THE JOURNAL OF FINANCE • VOL. LXIV, NO. 3 • JUNE 2009

Oil Futures Prices in a Production Economy with Investment Constraints

LEONID KOGAN, DMITRY LIVDAN, and AMIR YARON*

ABSTRACT

We document a new stylized fact, that the relationship between the volatility of oil futures prices and the slope of the forward curve is nonmonotone and has a V-shape. This pattern cannot be generated by standard models that emphasize storage. We develop an equilibrium model of oil production in which investment is irreversible and capacity constrained. Investment constraints affect firms' investment decisions and imply that the supply elasticity changes over time. Since demand shocks must be absorbed by changes in prices or changes in supply, time-varying supply elasticity results in time-varying volatility of futures prices. Estimating this model, we show it is quantitatively consistent with the V-shape relationship between the volatility of futures prices and the slope of the forward curve.

In recent years commodity markets have experienced dramatic growth in trading volume, the variety of contracts, and the range of underlying commodities. There also has been a great demand for derivative instruments utilizing operational contingencies embedded in delivery contracts. For all these reasons there is widespread interest in models for pricing and hedging commodity-linked contingent claims. Commodities offer a rich variety of empirical properties that make them strikingly different from stocks, bonds, and other conventional financial assets. Notable properties of commodity futures include: (i) commodity futures prices are often "backwardated" in that they decline with time-to-delivery (Litzenberger and Rabinowitz (1995)); (ii) spot and futures prices are mean-reverting for many commodities; (iii) commodity prices are strongly heteroscedastic (Duffie and Gray (1995)) and price volatility is correlated with the degree of backwardation (Ng and Pirrong (1994) and Litzenberger and

*Kogan is at the Sloan School of Management, Massachusetts Institute of Technology and NBER. Livdan is at the Haas School of Business, University of California, Berkeley. Yaron is at the Wharton School, University of Pennsylvania and NBER. We would like to thank two anonymous referees, Kerry Back, Darrell Duffie, Pierre Collin-Dufresne, Francis Longstaff, and Craig Pirrong for detailed comments. We also thank seminar participants at University of California, Berkeley, Northwestern University, Texas A&M University, Stanford University, the 2004 Western Finance Association meeting, the 2004 Society of Economic Dynamics meeting, and the 2004 European Econometric Society meeting for useful suggestions. We also thank Krishna Ramaswamy for providing us with the futures data and Jeffrey R. Currie and Michael Selman for discussion and materials on the oil industry. Financial support from the Rodney L. White Center for Financial Research at the Wharton School is gratefully acknowledged.

Rabinowitz (1995)); and (iv) unlike financial assets, many commodities have pronounced seasonalities in both price levels and volatilities.

The theory of storage of Kaldor (1939), Working (1948, 1949), and Telser (1958) has been the foundation of the theoretical explorations of futures/forward prices and convenience yields (value of the immediate ownership of the physical commodity). Based on this theory researchers have adopted two approaches to modelling commodity prices. The first approach is mainly statistical in nature and requires an exogenous specification of the "convenience yield" process for a commodity (e.g., Brennan and Schwartz (1985), Brennan (1991), and Schwartz (1997)). The second strand of the literature derives the price processes endogenously in an equilibrium valuation framework with competitive storage (e.g., Williams and Wright (1991), Deaton and Laroque (1992, 1996), Routledge, Seppi, and Spatt (2000)). The appealing aspect of this approach is its ability to link futures prices to the level of inventories and hence derive additional testable restrictions on the price processes.

From a theoretical perspective the models based on competitive storage ignore the production side of the economy, and consequently they suffer from an important limitation. Inventory dynamics have little if any impact on the long-run properties of commodity prices, which in such models are driven mostly by the exogenously specified demand process. In particular, prices in such models tend to mean revert too fast relative to what is observed in the data (see Routledge et al. (2000)), and more importantly these models cannot address the rich dynamics of the term structure of return volatility.

