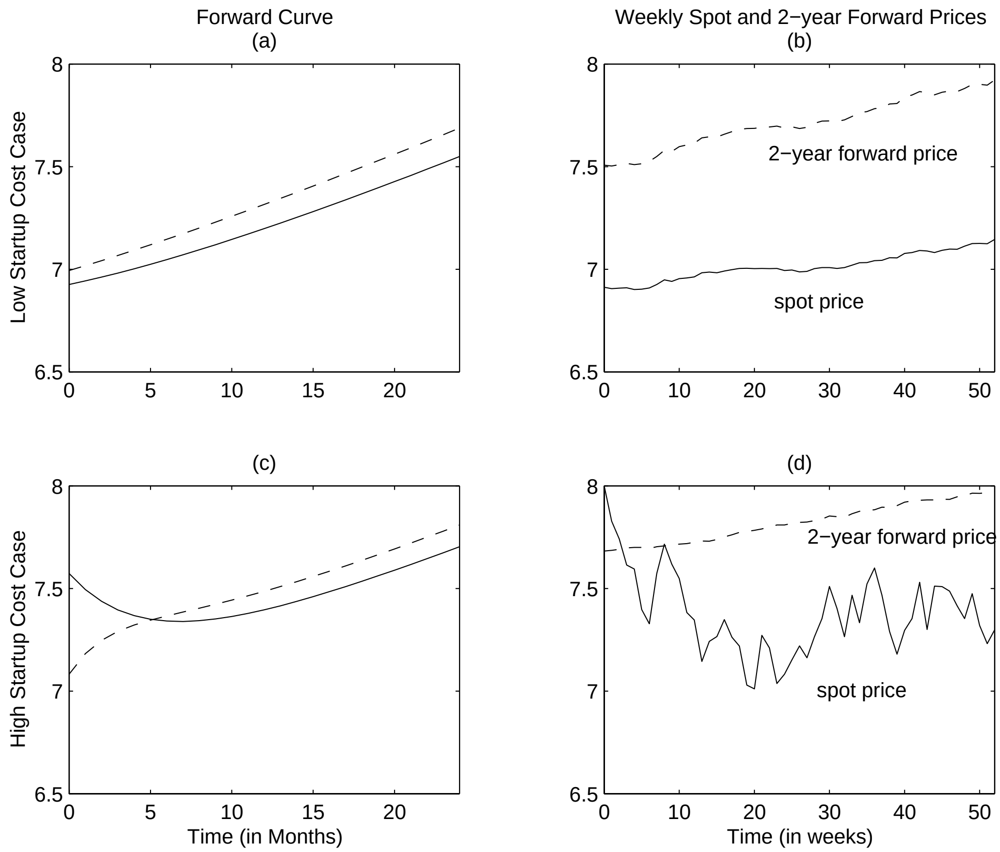

Figure 4: **Dynamics of forward prices.** On the left, forward curves from the model with zero and high startup costs are presented. Panel (a) shows that without startup costs forward curves at all dates are parallel. This is not the case when startup costs are high as shown in panel (c). Simulated spot and 2-year forward prices are displayed on the right. Panel (b) shows that without startup costs, spot prices and forward prices are equally variable. With high startup costs the spot price is much more volatile than the forward prices as shown in panel (d).

