

Equilibrium Exhaustible Resource Price Dynamics

MURRAY CARLSON, ZEIGHAM KHOKHER, and SHERIDAN TITMAN*

ABSTRACT

We develop equilibrium models of exhaustible resource markets with endogenous extraction choices and prices. Our analysis demonstrates how adjustment costs can generate oil and gas forward price dynamics with two factors, consistent with the behavior these commodities exhibit in the Schwartz and Smith (2000) calibration. Our two-factor model predicts that stochastic volatility will arise in these markets as a natural consequence of production adjustments, however, and we provide supporting empirical evidence. Differences between endogenous price processes from our general equilibrium model and exogenous processes in earlier papers can generate significant differences in both financial and real option values.

CONTINGENT CLAIMS ANALYSIS IS CURRENTLY USED extensively in natural resource industries. For example, energy traders often use models suggested by Black (1976), Brennan and Schwartz (1985), Schwartz (1997) and others for risk management as well as for valuing financial contracts and real investments. These reduced-form models, which assume an exogenous process that is typically calibrated to some combination of historical prices and observed forward and option prices, have been successful in valuing and hedging relatively short-term financial contracts. Such models can be viewed as tools for interpolating among comparable assets that trade in relatively liquid markets. They are less applicable, however, in situations where prices for directly comparable assets do not exist. For example, the valuation of an investment to exploit large oil and gas reserves requires estimates of long-dated forward and option prices for oil and gas. As we show, when these long-dated contracts are illiquid, there are problems associated with using existing pricing models to extrapolate their values from the observable prices of more liquid shorter-term contracts.

To explore these issues in greater detail we develop a general equilibrium model of an extractable resource market wherein both prices and extraction

*Sauder School of Business at the University of British Columbia, Richard Ivey School of Business at the University of Western Ontario, and McCombs School of Business at the University of Texas at Austin. We would like to thank an anonymous referee, Harjoat Bhamra, the editors, Ali Lazrak, Tan Wang, and seminar participants at the University of Texas, the Real Options Conference at Cambridge University, Texas A&M University, and the Western Finance Association Meetings in Park City for their comments. All errors are our sole responsibility.

choices are determined endogenously.¹ The fundamental sources of uncertainty in our model arise because of fluctuations in aggregate demand and changes in technology. Aggregate demand is assumed to follow a mean-reverting process while changes in technology, which affect the price of a potential future substitute for the commodity, fluctuate randomly.² Price responses to both sources of uncertainty are determined in part by endogenous supply responses (i.e., how production levels respond to changes in aggregate demand), and these responses are determined in turn by the nature of the technology for extracting the commodity.

We show that temporary demand shocks have a small but permanent effect on prices when producers can costlessly increase or decrease supply. Conversely, when the costs of altering current production are sufficiently high, temporary demand shocks have a disproportionately larger effect on current prices than on future prices, and the spot price will fail to respond to shocks that affect the cost of the future substitute. Hence, for the equilibrium price process to exhibit the long-term and short-term effects observed in historical data,³ producers must be able to alter production at a cost that is significant but not prohibitive.⁴

Our analysis is particularly close in spirit to Weinstein and Zeckhauser (1975) and Pindyck (1980), who add demand uncertainty to the seminal Hotelling (1931) model that describes how the prices of exhaustible resources evolve over time. These papers show that if competitive risk-neutral producers with zero marginal extraction costs can make costless supply adjustments, then expected prices (or equivalently, forward prices) rise at the riskless interest rate. The predictions of this earlier literature are clearly inconsistent with the data since, in reality, forward curves from oil and natural gas markets can be either backwardated or in contango.⁵ In addition, since these models assume that changing

¹ Deaton and Laroque (1996) and Routledge, Seppi, and Spatt (2000) analyze the effect of storage on commodity price dynamics when production is exogenous. We take the opposite approach, focusing instead on the effect that endogenous production decisions have on commodity price dynamics when no inventories are present.

² For example, oil shale is a potential future substitute for conventional oil.

³ Schwartz (1997) and Schwartz and Smith (2000) show that empirical models of oil prices with multiple factors, some mean-reverting and others permanent, outperform single-factor models with only short-term or long-term effects. Pindyck (1999) utilizes longtime series of spot prices for energy commodities, including natural gas, to estimate variance ratio statistics that are consistent with short-term and long-term components in prices. Fama and French (1988) and Bessembinder et al. (1995) show that many commodity prices, in particular, crude oil, have a mean-reverting component.

⁴ Adjustment costs are used in a number of studies, including Scarf (1960), Grossman and Laroque (1990), and Caballero and Engel (1999), to describe a variety of economic phenomena. In addition, Casassus, Collin-Dufresne, and Routledge (2004) as well as Kogan, Livdan, and Yaron (2004) study the effects of adjustment costs on futures prices in a production economy with irreversible investment.

⁵ See, for example, Litzenberger and Rabinowitz (1995). If futures prices are below the current spot price, the futures curve is said to be backwardated. Litzenberger and Rabinowitz 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.

the extraction rate is costless, they predict that prices are subject to only permanent shocks (i.e., price changes follow a random walk), whereas existing empirical evidence documents that prices of exhaustible commodities exhibit both permanent and temporary shocks.

Our model is also related to Litzenberger and Rabinowitz (1995), 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. Their model implies that forward prices are always weakly backwarddated, which is true on average for both oil and gas prices but is quite often violated for both commodities.⁶ As mentioned above, our model predicts periods during which forward curves will be in weak contango as well as in backwardation, and is therefore consistent with this aspect of the data.

Our model is similar to that of Litzenberger and Rabinowitz in that we both consider the possibility that volatility changes over time and examine the relation between volatility and the slope of the forward curve. In the Litzenberger and Rabinowitz model the volatility of demand for the commodity changes exogenously, which in turn causes the forward curve to change. When volatility is high, the value of delaying production increases, causing current prices to increase relative to future prices. The volatility of demand is constant in our model but production adjustment costs give rise to endogenous extraction choices, which in turn cause the volatility of resource prices to be high when demand is either high or low. These differences between the Litzenberger and Rabinowitz partial equilibrium model and our general equilibrium model are empirically testable. We predict a U-shaped relationship between the futures slope and volatility, where volatility is high when forward curves are either steeply backwarddated or in contango, whereas they predict a monotonic relationship. Using oil and natural gas price data, we construct a simple test and reject their hypothesis in favor of ours.

In addition to making the above theoretical contribution, the structural model developed in this paper also offers a practical contribution. In particular, we provide a method by which one can explicitly incorporate information about both supply variables, such as production costs and the costs of close substitutes, and demand variables, such as elasticities and income growth rates, into a model that can be used to value both financial and real investments. To illustrate the importance of incorporating this kind of information into a valuation model we

⁶ Empirical evidence from crude oil futures markets is somewhat consistent with their prediction during the time period considered: "Between February 1984 and April 1992 the nine months futures price was strongly backwarddated 77 percent of the time and weakly backwarddated 94 percent of the time" (Litzenberger and Rabinowitz (1995), page 1517). In the 1990s, however, oil futures prices were often in contango. Between April 1991 and June 1999 the 12-month futures price was strongly (weakly) backwarddated only 56% (75%) of the time. Natural gas is a depletable resource that is arguably less susceptible to direct price manipulation by producers. Between April 1991 and June 1999 the 12-month futures price for natural gas was strongly (weakly) backwarddated only 45% (60%) of the time. Thus, the more recent price data are somewhat at odds with a model that does not explain the frequent occurrence of weak contango in exhaustible resource prices.

compare the option prices generated by our structural model with the prices generated by the Schwartz and Smith reduced-form model, calibrated to a time series of forward prices generated by simulations from our model.

As we show, using plausible parameters our model generates prices that are roughly consistent with observed forward prices for oil as well as with the price processes calibrated in Schwartz and Smith. However, the subtle differences between the endogenous price process determined within our general equilibrium model and the exogenous processes assumed in earlier papers can generate significant differences in both financial and real option values. For example, although the endogenous price process generated by our model is qualitatively similar to the price process assumed by Schwartz and Smith, the functional form of the drift is generally nonlinear and it generates equilibrium price paths with less extreme realizations than would obtain in their 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 outline of the paper is as follows. In Section I, we analyze models in which the production choice is completely flexible. We find that such models are inconsistent with oil and natural gas spot and forward prices. In Section II we present a general model of the resource market, and we present implications of the equilibrium model for production decisions, forward prices, and price volatilities. Finally, Section III discusses empirical implications of the model and compares our predictions to those of Schwartz and Smith (2000).

I. The Resource Extraction Problem with Flexible Production

This section analyzes four closely related models of equilibrium price determination in exhaustible resource markets assuming production rates are flexible and may be changed at no cost. Each model relies on the same intuition, namely, that producers of a resource will shift output across time periods so as to maximize the resource's value. This principle has two important consequences. First, resource prices have only permanent components regardless of whether demand shocks are permanent or temporary, that is, price changes follow random walks. This follows from the fact that in the case of an exhaustible good it does not matter whether consumption is motivated by permanent or transitory shocks since current consumption has a permanent effect on the remaining supply and, therefore, on all future prices. Second, when demand shocks are temporary, optimal supply responses may exhibit nonconstant elasticity with respect to demand, giving rise to endogenous stochastic volatility in the resource price. This second effect does not arise in the classical equilibrium models of exhaustible resource prices.

The following assumptions hold throughout this section. First, inverse demand is given by

$$p_t = \frac{e^{a+y_t}}{q_t^\gamma}, \quad (1)$$

where q_t is the instantaneous aggregate production rate and $\gamma \geq 1$ determines demand elasticity. The state variable y_t drives demand dynamics. Second, aggregate reserves are known and finite with endogenous dynamics given by

$$dR_t = -q_t dt. \quad (2)$$

The optimal production rate q_t may depend on time, the demand state y_t , and remaining reserves R_t . Third, we assume that all individuals are risk neutral.⁷ Finally, producers are assumed to operate in a competitive market where marginal extraction costs are zero.