In this paper we document an important new stylized fact regarding the property of the term structure of volatility of futures prices. We demonstrate that the relation between the volatility of futures prices and the slope of the forward curve (the basis) is nonmonotone and convex, that is, it has a V-shape. Specifically, conditional on a negatively sloped term structure, the relation between the volatility of futures prices and the slope of the forward curve is negative. On the other hand, conditional on a positively sloped term structure, the relation between the volatility and the basis is positive. This aspect of the data cannot be generated by basic models that emphasize storage, since such models imply a monotone relation between futures price volatility and the slope of the forward curve (see Routledge et al. (2000)).

In light of the aforementioned stylized fact, we explore an alternative model characterizing the mechanism of futures price formation. Futures prices are determined endogenously in an equilibrium production economy featuring constraints on investment, namely, irreversibility and a maximum investment rate. These investment constraints lead to investment triggers that affect firms' investment decisions, which in turn determine the dynamic properties of their output. More specifically, if the capital stock is much higher than its optimal level, given the current level of demand, firms find it optimal to postpone investment and the irreversibility constraint binds. On the other hand, when the capital stock is much lower than the optimal level, firms invest at the maximum possible rate and the investment rate constraint binds. In either case, the supply of the commodity is relatively inelastic and futures prices are relatively volatile. Since futures prices of longer-maturity contracts are less sensitive to the current value of the capital stock than the spot price, the slope of the forward curve tends to be large in absolute value when the capital stock is far away from its long-run average value. Thus, the absolute value of the slope of the term structure of futures prices is large exactly when the investment constraints are binding. Hence, the model predicts that the volatility of futures prices should exhibit a V-shape as a function of the slope of the term structure of futures prices. Stated differently, because of the binding constraints on investment, supply-elasticity of the commodity changes over time. Since demand shocks must be absorbed either by changes in prices or by changes in supply, time-varying supply elasticity results in time-varying volatility of futures prices. In our calibration below we show that the model can also generate these patterns in a manner that is quantitatively similar to the data.

There exists very little theoretical work investigating the pricing of commodities futures using a production economy framework. Grenadier (2002) and Novy-Marx (2007) also consider futures prices in a production economy, and discuss how the proximity of the state variable to the investment threshold governs the slope of the forward price curve. Both these papers, which include investment irreversibility, do not include an investment rate bound and thus cannot generate the price volatility predictions in backwardation. Casassus, Collin-Dufresne, and Routledge (2004) also analyze spot and futures oil prices in a general equilibrium production economy but with fixed investment costs and two goods. While also a production economy, the structure and implications of their model are quite different. A recent paper by Carlson, Khoker, and Titman (2006) also considers an equilibrium model with production. While we assume that oil reserves are infinite, their model emphasizes exhaustibility of oil reserves. Their model also gives rise to the nonmonotone relation between the futures price volatility and the slope of the forward curve. This implication is driven, as in our model, by adjustment costs in the production technology. This provides further evidence of the theoretical robustness of our finding—the exact structure of the model is not particularly important, as long as adjustments of production levels are limited in both directions.

The rest of the paper is organized as follows. In Section I, we describe our data set and document empirical properties of futures prices. Section II develops the theoretical model. In Section III, we study quantitative implications of the model. Section IV provides conclusions.

I. Empirical Analysis

We concentrate our empirical study on crude oil. Kogan, Livdan, and Yaron (2005) report qualitatively similar findings for heating oil and unleaded gasoline. We choose however to focus on the crude oil contract because: (i) it represents the most basic form of oil, for which the investment constraints we highlight seem to be the most relevant and (ii) the other contracts clearly use crude oil as an input and thus their analysis may require a more specific "downstream" industry specification. Our data consist of daily futures prices for the

NYMEX light sweet crude oil contract (CL) for the period from 1982 to 2000. Following previous work by Routledge et al. (2000), the data are sorted by contract horizon with the "1-month" contract being the contract with the earliest delivery date, the "2-month" contract having next-earliest delivery date, etc.[1] We consider contracts up to 12 months to delivery since liquidity and data availability are good for these horizons.[2] Given we are using daily data, our data set is sufficiently large, ranging from 2,500 to 3,500 data points across different maturities.