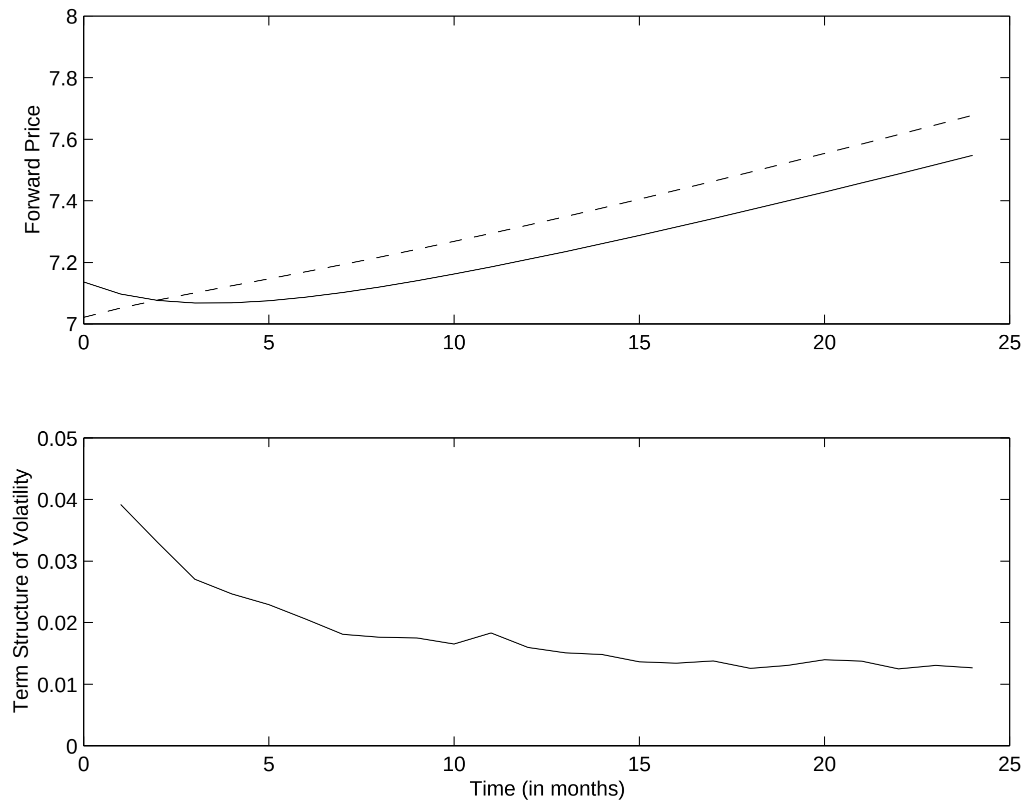

Figure 5: **Forward prices and the term structure of volatility: the base case.** The top panel presents two forward curves from the model under the base case parameterization. Forward curves may be backwardated or in contango. The lower panel displays the term structure of volatility for the base case. The magnitude of the volatility is low and declines with time.