In a competitive market with value-maximizing producers, prices are determined such that there is no incentive to shift production across periods. Under our assumptions this implies that prices are expected to grow at the constant riskless rate r . Additionally, the equilibrium aggregate production policy must result in the eventual extraction of all reserves.⁸

Forward prices are determined in equilibrium by risk-neutral traders. We denote by $f_{t,u}$ the forward price at date t for a unit of the commodity to be paid for and delivered at date $u > t$ and assume that speculators compete expected profits to zero.⁹ This condition implies the following characterization of forward prices in terms of current expectations of future spot prices:¹⁰

$$f_{t,u} = E_t(p_u). \quad (3)$$

A. Demand Dynamics with Riskless Innovations

We first consider a demand state variable with riskless dynamics. Although this model is well understood, it serves to develop the solution methods we employ in the more interesting cases that follow. As in Hotelling (1931), production

⁷ Our model does not explicitly incorporate a specification of risk premia since our explanation of the empirical phenomena discussed earlier is not risk based. Under certain specifications of the market price of risk (e.g., a constant) all our theoretical results will hold, with expectations interpreted as being calculated with respect to the risk-neutral probabilities. The assumption of risk neutrality eliminates the need to empirically estimate the risk premium in Section III, but the point we make there is valid regardless of the nature of risk premia.

⁸ In a few cases this condition will be satisfied in finite time, but in many cases the constraint will hold only as time approaches infinity. Assuming the demand curve is fixed ($g = 0$ in our setting), Hotelling shows that "whether the time until exhaustion will be finite or infinite turns upon whether a finite or an infinite value of p will be required to make q vanish" (p. 142). For the above specification with $\gamma = 1$, $q_t = \frac{e^{a+\gamma t}}{p_t}$, which implies that extraction will continue forever. The resource will be extracted in finite time if, for example, $q_t = a - p_t$.

⁹ We do not make a distinction between forward and futures prices since interest rates are nonstochastic in our setting.

¹⁰ Given this relationship, forward prices will depend on the level of all state variables relevant for forecasting future spot prices. In some cases, the current spot price will be a sufficient statistic for this forecast and there will be a straightforward relationship between the current spot price and forward prices. This is not the case in general, however, and forward prices are typically affected by information other than current price.

is proportional to remaining reserves and the resource is depleted at a rate that causes equilibrium prices to increase at the riskless rate.

PROPOSITION 1: *If demand in equation (1) is driven by the state variable y_t with dynamics $dy_t = g dt$, where $g < r$, then there exists an equilibrium in which the resource is depleted at a rate that is proportional to remaining reserves,*

$$q_t = \frac{r - g}{\gamma} R_t, \quad (4)$$

with the level of reserves at any time given by $R_t = R_0 \exp(-\frac{r-g}{\gamma}t)$.

Proof: We show the existence but not the uniqueness of an equilibrium in which the quantity produced is linear in the level of reserves. Assume that the optimal production policy has the form $q_t = \beta R_t$. Equilibrium price dynamics must be given by $dp_t = r p_t dt$. Furthermore, the dynamic equation for prices is implied by a differential equation that incorporates the functional form of inverse demand, the dynamics of the state variable, and the production policy

$$dp_t = p_t dy - \gamma \beta \frac{p_t}{\beta R_t} dR_t \quad (5)$$

$$= (g + \gamma \beta) p_t dt. \quad (6)$$

This differential equation identifies the optimal extraction rate, since, for prices to grow at the riskless rate, $g + \gamma \beta = r$ or equivalently $\beta = (r - g)/\gamma$. Reserves dynamics are given by $dR_t = -q_t dt = -\frac{r-g}{\gamma} R_t dt$ and the stated relationship between reserves and time solves this differential equation. The second requirement for equilibrium production is satisfied since, by inspection, reserves approach zero in the limit as time approaches infinity. Q.E.D.

In this setting forward prices follow from equation (3), which requires that the expected contract value at initiation is zero. For a contract established at date t expiring at date $t + m$ the forward price is $f_{t,t+m} = p_t e^{rm}$. The “slope” of the forward curve is thus constant and equal to the riskless interest rate.

B. Demand Dynamics with Risky, Permanent Innovations

We now consider a demand state variable whose dynamics incorporate only permanent components, which is a special case of a problem previously analyzed by Pindyck (1980). Unlike previous work, however, under our specific set of assumptions we are able to explicitly solve for price, production, and reserve dynamics.

PROPOSITION 2: *If demand in equation (1) is driven by the state variable y_t with dynamics $dy_t = g dt + \sigma dW_t$, where $r > g + \frac{1}{2}\sigma^2$ and dW are increments to a standard Brownian motion, then in equilibrium the resource is depleted at a rate that is proportional to remaining reserves,*

$$q_t = \frac{r - g - 1/2\sigma^2}{\gamma} R_t, \quad (7)$$

with the level of reserves at any time given by the deterministic function $R_t = R_0 \exp(-\frac{r-g-1/2\sigma^2}{\gamma}t)$.

Proof: As in the certainty case, we verify that the equilibrium production policy has the form $q_t = \beta R_t$. Applying Ito's lemma, price dynamics are given by

$$dp_t = p_t g dt + p_t \sigma dW_t + \frac{1}{2} p_t \sigma^2 dt + \gamma \beta p_t dt \quad (8)$$

$$= \left(g + \frac{1}{2} \sigma^2 + \gamma \beta \right) p_t dt + \sigma p_t dW_t. \quad (9)$$

In order for prices to increase at the riskless rate, optimal extraction must solve $g + \frac{1}{2} \sigma^2 + \gamma \beta = r$ so that production at any date is given by $q_t = \frac{r-g-1/2\sigma^2}{\gamma} R_t$. The formula relating reserves and time is determined by integration. Again, reserves approach zero as time approaches infinity and the production policy specified above is an equilibrium. Q.E.D.

In this setting, the production policy and the associated reserve dynamics are deterministic even though prices are stochastic. Permanent demand shocks imply that when demand is currently high, it is expected to be high in all future periods. Thus, there is no need to respond to higher demand by increasing production. Production is therefore equal to that in the certainty case, with a minor adjustment to account for the impact that convexity in inverse demand has on expectations of future prices.

Futures prices are proportional to spot prices as in the certainty case ($f_{t,t+m} = p_t e^{rm}$) and shocks to current demand move the entire forward curve up or down without any effect on its slope; hence, the elasticity of futures prices with respect to spot prices is equal to one and, using Bessembinder et al.'s (1995) definition, no mean-reversion is present in the commodity price. This in turn implies that the volatility of futures prices and spot prices are constant and equal, and thus the volatility of futures prices is constant for all maturities. In other words, the Samuelson (1965) effect does not arise in this setting.

C. Demand Dynamics with Risky, Temporary Innovations

We now assume that the demand state variable has Ornstein–Uhlenbeck dynamics, whereby temporary shocks to demand decay at an exponential rate. In this case, although the optimal production policy has no closed-form solution, it is possible to determine some of its basic properties. The following proposition establishes that the extraction rate is proportional to reserves and that an ordinary differential equation (ODE) characterizes the dependency on demand:

PROPOSITION 3: *If demand in equation (1) is driven by the state variable y_t with dynamics $dy_t = -\kappa y_t dt + \sigma dW_t$, where $\kappa > 0$ and dW are increments to*

a standard Brownian motion, then in equilibrium the resource is depleted at a rate that is proportional to remaining reserves, $q_t = e^{\beta(y_t)} R_t$, where the function $\beta(y_t)$ solves the second-order ODE

$$-\kappa y_t(1 - \gamma\beta'(y_t)) + \frac{1}{2}\sigma^2(-\gamma\beta''(y_t) + (1 - \gamma\beta'(y_t))^2) + \gamma e^{\beta(y_t)} = r. \quad (10)$$

Proof: We verify that the equilibrium production policy is of the form $q_t = e^{\beta(y_t)} R_t$. In this case production is linear in reserves and a nonconstant, nonlinear function of the demand state. Applying Ito's lemma, price dynamics are given by

$$\begin{aligned} dp_t &= [p_t - \gamma\beta'(y_t)p_t] dy + \frac{1}{2}\sigma^2[-\gamma\beta''(y_t)p_t + (1 - \gamma\beta'(y_t))^2 p_t] dt - \gamma \frac{p_t}{R_t} dR_t \\ &= \left[-\kappa y_t(1 - \gamma\beta'(y_t)) + \frac{1}{2}\sigma^2(-\gamma\beta''(y_t) + (1 - \gamma\beta'(y_t))^2) + \gamma e^{\beta(y_t)} \right] p_t dt \\ &\quad + \sigma(1 - \gamma\beta'(y_t))p_t dW_t. \end{aligned} \quad (11)$$

Equilibrium requires that $dp_t = rp_t dt + \sigma(y_t)dW_t$. Equating the drift terms gives rise to equation (10), a nonlinear second-order differential equation in y_t only, with solution $\beta(y_t)$, which characterizes the equilibrium production rate. A boundary condition is required to ensure that the resource is exhausted in the limit ($\int_0^\infty e^{\beta(y_t)} R_t dt = R_0$) so that the resulting production policy holds in equilibrium. The second boundary condition ensures that there is a solution to the differential equation for all levels of the state variable y . This is achieved by requiring $\lim_{y_t \rightarrow -\infty} 1 - \gamma\beta'(y_t) = 0$, a necessary condition for the first term of the differential equation (10) to approach r as $y \rightarrow -\infty$. Intuitively, this condition ensures that prices, which are proportional to $e^{y-\gamma\beta}$, become less sensitive to changes in demand when the drift of y becomes large, allowing prices to grow at the rate r . Q.E.D.