Instead of directly using futures prices, $P(t, T)$, we use daily percent changes, $R(t,T) = \frac{P(t,T)}{P(t-1,T)}$. Percent price changes are not as susceptible as price levels to seasonalities and trends, and therefore their volatility is more suitable for empirical analysis. We then proceed by constructing the term structure of the unconditional and conditional volatilities of daily percent changes on futures prices. In calculating conditional moments, we condition observations on whether the forward curve was in backwardation or in contango at the end of the previous trading day (based on the third-shortest and sixth-shortest maturity prices at that time). Figure 1 shows the conditional and unconditional daily volatilities for futures price percent changes. Unconditionally, the volatility of futures price changes declines with maturity, consistent with the Samuelson (1965) hypothesis. The behavior of crude oil (CL) contracts was previously studied by Routledge et al. (2000). We find, as they do, that the volatility of futures prices is higher when the forward curve is in backwardation. This has been interpreted as evidence in favor of the standard storage theories, emphasizing the effect of inventory stock-outs on price volatility.

Next, we study the patterns in volatility of futures prices in more detail. Specifically, we estimate a functional relation between the futures price volatility and the 1-day lagged slope of the forward curve. Following the definition of conditional sample moments, the time series of the slope of the forward curve, $SL(t)$, is constructed as a logarithm of the ratio of the futures price of the sixth-shortest maturity in months available on any day t, $P(t, 6)$, to the future price of the third-shortest maturity, $P(t, 3)$, available on the same day

$$SL(t) \equiv \ln\left[\frac{P(t-1,6)}{P(t-1,3)}\right]. \tag{1}$$

We use the demeaned slope,

$$\widetilde{SL}(t) \equiv SL(t) - \mathrm{E}\left[SL(t)\right], \tag{2}$$

[1] In our data set on any given calendar day there are several contracts available with different time to delivery measured in days. The difference in delivery times between these contracts is at least 32 days or more. We utilize the following procedure for converting delivery times to the monthly scale. For each contract we divide the number of days it has left to maturity by 30 (the average number of days in a month), and then round off the result. For days when a contract with time to delivery of less than 15 days is traded, we add 1 month to the contract horizon obtained using the above procedure for all contracts traded on such days. The data are then sorted into bins based on the contract horizon measured in months.

[2] We refer to this time to delivery as time to maturity throughout the paper.