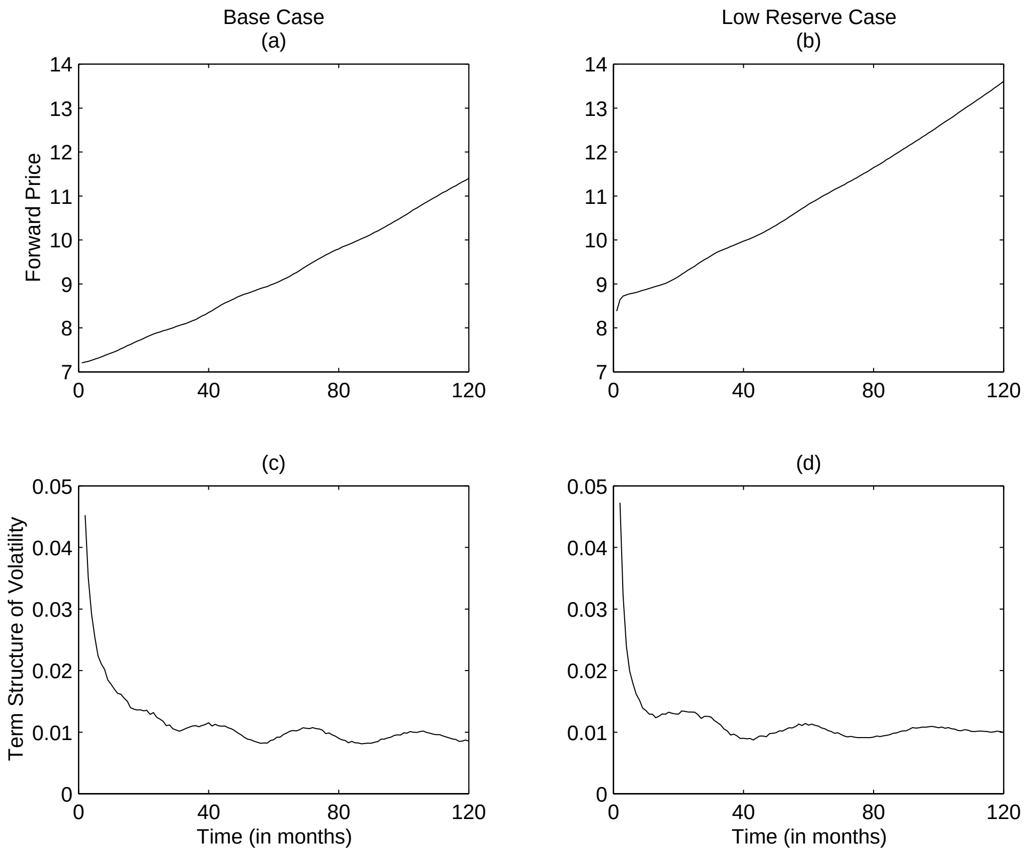

Figure 6: **Analysis of a change in the level of reserves.** Panels (a) and (b) show that when reserves drop forward prices rise. Panels (c) and (d) show that the term structure of volatility is insensitive to the amount of reserves.