In this equilibrium, prices are expected to rise at the riskless interest rate and, as in the case in which demand shocks are permanent, forward prices are stochastic. The slope of the forward curve is not stochastic, however, because of the effect of production responses that convert temporary demand shocks to permanent price shocks. Panel A in Figure 1 presents an example of an optimal production policy given one parameterization of the model.¹¹ Holding reserves constant, when demand is high production is high. In these states the sensitivity of production to the demand state variable is also high. It is this sensitivity, as measured by the slope and convexity of (log) production, that gives rise to an endogenous price process with a constant drift r , a point made formal by the ODE in equation (10).

¹¹ In this example, $\kappa = 0.9$, $\sigma = 0.5$, $\gamma = 1$, and $r = 0.05$.




Figure 1. Optimal production policy and endogenous diffusion with mean-reverting demand. Producers respond to mean-reversion in demand by optimally adjusting production rates. Panel A displays the relationship between the demand state y_t and the production rate $e^{\beta(y_t)}$. Panel B displays the relationship between demand and volatility of log spot price changes. Volatility is stochastic and mean-reverting because of its monotonic relationship with y_t .

A key difference between this equilibrium and those in the previous subsections is that the volatility of changes in prices are stochastic:

COROLLARY 1: *The diffusion of the log price process is related to the demand state variable by the following equation:*

$$\sigma_p(y_t) = \sigma(1 - \gamma\beta'(y_t)). \quad (12)$$

Proof: Follows from inspection of equation (11). Q.E.D.

The state dependence of volatility is illustrated in Panel B of Figure 1, which plots the diffusion equation (12) relating the stochastic demand state variable to price volatility. Production responses are again responsible for this phenomenon. The intuition follows if one recognizes the derivative $\beta'(y_t)$ as the elasticity of production with respect to demand. Equation (12) then states that (constant volatility) changes in demand are converted into highly volatile equilibrium price changes in states in which production is most sensitive to the demand state variable. The reverse is true when the elasticity of production with respect to demand is low. We therefore have a structural model of resource price dynamics with mean-reverting stochastic volatility, since $1 - \gamma\beta'(y_t)$ is monotonic in the mean-reverting y_t state variable.

The predictions of this model are consistent with empirical characterizations of price dynamics for precious metals, commodities for which supply is relatively flexible and storage costs are low relative to inventory values. As Fama and French (1987, 1988) note, the slope of the gold and silver futures term structure is well described by the term structure of riskless interest rates. Consistent with the prediction of mean-reverting volatility, price changes for these commodities also exhibit GARCH effects (Ng and Pirrong (1994)).¹² Thus, as Corollary 1 predicts, price dynamics with only permanent components can be coupled with mean-reverting stochastic volatility in markets with flexible production.

Prices of other commodities, in particular oil and natural gas, have more complex dynamics exhibiting both permanent and temporary components (e.g., Schwartz (1997)). Engineers must deal with the complex physics of fluid dynamics when extracting these commodities, and such considerations place restrictions on the flexibility of supply. Section II considers these restrictions by adding an adjustment cost that is determined by historical production rates. These adjustment costs limit production responses and restrict their ability to transform temporary demand shocks into permanent price shocks.

D. Demand Dynamics with Risky, Independent Innovations

We now consider a simplified demand process that allows for closed-form solutions in discrete time. Specifically, we measure the demand state process y_t

¹² Empirical results not reported here confirm that these effects remain in data covering a longer and more recent time period.

at regular time intervals and assume these observations are independent and identically distributed. These demand shocks, which are an extreme example of temporary shocks, can be interpreted as the limiting case of the previous class of mean-reverting shocks in which the rate of mean-reversion, κ , is large.

The timing of the information and decisions is as follows. At the beginning of each decision period t , the current level of reserves is known to be R_t . Producers observe a shock to the demand curve y_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 are made, the level of reserves drops to $R_{t+1} = R_t - q_t$.

The following proposition characterizes the equilibrium price dynamics in this simplified case.

PROPOSITION 4: *If inverse demand is given by equation (1), where $\{y_t\}_{t=0}^{\infty}$ are independent random variables with $E(e^{y_t}) = 1 \forall t$ and $\gamma = 1$, then discounted prices in a competitive equilibrium are martingales. Thus, for $u > t$,*

$$E_t(e^{-ru} p_u) = e^{-rt} p_t. \quad (13)$$

Moreover, the price of the commodity at an arbitrary point in time is a function of two random state variables, y_t and R_t

$$p_t = (k + e^{y_t}) \frac{e^a}{R_t}, \quad (14)$$

where $k = \frac{1}{e^r - 1}$.

Proof: See the Appendix.

Again, the discounted expected value of the future spot price is the current spot price, and at every point in time prices are expected to rise at the riskless interest rate.¹³ Thus, the forward curve is given by

$$f_{t,u} = E_t(p_u) = e^{r(u-t)} p_t. \quad (15)$$

This illustrates again that uncertainty cannot, by itself, generate the backwardation result in Litztenberger and Rabinowitz (1995). Indeed, supply responses turn temporary demand shocks into permanent price shocks. Prices are martingales because shocks to demand are met by an immediate change in quantity, which is then transmitted to all forward prices through the impact on reserves.¹⁴

¹³ This extension of Hotelling's (1931) result is also noted in Weinstein and Zeckhauser (1975) and Pindyck (1980).

¹⁴ This effect is analogous to the permanent income hypothesis in which transitory income shocks are capitalized into permanent increases in consumption. However, in our setting a transitory demand shock is "capitalized" through its impact on reserves and hence the permanent price factor, but the instantaneous consumer surplus, our analogue of consumption, reverses itself as the transitory demand factor declines.

Given the equilibrium price function (14), it is easy to characterize the variance of both spot and forward prices.

PROPOSITION 5: *At any point in time the conditional variance of next period's spot price is given by*

$$\text{Var}_t(p_{t+1}) = \frac{e^{2a} \text{Var}(e^{y_{t+1}})}{R_{t+1}^2} \quad (16)$$

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

$$\text{Var}_t(\log p_{t+u}) = \sigma_y^2 + (u - t)\sigma_\eta^2, \quad (17)$$

where σ_y and σ_η are constants.

Proof: See the Appendix.

Remember that $R_{t+1} = R_t - q_t$ is in the information set at time t . The first part of the proposition shows 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 that case, the variance of next period's price would be $e^{2a} \text{Var}(e^{y_{t+1}})/q_t^2$, which is clearly higher than in equation (16) since current production is much lower than the total remaining reserves. The second part of Proposition 5 shows that supply responses cause the variance of the log of the future spot price to be linear in the holding period. This again reflects the fact that equilibrium prices have only permanent components.

II. The Resource Extraction Problem with Adjustment Costs

We now introduce and analyze a model in which adjustments to production are costly.¹⁵ As we will show, this modification causes stochastic resource prices to endogenously exhibit both temporary and permanent factors, which is

¹⁵ This is a different type of friction than that considered in Litzenberger and Rabinowitz (1995). In their model, producers are not able to extract all of an oil well's reserves at an arbitrary point in time. That is, although some portion of the reserves can be extracted at will, there is effectively an upper bound on production rates that ensures all wells will have some productive reserves available in the future (footnote 12, page 1523). This assumption is key to their backwardation result (Theorem 1 in their paper), which states that the amount of weak backwardation in forward prices is equal to the value of a put option on oil with a strike price equal to the extraction costs of the marginal producer. It is instructive to consider what happens in their setting absent the production constraint. If infra-marginal producers can extract all their reserves at any given point in time (i.e., no production frictions), then the future price of oil will be bounded below by the marginal producer's extraction costs. In this case, the relevant put option price is zero and forward prices will not exhibit backwardation.

consistent with the empirical findings of Schwartz and Smith (2000).¹⁶ Finally, our model with adjustment costs generates stochastic volatility that is related to the slope of the forward curve, an empirically relevant feature that is not currently present in reduced form pricing models.

A. The Economy

The economy is defined in continuous time with an infinite horizon. Instantaneous borrowing and lending are possible at a constant interest rate r . There is a finite reserve R_0 of a commodity, owned by a continuum of price-taking producers, and an inexhaustible supply of a substitute good. Once extracted, we assume that the commodity cannot be stored. 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.¹⁷

We also assume that producers incur a cost when aggregate production rates increase but not when they decrease. Although the study of more general adjustment costs is possible, we assume that this cost is proportional to the magnitude of the increase in production over its historical average, that is,

$$A(q_t; z_t) = \delta \max\{q_t - z_t, 0\} \equiv \delta(q_t - z_t)^+, \quad (18)$$

where δ is a constant, q_t is the chosen aggregate production rate, and z_t is the historic weighted average production rate,

$$z_t = \phi \int_{-\infty}^t e^{\phi(u-t)} q_u du,$$

with deterministic dynamics

$$dz_t = \phi(q_t - z_t) dt. \quad (19)$$

The form of this cost function is meant to capture the cost of developing new reserves in a reduced form.¹⁸

¹⁶ Our results emphasize the role of adjustment costs. Other channels can be used to generate temporary and permanent components, for example, mean-reversion in either marginal extraction costs or in risk premia. However, these alternatives will not give rise to the Markov-switching behavior we describe below.

¹⁷ Pindyck (1980) uses this specification of extraction costs in his model.

¹⁸ In the limit as $\phi \rightarrow -\infty$ this average approaches instantaneously lagged production. We have in mind an application in which reserves are equally distributed among several identical sources, some of which have been developed and all of which yield identical production flows. In this case, the economy-wide production rate can be increased either by increasing the number of developed sources or by increasing the production flow from each developed source. Hence, it is possible to interpret our specification of adjustment costs as either a cost associated with increasing the production flow from existing sources or with developing new sources. The latter interpretation is only loosely true, however, because in states in which developed reserves are optimally exhausted first, costs would need to be incurred in order to maintain current production.