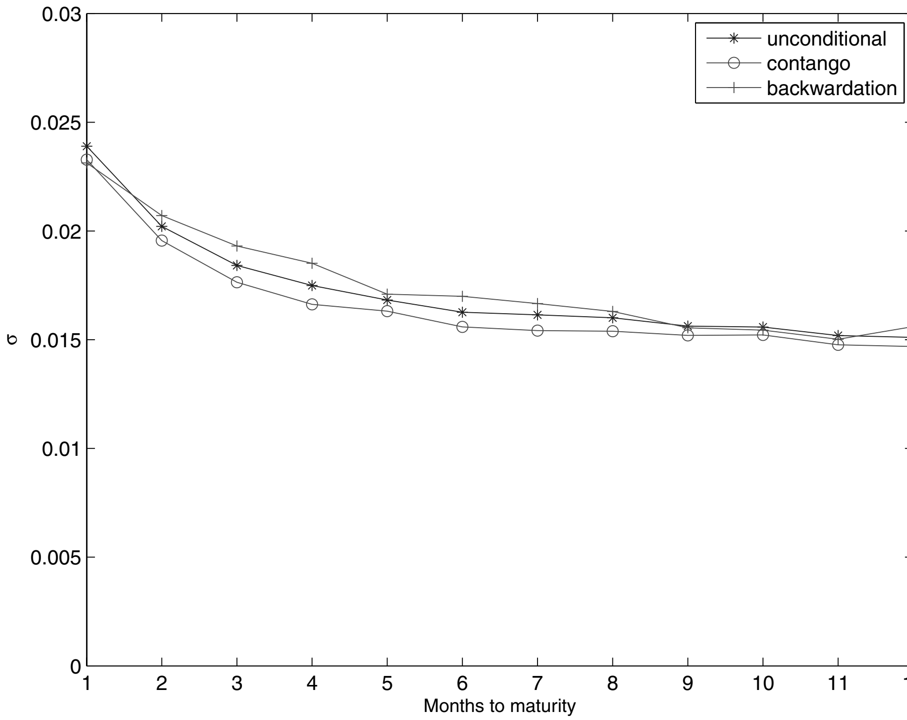

Figure 1. Term structure of volatility, crude oil futures. The data are daily percentage price changes on NYMEX crude oil (CL) futures from 1983 to 2000. σ denotes the standard deviation of daily percentage price changes. The time to maturity is defined as the number of months left until the delivery month. The unconditional standard deviation is constructed using the sample's first and second moments, while standard deviations conditional on backwardation and contango are conditioned on the shape of the forward curve one day prior.

in our analysis. We start by using the demeaned lagged slope as the only explanatory variable for realized volatility:

$$|R(t,T)| = \alpha_T + \beta_T \widetilde{SL}(t-1) + \varepsilon_T(t). \tag{3}$$

Since we are now estimating a different functional form, note that the relation (3) can potentially yield different information from that contained in Figure 1, which was obtained by simply splitting the sample based on the slope of the forward curve. The term structure of β_T as well as the corresponding t-statistics are shown in Figure 2. We also report these results in Table I for T equal to 1, 5, and 10 months. The negative sign of β_T for all times to maturity is a notable feature of these regressions. This result seems to be at odds with the relations shown in Figure 1, where volatility conditional on backwardation is for the most part higher than the unconditional volatility.

The apparent inconsistency becomes less puzzling in light of the intuition of the model we present below. In particular, our theoretical results motivate one to look for a nonmonotone relation between the volatility of future prices and the slope of their term structure. To do so we decompose the lagged demeaned

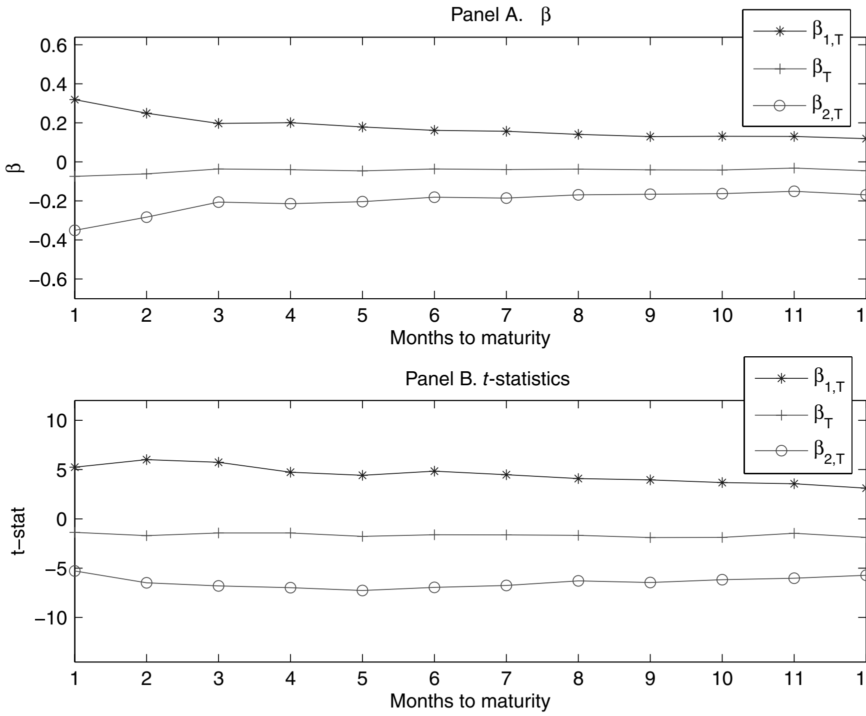

Figure 2. Conditional volatility, crude oil futures. The data are daily percentage price changes on NYMEX crude oil (CL) futures from 1983 to 2000, denoted by $R(t, T)$. Two different specifications are used to relate the instantaneous volatility of daily percentage price changes to the beginning-of-period slope of the forward curve defined as