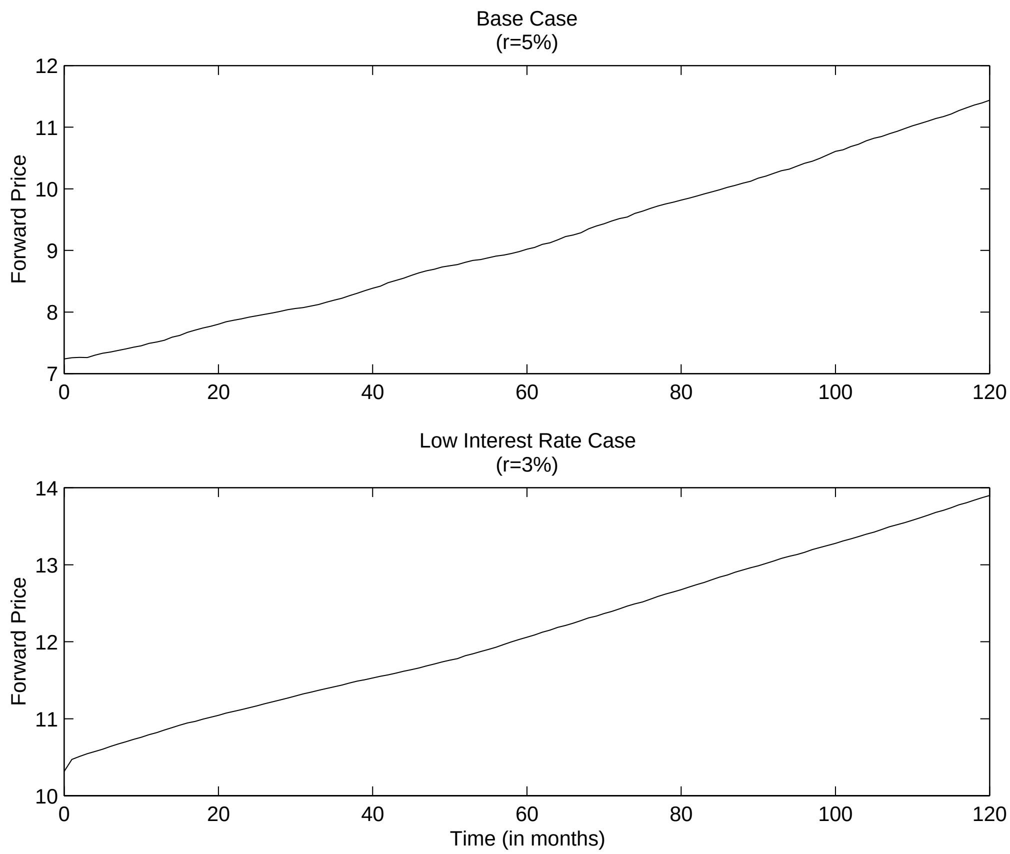

Figure 7: **Analysis of a change in the rate of interest.** Panels (a) and (b) show that when the interest rate decreases forward prices rise.

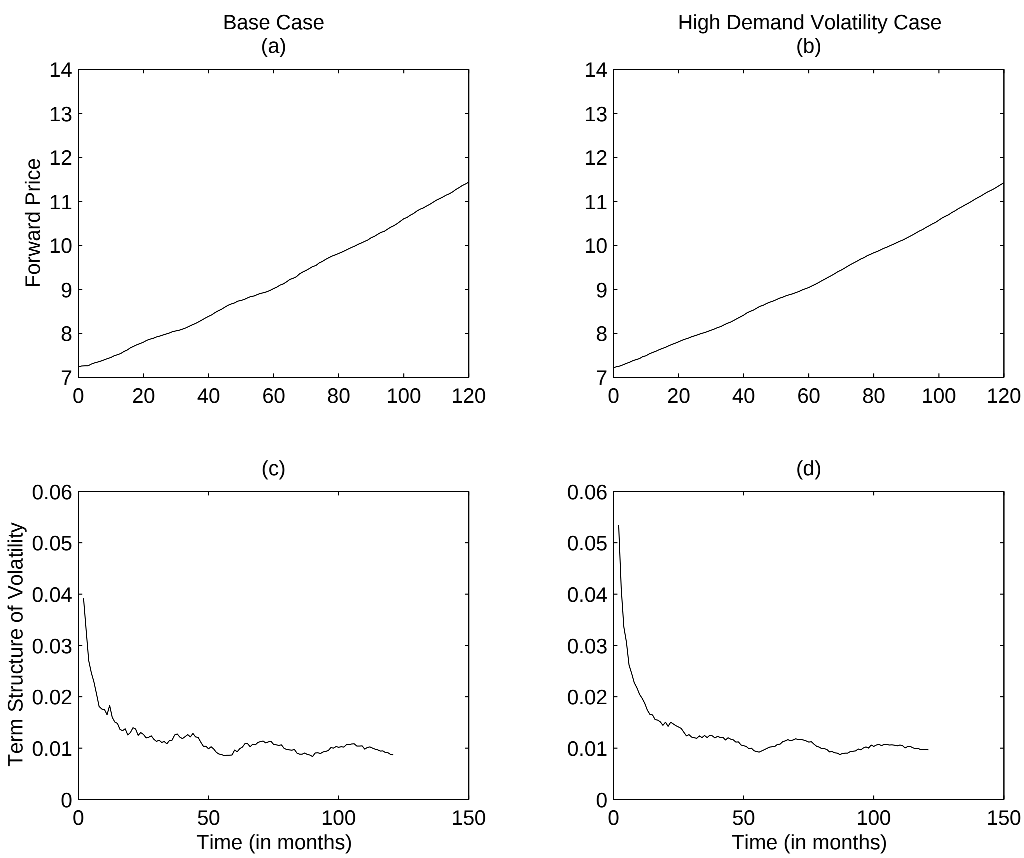

Figure 8: **Analysis of a change in the volatility of the demand shock.** Panels (a) and (b) show that forward prices do not change when the volatility of the demand shock increases. Panels (c) and (d) show that higher demand shock volatilities result in higher price volatilities.