The dynamics of the reserve process define how the reserves are depleted over time and can be expressed as

$$dR_t = -q_t dt, \quad (20)$$

where q_t is the production process and $R(0) = R_0$. Note that there is no exogenous uncertainty in this process.¹⁹ However, the reserves process will be random since production rates will depend on the stochastic demand state variable. Given a production policy, the time to exhaustion of the reserves, τ , is defined implicitly by

$$R_0 = \int_0^\tau q_t dt. \quad (21)$$

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

The (inverse) demand function for the commodity is assumed to be of the form $p_t = g(q_t; y_t)$. The parameter y_t characterizes intertemporal demand shocks that arrive according to the process

$$\frac{dy}{y} = \mu_y(y) dt + \sigma_y(y) dW_y. \quad (22)$$

We focus on the case in which this process is mean-reverting with a constant volatility, so that $\mu_y(y) = \kappa_y(\mu_y - \ln(y))$ and $\sigma_y(y) = \sigma_y$.

We assume that a substitute for the commodity exists with effectively infinite reserves. For instance, one might want to think of the commodity that we examine as oil and the substitute commodity as a high cost alternative to conventional reserves, such as oil shale. The substitute may not be currently produced because of its excessive marginal extraction costs, s_t . We specify a high price for the substitute good to ensure that the marginal value of reserves is sufficiently large to provide an incentive to delay extraction, as we have in mind a setting in which its predominant use will be in the distant future. Innovations arrive stochastically and affect this cost according to

$$\frac{ds}{s} = \mu_s(s) dt + \sigma_s(s) dW_s. \quad (23)$$

We focus on the case in which this process is a geometric Brownian motion with a constant drift, $\mu_s(s) = \mu_s$, and volatility, $\sigma_s(s) = \sigma_s$. This uncertainty may be driven by technological factors that reduce costs and, for example, environmental externalities that raise them.

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

¹⁹ Pindyck (1980) and Sundaresan (1984) analyze models with random reserve processes. These authors show that uncertain supply can give rise to backwardation in forward prices. We have chosen to focus instead on the effects of randomness in demand and in the marginal cost of the substitute good.

$$p(q; y, s) = \min \left(s, \frac{y}{q} \right), \quad (24)$$

where q is the current amount produced from conventional reserves.

B. Equilibrium in the Economy

Producers, who are assumed to be pricetakers, make output decisions that maximize the market value of their reserves net of the expected costs of extraction. Since the market value of reserves is a function of the equilibrium price, optimal production decisions and market clearing prices are 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 decisions.

To solve for equilibrium prices and quantities, we solve the related problem of a social planner who maximizes the discounted expected consumer plus producer surplus. At a given point in time this social surplus, SS , is defined as

$$SS(q_t; y_t, s_t, R_t, z_t) = \int_0^{q_t} p(x; y_t, s_t) dx - C(R_t)q_t - A(q_t; z_t) \quad (25)$$

and the social planner chooses production rates to maximize her discounted expected value,

$$V(R_t, y_t, s_t, z_t) = \max_{q_u \geq 0} E_t \int_t^\tau e^{-r(u-t)} SS(q_u; y_u, s_u, R_u, z_u) du, \quad (26)$$

subject to the dynamic equations for y , s , r , and z and where τ is a stopping time indicating the date at which reserves are fully depleted. Under conditions outlined in Dixit and Pindyck (1994), the solution to this problem coincides with production policies generated within a competitive equilibrium.²⁰ By casting the problem in terms of maximizing social welfare, traditional dynamic programming techniques can be applied to solve the problem numerically.

C. Computation and Calibration of the Equilibrium

The equilibrium, characterized by the solution to the constrained social planner's problem defined by equation (26), is conceptually straightforward to solve using the standard recursive techniques of dynamic programming. Specifically, given an initial estimate for the value function in any state, $V_0(R, y, s, z)$, we apply policy iteration techniques in order to converge to the fixed point that characterizes the production policy associated with the optimum (see, e.g., Puterman (1994)). Using 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.

²⁰ See also Weinstein and Zeckhauser (1975).

Forward prices and volatilities may be determined from state-dependent simulations of spot price paths. Applying the definition in equation (3), cross-sectional averages of the simulated future spot prices provide estimates of forward prices. Note that forward term structures are computed in this manner. In addition, and as is standard, the term structure of volatility is defined by

$$\text{TSOV}(u) \equiv \sqrt{\frac{\text{Var}_t(\log(p_{t+u}))}{u}}. \quad (27)$$

We calculate this function, again using simulated data, by averaging the squared differences between realized future spot prices and the associated forward price.

Although no complex theoretical issues arise in solving for the equilibrium, considerable practical problems must be addressed to numerically implement the solution due to the fact that our problem has four state variables (R , y , s , z), and one continuous choice variable, the production rate. The Appendix describes how we deal with the “Curse of Dimensionality” and provides details on our numerical technique.

To parameterize the model we proceed as follows.²¹ First, our model implies a region in which quantities are constant so that price dynamics exactly mimic those of the demand variable y . We therefore choose a rate of mean-reversion for demand, κ_y , and an instantaneous variance, σ_y , that approximate those reported for resource prices in the empirical literature (see, e.g., Casassus and Collin-Dufresne (2005)). We also choose the mean level to which (log) demand mean-reverts, μ_y , to reflect prices consistent with a commodity like oil. Second, since we have in mind an application such that the use of the substitute good is reserved for the distant future, there is little directly measurable evidence on which to base its calibration. We set its drift, μ_s , to zero and its diffusion, σ_s , to 5% per year. Finally, given the choice of the risk-free rate, we set the weight on historic production and the cost of increasing production to generate futures backwardation and contango roughly consistent with what is empirically observed. Table I summarizes these parameter choices.

D. Optimal Production with Adjustment Costs

In this subsection, we utilize our numerically solved model and analytically derived expressions to demonstrate the properties of endogenous supply responses when adjustment costs are present. This analysis leads to empirically relevant predictions regarding the dynamics of resource prices which, in turn, affect values of observable financial derivatives (e.g., futures and options prices) and real assets (e.g., natural gas wells).

²¹ We would ideally estimate the model's parameters using the Simulated Method of Moments. In light of the computational burden, however, we are forced to adopt an efficient calibration scheme that does not require repeated evaluations of the model's empirically relevant moments.

Table I
Parameter Values for the Base Case

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
Weight on historic production	ϕ	1.00
Cost of increasing production	δ	0.50
Extraction cost	C	0.00

We begin the analysis with the Hamilton–Jacobi–Bellman (HJB) equation for the social planner’s problem, which characterizes the value of the resource, V

$$rV = \max_q SS(q) - qV_R + \phi(q - z)V_z + \mu_y V_y + 1/2\sigma_y^2 V_{yy} + \mu_S V_S + 1/2\sigma_S^2 V_{SS}. \quad (28)$$

Dependencies on the state (R_t, y_t, s_t, z_t) are suppressed to enhance readability and subscripts denote partial derivatives.

We summarize necessary conditions for an optimum in the following proposition

PROPOSITION 6: *At each point in the state space one of the following three conditions will hold:*

(a) *Output satisfies $q_t < z_t$, with*

$$p(q_t; y_t, s_t) = V_R(R_t, y_t, s_t, z_t) - \phi V_z(R_t, y_t, s_t, z_t). \quad (29)$$

(b) *Output satisfies $q_t > z_t$, with*

$$p(q_t; y_t, s_t) - \delta = V_R(R_t, y_t, s_t, z_t) - \phi V_z(R_t, y_t, s_t, z_t). \quad (30)$$

(c) *Output satisfies $q_t = z_t$, with*

$$p(z_t; y_t, s_t) - \delta < V_R(R_t, y_t, s_t, z_t) - \phi V_z(R_t, y_t, s_t, z_t) < p(z_t; y_t, s_t). \quad (31)$$

Proof: The proof follows from differentiating the HJB equation in (28) to obtain necessary conditions for optimal production. Q.E.D.

Figure 2 illustrates this proposition under the parameterization in Table I. The downward-sloping discontinuous solid line represents (net) price as a




Figure 2. First-order conditions for optimal production. This figure illustrates the effect of changing the exogenous state variables on the first-order conditions specified in Proposition 6.

function of output quantity and the upward-sloping curve $V_R - \phi V_z$ represents the marginal cost of output as a function of its historical average. If current demand is low (see the dashed curve labeled “ $p(q)$ when y is low”), then production is reduced relative to its historic average z_t and the first-order condition specified in equation (29) is in effect. In this case, the marginal benefit of producing one unit of the resource is its price and the first-order condition equates this with the marginal cost ($V_R - \phi V_z$), which has two components that relate to the effect of production on the state variables R and z . (These mechanics are illustrated in the figure by the arrows originating at $(z_t, V_R - \phi V_z)$ pointing left and down.) Alternatively, if current demand is sufficiently high (see the dashed curve labeled “ $p(q) - \delta$ when y is high”), then production is increased, which implies that the first-order condition specified in equation (30) must be satisfied. Adjustment costs are incurred in these states so that the marginal benefit of producing one unit is the price *less* the adjustment cost. (These mechanics are illustrated in the figure by the arrows pointing right and down from the point $(z_t, V_R - \phi V_z)$.) Finally, at intermediate levels of demand, the state variables y and s may be in a region described by the inequalities (31). Within this region, production is set equal to z_t since the benefit of producing at a lower rate is

high relative to the implicit cost, and the benefit of producing at a higher rate is too small.