$$SL(t) = \ln\left[\frac{P(t-1,6)}{P(t-1,3)}\right].$$

The first specification is

$$|R(t,T)| = \alpha_T + \beta_T \widetilde{SL}(t-1) + \varepsilon_T(t),$$

where $\widetilde{SL}(t)$ is the demeaned slope. The second specification decomposes the slope into positive and negative parts,

$$|R(t,T)| = \alpha_T + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t),$$

where $(X)^\pm$ denotes the positive (negative) part of X. The figure shows the estimates of all three coefficients (Panel A) and their respective t-statistics (Panel B) for different times to maturity. All t-statistics are White-adjusted.

slope into positive and negative parts and use them as separate explanatory variables (i.e., use a piecewise linear regression on the demeaned slope of the term structure),

$$|R(t,T)| = \alpha_T + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t), \tag{4}$$

where $(X)^\pm$ denotes the positive (negative) part of X. Figure 2 as well as Table I illustrate our results. Both $\beta_{1,T}$ and $\beta_{2,T}$ are statistically and economically significant for most maturities. More importantly, $\beta_{1,T}$ and $\beta_{2,T}$

Table I

Conditional Volatility, Crude Oil Futures

This table reports results for two different regressions. The data are daily percentage price changes on NYMEX crude oil (CL) futures from 1983 to 2000. The specifications in both panels are the same as in Figure 2. All results are reported for times to maturity equal to 1, 5 and 10 months. All t-statistics are White-adjusted for conditional heteroscedasticity.

	1 Month	5 Months	10 Months
	$\lvert R(t,T)\rvert = \alpha_T + \beta_T \widetilde{SL}(t-1) + \varepsilon_T(t)$		
β_T	−0.0743	−0.0462	−0.0419
(t-stat)	(−1.37)	(−1.77)	(−1.89)
R^2	0.0122	0.0111	0.0117
	$\lvert R(t,T)\rvert = \alpha_T + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t)$		
$\beta_{1,T}$	0.3191	0.1791	0.1308
(t-stat)	(5.25)	(4.42)	(3.68)
$\beta_{2,T}$	−0.3505	−0.2039	−0.1632
(t-stat)	(−5.29)	(−7.26)	(−6.17)
R^2	0.1320	0.1058	0.0790

differ in sign: $\beta_{1,T}$ is positive and $\beta_{2,T}$ is negative. Therefore, the relation between the volatility of futures prices and the slope of the term structure of prices is nonmonotone and has a V-shape: Conditional volatility declines as a function of the slope when the latter is negative, and increases when the latter is positive.[3]

One potential concern is that the piecewise linear regression may artificially lead our estimates to highlight the "V-shape" pattern we report. To allow for a more flexible form for the relationship between volatility and the slope we use a nonparametric regression (our specific implementation is based on Atkeson, Moore, and Schaal (1997)). Figure 3 shows the results of the nonparametric regression for horizon T equal to 2, 4, 6, and 8 months. For all maturities it reveals a clear nonmonotone V-shap relationship between the volatility of futures prices and the slope of the term structure of prices.

We perform several additional robustness checks. First, we estimate conditional variances instead of conditional volatility by using the square of daily price changes instead of their absolute value. We find that the conditional variance leads to very similar conclusions. In most cases, both $\beta_{1,T}$ and $\beta_{2,T}$ remain statistically significant. Next, we fit a GARCH(1,1) model to the daily percentage price changes of maturity T to obtain the fitted volatility time series $\sigma_{GARCH}(t,T)$ and use it as an independent variable in (4). The results are reported in Panel 1 of Table II and show that both $\beta_{1,T}$ and $\beta_{2,T}$ remain statistically significant. To avoid concerns about look-ahead bias, we also predict volatility based on a rolling GARCH(1,1), which we denote $\widehat{\sigma}_{GARCH}(t,T)$. It is the fitted value of a GARCH(1,1) model to percentage price changes based on

[3] An alternative parametric form would be quadratic. While the quadratic specification does not allow us to test for the conditional sign of the relationship between the slope of the forward curve and conditional volatility of futures returns, it provides a test for the convexity of that relationship. The estimated relationship (not reported here) shows that the second-order term is highly statistically significant.