The form of the optimal production policy and, in particular, the presence of a no-response region has important implications for output and price dynamics, which translate into predictions for the state dependence of forward prices and price volatility. In contrast to the models without adjustment costs, prices are expected to grow at the riskless rate only in states in which production flexibility has an economically insignificant impact on the potential of incurring future adjustment costs. This may occur, for example, when current output is significantly below its historic average, so that the term V_z from equation (29) is small. In such states, this first-order condition equates prices with the marginal value of reserves, just as in the case for the models analyzed in Section I.

Adjustment costs thus give rise to interesting state dependencies in the level and shape of spot and forward prices. Furthermore, because they endogenously restrict production flexibility in certain states, adjustment costs also affect the dynamics of price volatility. We explore these implications in the following two subsections, which undertake a numerical analysis of the equilibrium and then analyze the model's time-series properties by utilizing impulse response functions.

E. A Numerical Analysis of the Equilibrium

We begin by demonstrating that equilibrium forward prices are qualitatively consistent with the empirical specification of Schwartz and Smith (2000) under our base-case calibration.

OBSERVATION 1 (Forward Curves): *The forward curves in the economy can be in backwardation or in contango (see Figure 3).*

The forward curves are in backwardation or contango depending on whether the demand shock process is above or below its long-run mean. Backwardation occurs because producers are (optimally) reluctant to increase output in some high-demand states. A less obvious effect occurs because producers also foresee that reducing current production when demand is low will increase the possibility of incurring adjustment costs if high demand is realized in the future, so in these states forward curves may be in contango. The result is that equilibrium prices may inherit some of the properties of the exogenous demand shock, a prediction that contrasts with those made by models with flexible production.

OBSERVATION 2 (Reserve Levels): *All forward prices rise as reserves are consumed (see Figure 4, Panel A).*

Intuitively, as reserves are consumed we would expect to see the level of prices increase. This is indeed the case as shown in Figure 4, Panel A, where we depict




Figure 3. 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.

forward curves at high and low reserve levels. Notice that prices at both the short and long ends of the forward curve are higher when reserves are low.

OBSERVATION 3 (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 4, Panel B).*

This observation is consistent with 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. 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. 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.

We can further clarify the dynamics of the forward curves if we compare the spot price process to the forward price process. When adjustment




Figure 4. Analysis of changes in the level of reserves and interest rates. Panel A shows that when reserves drop, forward prices rise. Panel B shows that when interest rates decrease, forward prices rise.

costs are present, the spot price process may be considerably more volatile than the forward price process, indicating that prices have a mean-reverting tendency.

OBSERVATION 4 (Term Structure of Volatility): *The term structure of volatility is downward sloping at short-to-intermediate horizons (see Figure 3) and upward sloping at very long horizons (see Figure 5).*

The reason for the higher short-run volatility is that current supply responses are constrained and hence exogenous shocks cause increased volatility at the short end of the curve. At intermediate horizons the curve exhibits lower volatility since the effect of exogenous shocks is dampened by producers' supply responses. At very long horizons, when reserve levels are likely to be low, the volatility of the future price of the substitute good drives the term structure of volatility. In the limit, spot price volatilities rise to the volatility of the marginal




Figure 5. Analysis of the long-run sensitivity of the term structure of volatility to the volatility of the alternative technology. This figure shows that term structure of volatility rises as reserves approach exhaustion.

cost of the substitute good, provided the volatility of the substitute good is sufficiently high.²²

OBSERVATION 5 (Demand Volatility): *A decrease in demand volatility has a small effect on forward prices and causes price volatilities to decrease (See Figure 6).*

In theory, price levels depend on exogenous demand volatility (as shown in Section I). With the current parameters the magnitude of this effect is small. Indeed, Panel A shows that forward prices are insensitive to a change in demand volatility from 15% to 10% per year. There is, however, a direct and intuitive effect on the term structure of volatility as is illustrated in Panel B.

OBSERVATION 6 (Volatility of Alternative Technology): *A decrease in the volatility of the alternative technology has a small effect on forward prices and causes the long-run price volatility to decrease (see Figure 7).*

²² As we show in Observation 6, our model generates an (inverted) humped term structure of volatility even when the substitute good's volatility is reduced.




Figure 6. Analysis of a change in the volatility of the demand shock. Panel A shows that forward prices do not change when the volatility of the demand shock decreases. Panel B shows that lower demand shock volatilities result in lower price volatilities.

Panel A shows that forward prices are insensitive to a change in the volatility from 5% to 2% per year. In contrast, as illustrated in Panel B, the long-maturity forward contract volatilities are sensitive to this parameter. Just like the base case, as conventional reserves are exhausted, the alternative source becomes more important and the term structure rises. However, with less uncertainty in the price at which this alternative will become available, there is a smaller long-run rise in the term structure of volatility.

F. The Time-Series Properties of Prices

To improve our understanding of the mechanics underlying the model, we study quantity and price dynamics by applying one-time shocks to the state variables and considering their impact over time. This analysis provides insights into the permanent versus temporary components of these shocks and thereby sharpens our predictions about the dynamics of forward curves. Our analysis also highlights how state variables influence price volatility under




Figure 7. Analysis of a change in the volatility of the alternative technology. Panel A shows that forward prices do not change when the volatility of the alternative technology decreases. Panel B shows that lower alternative technology volatilities result in lower price volatilities.

three different regimes. In the first regime production is flexible and costless (as described by equation (29)), in the second, production is flexible and adjustment costs are incurred (as described by equation (30)), and in the third, production is sticky (as described by equation (31)).

F.1. The Impulse Response Function for y

To illustrate the effects of switching across the model's three regimes, we choose the steady state μ_y as a starting value for y_t , and set the other three state variables (R , s , z) such that the system is within the region defined by equation (31) where production is unresponsive to small shocks.²³ We focus on the impact of an increase in y . This variable mean-reverts, so it will drift down following such a shock; however, since the inverse demand curve is directly

²³ It may be helpful to refer to Figure 2 while working through this subsection and the next. Impulse response functions are qualitatively similar to what is described here when historical production is high or low relative to current production.

proportional to y , it will shift up and then drift down. While the improvement in current demand conditions provides an incentive to increase production, to understand the response we must also consider the change in the marginal value of reserves and historical production, which is the right-hand side of equations (29) and (30). Here, we must rely on numerical results to determine the impact, since the marginal value of R will increase when demand rises, but so will the marginal value of z and intuition alone cannot predict which effect will dominate.

To undertake this exercise, we solve for the optimal policy using the procedure described in Section II.C. Next, using the numeric output linking the state space to the optimal policy, we identify specific points at which to perform the analysis.²⁴ We then trace out the path followed by (R, y, s, z) when no shocks are applied to the dynamic system, and record the associated time series for optimal quantities and prices, (q_t, p_t) . Finally, we apply a one-time shock to y ; observe the new values, (q'_t, p'_t) , generated by the procedure; and represent impulse response functions as the difference between the two paths.

Figure 8 presents two such impulse response functions following small and large increases in y . The top panel traces the change in quantity resulting from the shock and the bottom panel plots the impulse response of prices. The dashed line applies for shocks to y that are relatively small (0.05%). In this case, no change in output is required and the necessary conditions in inequality (31) continue to hold. Prices temporarily rise, due to the immediate shift in demand, but then fall, as y reverts back to its mean. In this sense, prices are locally mean-reverting.²⁵

More interesting mechanics underlie the response to larger shocks (0.5%), illustrated by the solid lines in Figure 8. In this case, the first-order condition in equation (30) will determine the optimal amount by which production increases after the shock is applied and the immediate direct effect of the increase in y is dampened.²⁶

Now consider the impact that remains after some discrete amount of time when z will have increased, in accordance with its dynamic equation (19). Optimal production at this point will be above its pre-impulse level and there will have been an increase in the state variable z . Thus, an innovation in the temporary demand variable y can imply an upward shift in quantities, q , and a downward shift in prices, p , even when the demand state y has returned to its long-run mean.²⁷ Negative correlation between short- and

²⁴ Specifically, we set $y_t = 40$, $s = 40$, $z_t = 1.8$, and $R_t = 18$.

²⁵ In this case the temporary price factor responds to small y shocks, but the permanent price factor does not, resulting in a local correlation close to zero.

²⁶ In addition to "local" mean-reversion, prices will also revert in response to an accumulation of shocks, in which case a production response occurs and prices react slowly to the direct effect of the y shock. The mechanics of this mean-reversion differ from "local" mean-reversion: The first is driven by a production response, while the latter arises due to the absence of an immediate response.

²⁷ Approximately 90% of the initial impulse to y dissipates over the 2-year impulse horizon depicted in Figure 8.




Figure 8. Impulse response of quantity and price to a shock in y . This figure shows the effect of an increase in the exogenous state variable y . The top panel illustrates the differential impact on quantities and the bottom panel illustrates the impact on prices, where dashed (solid) lines apply to an increase of 0.05% (0.5%).

long-term price factors may be offset, however, because higher depletion rates result in lower eventual reserves, which causes a permanent upward shift in prices.²⁸

Note that output quantities initially overreact, rising dramatically and then subsequently falling, and prices initially underreact to the shock. This effect is partially due to the incentive to minimize adjustment costs. Recall that these costs are incurred only when quantities are above their historical average, which follows current production with a lag. A cost-efficient way to respond to the shock is to increase production, q , above its historic average for a short

²⁸ The nonmonotonicity in the impulse response function plotted in Figure 8 arises because the reference quantity q is associated with a smaller level of historical production z and hence leaves the no-response region earlier than the impacted quantity q' . This effect gives rise to an apparent increase in relative quantities beginning approximately at post-impulse time $t = 0.5$.

period, during which time adjustment costs are incurred, and then to allow the rate of production to fall to a new but higher level of z .²⁹

In sum, the analysis in this subsection identifies three principal implications. First, prices are locally mean-reverting in response to small y shocks. Second, temporary demand shocks that overcome the adjustment cost hurdle can cause more persistent changes in production. Finally, prices may initially underreact to temporary demand shocks.

F.2. The Impulse Response Function for s

The impact of changes to the state variable s can be best understood in light of its economic interpretation as a proxy for the costs of supplying a competing substitute commodity (e.g., s could be the marginal cost of manufacturing oil from tar sands). An increase in this variable will increase the marginal value of reserves by causing the transfer of production to states where prices were previously bounded by the lower value of s .

Figure 9 plots this response when lagged output z equals current production, the mechanics of which can be understood using equation (29). The increase in s has no direct impact on current demand, but there is an upward shift in the marginal value of reserves as expressed by V_R . This will lead to a decrease in current production to a new level below its long-run average, which causes the state variable z to drift down. Shocks to s are permanent, so z will also shift permanently downwards, which reinforces this effect. The net impact on future prices is an upward shift at all dates, but in contrast to the model with permanent shocks and flexible production outlined in Section I, this shift will not be parallel. This implies that part of the shock is incorporated into prices as a temporary increment and the remainder as a positively correlated permanent increment.³⁰

F.3. The Dynamics of Volatility

Our analysis of responses to exogenous shocks in the preceding subsection gives rise to an intuitive explanation of the dynamics of volatility. Consider first the effect of demand volatility induced by y . In the no-response region, small shocks to y directly translate into price volatility, since there is no offsetting quantity change. However, since production is fixed and s has no influence on

²⁹ Note also that the price impact of y is dampened considerably, but not made completely permanent as in Section I. Inspection of Figure 8 shows that backwardation does exist following a positive shock, when defined using long-run futures prices less the current spot price, since the overall effect of the shock is a greater increase in the current price than the future price.

³⁰ As can be seen in Figure 9, the shock to s has a relatively large impact on the permanent price factor and thus the impulse gives rise to contango in futures prices. In the case when lagged output is below current production, a greater portion of the shock is permanent, causing a decline in the correlation between permanent and transitory price factors. This can be seen with reference to Figure 2. An increase in s when q is above z will shift the function $V_R - \phi V_z$ permanently upwards, with an accompanying and immediate permanent drop in quantity.




Figure 9. Impulse response of quantity and price to a shock in s . This figure shows the effect of an increase in the exogenous state variable s . The top panel illustrates the differential impact on quantities and the bottom panel illustrates the impact on prices, where dashed (solid) lines apply to an increase of 0.05% (0.5%).

demand, small shocks to s are not directly translated into price shocks in this region. Hence, price volatility in this region reflects only the constant volatility of the state variable y .

Volatility dynamics are considerably more interesting when the state variables y and s are outside the no-response regime. In particular, increases and decreases in y are met by corresponding increases and decreases in q , thereby dampening the effect of y on price volatility relative to the no-response regime. Quantity adjustments in response to the state variable s are transmitted to prices in this region, however, and this gives rise to a second source of price volatility. This response arises from changes in the marginal value of reserves and historical production as reflected in the right-hand sides of the equations in Proposition 6.

To summarize, the resource is produced at a constant rate within the no-adjustment region, the location of which depends on historic production decisions. If the state variable s hits a critical lower (upper) boundary, where the slope of the forward curve is negative (positive), production begins to vary and

Table II
Term Structure Slope and Volatility

This table presents the relationship between the slope of the term structure and volatility in the base-case model. We examine six points in the state space that differ in the amount of contango and backwardation that they exhibit. Backwardation and contango are measured by the percentage difference between the 12-month forward price and the spot price. For the six points in the table, volatility is calculated as the variance of future prices (approximately one month out) divided by the square root of time to maturity. The annualized volatilities are reported in Column 2, stated in units of percent per year. The table illustrates that the model can generate high price volatility, either when the forward curve is in contango or when it is backwardated.

12-Month Forward Slope	Volatility
−1.09	12.85
−0.95	10.96
0.01	0.39
0.92	9.40
1.88	12.49

the system moves into a region where prices respond to both s and y shocks. (Changes in the state variable y can also give rise to this behavior.) The production policy therefore gives rise to volatility behavior similar to that of a Markov switching model, but the slope of the forward curve provides information about the average level of volatility. Specifically, when the forward curve is steeply upward or downward sloping, volatility should be higher than when it is flat.

OBSERVATION 7 (U-shaped relationship between the slope of the forward curve and spot price volatility): *Volatility is stochastic. Specifically, price changes are relatively more volatile when the forward curve is backwardated or in contango (see Table II).*

To illustrate that the model can deliver this behavior, Table II presents the relationship between the slope of the term structure and volatility in the base-case model. We examine six points in the state space, each with different amounts of contango and backwardation. The degree of backwardation or contango is measured as the percentage difference between the 12-month forward price and the spot price. For the six points in the table, volatility is calculated as the variance of future prices (approximately 1 month out) divided by the square root of time to maturity. The annualized volatilities are reported in Column 2, stated in units of percent per year. The table illustrates that the model can generate high price volatility, either when the forward curve is in contango or when it is backwardated.

III. Empirical Evidence and Implications for Option Pricing

The models with flexible production in Section I show that state-dependent supply responses undo the effects of temporary demand shocks, implying that without frictions equilibrium prices have only permanent components. This

section examines the dynamics of oil and natural gas prices, a context in which temporary shocks have been shown to exist. The focus here, therefore, is on predictions of the model in the previous section, where the relevant friction is an adjustment cost incurred when production rates are increased beyond their historic average.

A. Stochastic Volatility in Oil and Natural Gas Prices

Daily observations of futures prices for NYMEX crude oil and natural gas futures contracts provide a basis for the analysis. We examine crude oil prices from April 1983 to June 2003 and natural gas prices from June 1990 to June 2003. Following a standard practice, we use the nearest-to-maturity futures to proxy for the spot price. Realized volatility for a month is calculated by summing the squared daily changes in log prices over the month. This results in two monthly time series of realized volatility, covering 243 months for crude oil and 157 months for natural gas.

The natural gas futures term structures and volatilities exhibit a seasonal variation. Further, they depend on short interest rates. We wish to focus on the relationship between futures prices and spot volatilities as predicted by a model with no seasonalities and constant riskless interest rates. We therefore remove these effects by first regressing the series on calendar-month dummy variables and on the 3-month T-bill rate. Although seasonal variation is much less evident in crude oil, we deseasonalize this data as well.

The prior literature provides evidence of stochastic volatility in crude oil and natural gas spot prices.³¹ We confirm these findings in Table III. Panel A reports results from estimating a GARCH model with monthly data and Panel B reports results from fitting an ARMA(1,1) model to the realized volatility series following the approach in Andersen et al. (2003).³² There is strong evidence supporting heteroskedasticity of the spot return series for both commodities. Lagged volatility and squared price innovations have a statistically significant impact on return innovations, as evidenced by the significant coefficients in the GARCH model. The ARMA model for realized volatility supports this finding, indicating a statistically significant role for lagged volatility.

To test our prediction that prices are more volatile when the futures term structure is either strongly backwardated or strongly contangoed, we regress realized volatility on the deseasonalized futures slope and its square.³³ As we

³¹ Duffie, Gray, and Hoang (1999) consider a variety of models for dynamic volatility in energy prices. Litzenberger and Rabinowitz (1995) provide evidence of a relation between crude oil backwardation and spot price volatility.

³² We use the GARCH specification of Bollerslev (1986), that is, $\sigma_t^2 = b_1 + b_2\epsilon_t^2 + b_3\sigma_{t-1}^2$, where $\sigma_t \equiv \text{Var}(\epsilon_{t+1})$.

³³ To check that the empirical significance of the squared slope term arises because of nonmonotonicity and not convexity, we verify that our result is robust to interacting the squared slope term with dummy variables for positive and negative slopes. The nonmonotonic relation is also confirmed by Kogan, Livdan, and Yaron (2004), who test our relationship by interacting the slope term with dummy variables for positive and negative slopes.

Table III
Stochastic Volatility of Spot Returns

This table reports evidence of stochastic volatility in deseasonalized oil and natural gas spot price returns. In Panel A, we report estimates for a GARCH (1,1) model. In Panel B, we report estimates for an ARMA (1,1) model of monthly realized spot return volatility. *t*-statistics are provided in parentheses.

	Crude Oil	Natural Gas
Panel A: GARCH (1,1) Model		
Constant	0.0003 (1.49)	0.0163 (2.39)
GARCH term	0.602 (9.69)	-0.019 (-0.07)
ARCH term	0.459 (4.54)	0.385 (2.12)
Panel B: ARMA (1,1) Model		
Constant	-0.026 (-0.20)	-0.004 (-0.06)
AR term	0.894 (29.94)	0.525 (4.70)
MA term	-0.331 (-4.48)	-0.008 (-0.05)

report in Table IV, we find a significantly positive coefficient on the second term, which is consistent with our model.³⁴ Newey–West *t*-statistics, using 12 lags, confirm that a significant relationship between spot volatility and the futures term structure exists for both crude oil and natural gas. Furthermore, the negative and statistically significant coefficients on the squared slope terms are consistent with the prediction of our model. Figure 10 illustrates this relationship in the data, confirming that volatilities are high, both in times of contango and backwardation.

B. Option Pricing

In prior sections, we explained why the endogenous price process from the equilibrium model has both temporary and permanent components. However, in contrast to the Schwartz and Smith (2000) two-factor model (SS2), our general equilibrium model generates a short-run price component whose drift is not always linear. As we saw, with adjustment costs, producers optimally increase (or decrease) production only when large demand shocks arrive; small

³⁴ Term structure slopes utilize the nearest and third-nearest contracts. This choice provides the longest possible time series, since both of these contracts have been trading simultaneously for the entire sample periods. Liquidity in the shorter maturity contracts is also high relative to, say, the 12-month contract, providing further justification for their use in our analysis.

Table IV
Futures Volatility and Term Structure Slope

Coefficient estimates for the regression of the realized volatility of log spot prices (deseasonalized) on the futures term structure slope (also deseasonalized) are provided. Newey-West *t*-statistics, utilizing 12 lags, are reported in parentheses.

	Crude Oil	Natural Gas
Constant	−0.119 (−1.32)	−0.033 (−0.59)
XSSlope	−1.44 (−1.18)	−0.777 (−2.62)
XSSlope2	66.470 (3.70)	2.702 (2.27)

demand shocks do not give rise to supply responses. As a result, the short-run component of the equilibrium price process has a drift that is “locally” linear since, if quantity supplied is constant, the drift of the endogenous price process reflects the linearity of the drift in the exogenous demand shock. However, overall the drift is nonlinear since large temporary demand shocks are met by nontrivial supply responses. An important consequence of the deviation from the SS2 dynamic specification is that the distribution of prices in our model will have tails that are truncated relative to those of the calibrated SS2 model. The resulting option prices predicted by the calibrated SS2 model will therefore be higher than the option prices generated by our model.³⁵

In this section, we calibrate the SS2 model to a time series of forward prices artificially generated by simulations from our model. We show that from a statistical perspective, the SS2 model does a good job of describing these forward prices. However, the SS2 model has an important source of misspecification that shows up when the calibrated model is used to price options.³⁶ Option prices predicted by the SS2 model are biased upwards from the “true” option prices generated by our model under the base-case parameterization described in Table I. We demonstrate this with the following experimental design. One hundred time series consisting of 3 years of weekly forward curves, each with 24 monthly contract prices, are artificially generated. This is done by simulating the demand and marginal cost state variables and then using information from the numerical solution to the equilibrium model to map these state variables to forward prices.³⁷ We then numerically calculate 10 option prices at the final calendar date, one maturing at the end of each of the 10 years following that

³⁵ Schwartz and Miltersen (1998) describe how to use information from a calibrated two-factor model to price options on commodities. The Fourier inversion approach of Duffie, Pan, and Singleton (1998) can also be applied to price options in this setting.

³⁶ It is well known that models consistent with the same forward curve can disagree on option prices. For instance, such pricing differences are observed within models of the short rate. Here our point is to clarify the source of such differences.

³⁷ Computational complexity limits the size of the experiment we can feasibly undertake.




Figure 10. The relationship between empirically measured log volatility and the slope of the forward curve. The figure illustrates that volatility is high when forward prices are either backwarddated or in contango for both crude oil (Panel A) and natural gas (Panel B).

date and struck at-the-money using the associated forward price. Next, we calibrate the SS2 model using the forward price data.³⁸ Finally, we calculate option prices from the SS2 model and compare them to those from our model.

³⁸ See Schwartz and Smith (2000) for the details of this calibration procedure.

Table V
Estimates of Parameters from the Schwartz
and Smith (2000) Calibration Procedure

We use simulated forward prices to estimate the parameters from the Schwartz and Smith (2000) two-factor model. Mean values and standard errors are derived from the distribution of the estimates from 100 independent simulations of 152 weekly forward prices extending out 24 months.

Parameter Name	Symbol	Value	
		Mean	Std. Error
Drift: long-run factor	μ	0.03	0.003
Diffusion: long-run factor	σ_{ξ}	0.01	0.002
Rate of mean-reversion: short-run factor	κ_{χ}	1.56	0.190
Diffusion: short-run factor	σ_{χ}	0.07	0.015
Correlation of factors	ρ	0.04	0.060

Table V describes the distribution of the parameter estimates resulting from the calibration exercise. The mean point estimates of the parameters are intuitively reasonable: The rate of mean-reversion of the temporary component ($\kappa = 1.56$) is close to that of the demand shock ($\kappa_y = 1.0$), the drift of the long-run component ($\mu_{\xi} = 0.03$) is close to the riskless interest rate ($r = 0.05$), the volatility of the short-run component ($\sigma_{\chi} = 0.07$) is somewhat less than that of the demand shock ($\sigma_y = 0.15$), and supply responses lead to a long-term component with low volatility ($\sigma_{\xi} = 0.01$). One can also see that the parameters in the SS2 model are measured very precisely; except for the parameter that measures the correlation between the long- and short-run factors, all the estimates are highly significant.³⁹ Based on these statistics alone, we would conclude that the SS2 model fits the simulated historical data very well. However, when the model is used out-of-sample the misspecification becomes very apparent.

Table VI demonstrates that the calibrated SS2 model overvalues a large class of options with maturities ranging from 1 to 10 years. The pattern of mispricing is nonmonotonic. For short-maturity options, the mispricing is low,⁴⁰ reflecting the fact that the price processes are well specified in terms of their “local” behavior. However, as the maturity of the options increases, the SS2 model overprices options by a significant amount, steadily increasing until it reaches a maximum at 5 years, then decreasing for options with maturities of 6 to 10 years. This occurs because, in the equilibrium model, when conventional reserves of the resource are depleted, the volatility of the marginal cost of the alternative technology becomes a more important component of the resource price process. Given the specification and parameterization of this price shock (see equation (23)), distant spot price volatilities are driven up. This effect is

³⁹ The unconditional permanent-transitory correlation might be close to zero because, as discussed in Section II.F, the conditional price correlation can be positive, zero, or negative. Additionally, the exogenous temporary and permanent state variables, y and s , are themselves uncorrelated.

⁴⁰ For maturities of less than a year, the two models yield very consistent option prices.

Table VI
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 (2000) model. In general, the option prices from the calibrated model are higher than the actual option prices.

Time to Maturity (years)	Median Option Price (\$)		Overpricing (%) Percentile		
	Equilibrium	Calibration	25	50	75
1	0.1321	0.1319	-12	1	14
2	0.1170	0.1351	2	15	28
3	0.0724	0.1370	63	86	104
4	0.0591	0.1384	113	135	158
5	0.0604	0.1396	114	138	161
6	0.0649	0.1414	104	123	141
7	0.0730	0.1447	86	101	117
8	0.0806	0.1477	73	186	99
9	0.0872	0.1506	64	74	87
10	0.0940	0.1526	55	65	77

absent from the SS2 model and, as a result, offsets the underpricing effect at very long horizons.

IV. Conclusion

This paper develops a general equilibrium model of exhaustible resource prices that extends the existing literature in a number of directions. Using several examples we show that uncertainty alone cannot explain the backwardation observed in resource markets. In fact, for resources with perfectly flexible production processes, forward prices rise at the rate of interest and temporary demand shocks are uniformly transmitted throughout the forward curve. In addition, in many of these settings the term structure of volatility is low and constant. Therefore, to explain the observed price behavior of commodities such as oil and gas, a supply adjustment cost is necessary. Although introducing this extra cost significantly complicates the analysis and necessitates a numerical solution, it generates endogenous price processes that can exhibit both backwardation in the forward curve and mean-reversion in spot prices.

As we mention in the Introduction, the model provides a practical framework for incorporating information about demand and supply functions into valuation problems. Our simulations suggest that this information is potentially quite important and can lead to very different option prices from Schwartz and Smith's (2000) reduced-form model, even when the Schwartz-Smith model provides a very good description of the process generating both forward and spot prices. In contrast to Litzenberger and Rabinowitz (1995), who take stochastic volatility as exogenous, our analysis shows that volatility of price changes

can arise as a natural consequence of the production decisions made by value-maximizing resource owners and that this volatility is related to the amount of backwardation as well as contango in prices. Our empirical analysis of oil and natural gas data is consistent with our unique explanation of stochastic volatility. Specifically, consistent with our model, the volatility of price changes for these commodities is higher when forward curves are both upward and downward sloping.

Two possible extensions of the model are left for future research. First, any serious attempt to apply this model to oil markets would require modeling the strategic interactions of producers with market power. Although extending the model to the case of a monopolist is straightforward, requiring only that we modify our objective function, important theoretical and computational issues arise in an oligopolistic market structure. In such a setting, production strategies depend on a producer's own reserve levels and the reserve levels of all other producers. This problem is especially challenging in the realistic case in which extraction costs vary among producers.

Second, storage is an important source of flexibility that we have ignored in our model. Routledge, Seppi, and Spatt (2000) show that storage has important implications for forward prices. If adjustment costs are small storage has little value in our context, since production flexibility is a perfect substitute for inventory. However, this is not the case when adjustment costs are high. Thus, it would be informative to analyze the joint optimal production and storage decisions in such cases.

Appendix

Proof of Proposition 4: We solve for the equilibrium by solving the social planner's problem as in Section II, with the simplifying assumption that the demand process is discrete. The following static variational problem characterizes the equilibrium:

$$\max_{q_t \geq 0} E \sum_{t=0}^{\infty} SS(q_t) \quad (\text{A1})$$

subject to

$$\sum_{t=0}^{\infty} q_t = R_0 \quad \text{a.s.}, \quad (\text{A2})$$

where SS is the social surplus function defined by equation (25). Equivalently, consider the unconstrained problem

$$\max_{q_t} \sum_{\omega} \left\{ \pi(\omega) \sum_t e^{-rt} \int_0^{q_t} p(x; \omega) dx - \lambda(\omega) \left[\sum_t q_t(\omega) - R_0 \right] \right\},$$

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

$$e^{-rt} p_t(\omega) \pi(\omega) = \lambda(\omega) \quad (\text{A3})$$

$$\sum_{t=0}^{\infty} q_t(\omega) = R_0, \quad \forall \omega. \quad (\text{A4})$$

Now, along any path ω , define $\hat{\lambda}(\omega) \equiv \frac{\lambda(\omega)}{\pi(\omega)}$. Thus, $\hat{\lambda}(\omega) = e^{-rt} p_t(\omega)$. Substitute $q_t = \frac{e^{a+y_t}}{p_t} = \frac{e^{a+y_t}}{\hat{\lambda}(\omega)e^{rt}}$ into equation (A4) and obtain

$$\hat{\lambda}(\omega) = \frac{\sum_{t=0}^{\infty} (e^{a+y_t}) e^{-rt}}{R_0}, \quad \forall \omega.$$

Let S be the set of ω such that $R_t = R$, $\varepsilon_t = \varepsilon$, and $p_t = p$ and use equation (A3) to sum over S , that is,

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

which implies

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

However, recall that $\hat{\lambda}(\omega) = e^{-rt} p_t(\omega)$. Thus, discounted prices are martingales.

To obtain the second part of the proposition note that

$$\begin{aligned} E_0(\hat{\lambda}) &= E_0 \left(\sum_{t=0}^{\infty} e^{-rt} \frac{e^{a+y_t}}{R_0} \right) \\ &= \left(e^{y_0} + \sum_{t=1}^{\infty} e^{-rt} \right) \frac{e^a}{R_0} \\ &= (k + e^{y_0}) \frac{e^a}{R_0}, \end{aligned}$$

where $k = \sum_{t=1}^{\infty} e^{-rt}$. Similarly, for any time t ,

$$E_t(\hat{\lambda}) = e^{-rt}(k + e^{y_t})\frac{e^a}{R_t}$$

and hence

$$p_t = (k + e^{y_t})\frac{e^a}{R_t}. \quad \text{Q.E.D.}$$

Proof of Proposition 5: Given the expression for price $p_t = (k + e^{y_t})\frac{e^a}{R_t}$, we clarify the form of the reserves process, R_t :

$$R_1 = R_0 - q_0 = R_0 \frac{ke^{-y_0}}{1 + ke^{-y_0}}.$$

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

$$p_t = \frac{e^{a+y_t}}{R_0 \prod_{i=1}^t \eta_i}.$$

This formula for prices allows computation of the term structure of volatility (17). Q.E.D.

The Numerical Solution of the Equilibrium with Adjustment Costs: The equilibrium is characterized by the solution to the constrained social planner's problem defined by equation (26). This problem is conceptually straightforward to solve using the standard recursive techniques of dynamic programming. The first step in solving these types of dynamic problems numerically is to form a discrete approximation to the continuous state space (see, for example, Kushner and Dupuis (1992)).⁴¹ In this case the (stochastic) processes consist of the two exogenous demand variables, namely, the reserve state variable and the lagged output, thus the state space is four dimensional. Each point in the state space can transition to eight "neighboring" points, two along each dimension.

Given the stochastic differential equations for the exogenous processes, the transition probabilities across the states are well known (see, for example, Kushner and Dupuis (1992)). For the lagged output and the reserve state

⁴¹ Another common but less direct approach for solving such dynamic programs involves iterating among approximating functions for the policy function or for the expectation arising in the first-order conditions (see, for example, Judd (1998)). In our setting, however, neither of these functions has a form that is known a priori. Furthermore, as we will show in later sections, they are not likely to be well approximated by a low-dimensional polynomial, as is required by these techniques.

variable, the transition probabilities depend upon the optimal production in a given state. Thus, we start by assigning a production level to each state. The transition probabilities for all processes including the lagged output and the reserve variable can then be computed from their stochastic differential equations. Corresponding to this initial production assignment we compute the initial value at each state, $V(R, y, s, z)$. This value consists of the reward at each node plus a probability-weighted average of the value at the nodes with positive transition probabilities. The reward at each node is computed using the current production policy as an input into the consumer surplus.

Having completed the initial value assignment to the state space, we need to update the optimal production policy at each node. The sparse nature of the transition matrix mitigates the problems associated with the “Curse of Dimensionality,” although the computational and storage requirements are still considerable. At each point in the state space we use a first-order condition to determine the optimal production response at each node. The first-order condition solves for the updated value at each node corresponding to a production policy. Recall that this value consists of the reward at each node plus a probability-weighted average of the values at nodes to which there is a positive transition probability. The optimal policy is then a function of the partial derivatives of the value function along the four dimensions, so numerical gradients need to be computed along these dimensions. The resulting optimal production implies a new “value” at each node, as described above. We then apply policy 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 (1994)). Given the optimal production policy, equilibrium prices are determined as a function of the state variables, which makes it possible to describe equilibrium price dynamics by using the transition density of the resulting Markov chain.

To test our algorithm, we verify that the numeric results without adjustment costs are consistent with the results from Section I. In this case forward prices grow at the rate of interest and the term structure of volatility is flat. In addition, supply responses considerably dampen demand shocks, resulting in price volatilities that are an order of magnitude smaller than demand volatility. Finally, average production decreases with time and quantities are about as volatile as demand shocks, indicating that changes in demand are matched by changes in quantity supplied.

For our base case, the exogenous structural parameters are chosen such that the endogenous parameters match empirical moments. We have given considerable thought to using a formal generalized method of moments calibration of the structural model to a panel of futures price data. However, this approach requires the computation of conditional moments in the four-dimensional state space, which must be estimated numerically at each empirical observation of the futures curve in the cross-section and over time. Moreover, the conditional moments must be repeatedly estimated for multiple candidate optimal parameters, rendering the time costs prohibitive. In light of the outlined

computational complexity, our efforts at a formal estimation have not been successful.

REFERENCES

- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Paul Labys, 2003, Modelling and forecasting realized volatility, *Econometrica* 71, 579–625.
- Bessembinder, Hendrik, Jay Coughenour, Paul Seguin, and Margaret Smoller, 1995, Mean reversion in equilibrium asset prices: Evidence from the futures term structure, *Journal of Finance* 50, 361–375.
- Black, Fisher, 1976, The pricing of commodity contracts, *Journal of Financial Economics* 3, 167–179.
- Bollerslev, Tim, 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.
- Brennan, Michael, and Eduardo Schwartz, 1985, Evaluating natural resource investments, *Journal of Business* 20, 135–157.
- Caballero, Ricardo J., and Eduardo M. R. A. Engel, 1999, Explaining investment dynamics in U.S. manufacturing: A generalized (S, s) approach, *Econometrica* 67, 783–826.
- Casassus, Jaime, and Pierre Collin-Dufresne, 2005, Stochastic convenience yield implied from commodity futures and interest rates, *Journal of Finance* 60, 2283–2331.
- Casassus, Jaime, Pierre Collin-Dufresne, and Bryan R. Routledge, 2005, Equilibrium commodity prices with irreversible investment and non-linear technologies, Working paper, Carnegie Mellon University.
- Deaton, Angus, and Guy Laroque, 1996, Competitive storage and commodity price dynamics, *Journal of Political Economy* 104, 896–923.
- Dixit, Avinash K., and Robert S. Pindyck, 1994, *Investment Under Uncertainty* (Princeton University Press, Princeton, NJ).
- Duffie, Darrell, Stephen F. Gray, and Philip Hoang, 1999, Volatility in energy prices, in Robert Jameson, ed.: *Managing Energy Price Risk*, 2nd ed. (Risk Publications, London).
- Duffie, Darrell, Jun Pan, and Kenneth J. Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.
- 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* 60, 55–73.
- Fama, Eugene F., and Kenneth R. French, 1988, Business cycles and the behavior of metals prices, *Journal of Finance* 43, 1075–1093.
- Grossman, Sanford J., and Guy Laroque, 1990, Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods, *Econometrica* 58, 25–51.
- Hotelling, Harold, 1931, The economics of exhaustible resources, *Journal of Political Economy* 39, 137–175.
- Judd, Kenneth, 1998, *Numerical Methods in Economics* (MIT Press, Cambridge, MA).
- Kogan, Leonid, Dmitry Livdan, and Amir Yaron, 2004, Futures prices in a production economy with investment constraints, Working paper, Wharton School.
- Kushner, Harold J., and Paul G. Dupuis, 1992, *Numerical Methods for Stochastic Control Problems in Continuous Time* (Springer Verlag, New York, NY).
- Litzenberger, Robert H., and Nir Rabinowitz, 1995, Backwardation in oil futures markets: Theory and empirical evidence, *Journal of Finance* 50, 1517–1545.
- Ng, Victor K., and S. Craig Pirrong, 1994, Fundamentals and volatility: Storage, spreads and the dynamics of metals prices, *Journal of Business* 67, 203–230.
- Pindyck, Robert S., 1980, Uncertainty and exhaustible resource markets, *Journal of Political Economy* 88, 1203–1225.
- Pindyck, Robert S., 1999, The long-run evolution of energy prices, *Energy Journal* 20, 1–27.
- Puterman, Martin L., 1994, *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley and Sons, Inc., New York, NY).
- Routledge, Bryan R., Duane Seppi, and Chester Spatt, 2000, Equilibrium forward curves for commodities, *Journal of Finance* 55, 1297–1338.

- Samuelson, Paul A., 1965, Proof that properly anticipated prices fluctuate randomly, *Industrial Management Review* 6, 41–49.
- Scarf, H., 1960, The optimality of (S, s) policies in the dynamic inventory problem, in K. Arrow, S. Karlin, and P. Suppes, eds.: *Mathematical Methods in the Social Sciences* (Stanford University Press, Stanford, CA).
- Schwartz, Eduardo S., 1997, The stochastic behavior of commodity prices: Implications for valuation and hedging, *Journal of Finance* 52, 923–973.
- Schwartz, Eduardo S., and Kristian 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.
- Schwartz, Eduardo S., and James E. Smith, 2000, Short-term variations and long-term dynamics in commodity prices, *Management Science* 46, 893–911.
- Sundaresan, Suresh M., 1984, Equilibrium valuation of natural resources, *Journal of Business* 57, 493–518.
- Weinstein, Milton C., and Richard J. Zeckhauser, 1975, The optimal consumption of depletable natural resources, *Quarterly Journal of Economics* 92, 371–392.

Copyright of Journal of Finance is the property of Blackwell Publishing Limited and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.