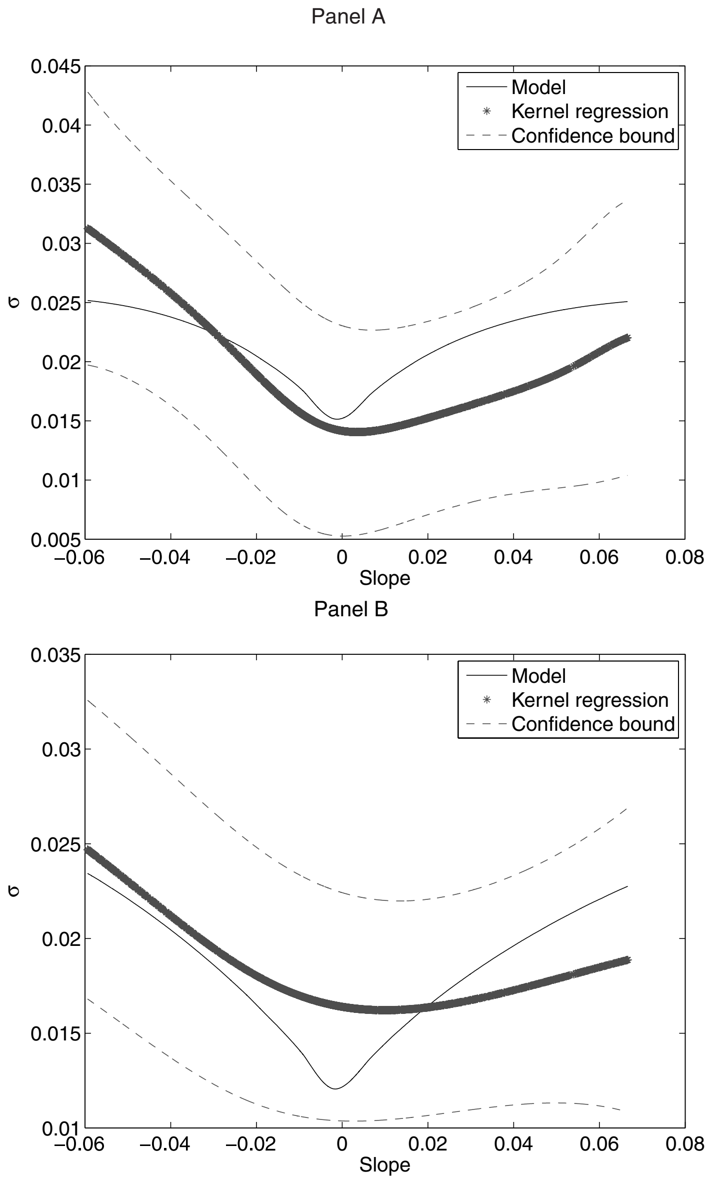

Figure 3. V-shape of volatility of futures prices, data and model output. Futures volatility is plotted as a function of the slope of the forward curve for different values of time to maturity T. Time to maturity is measured in months. In Panel A, $T = 2$; in Panel B, $T = 4$; in Panel C, $T = 6$; in Panel D, $T = 8$. We use Receptive Field Weighted Regression (RFWR), based on Atkeson, Moore, and Schaal (1997) to construct the functional form of volatility implied by the data. In both cases, data and model, volatility exhibits a V-shape pattern as a function of the slope. The slope is demeaned. Volatility is expressed in annual terms. Model parameters are given in Table III.

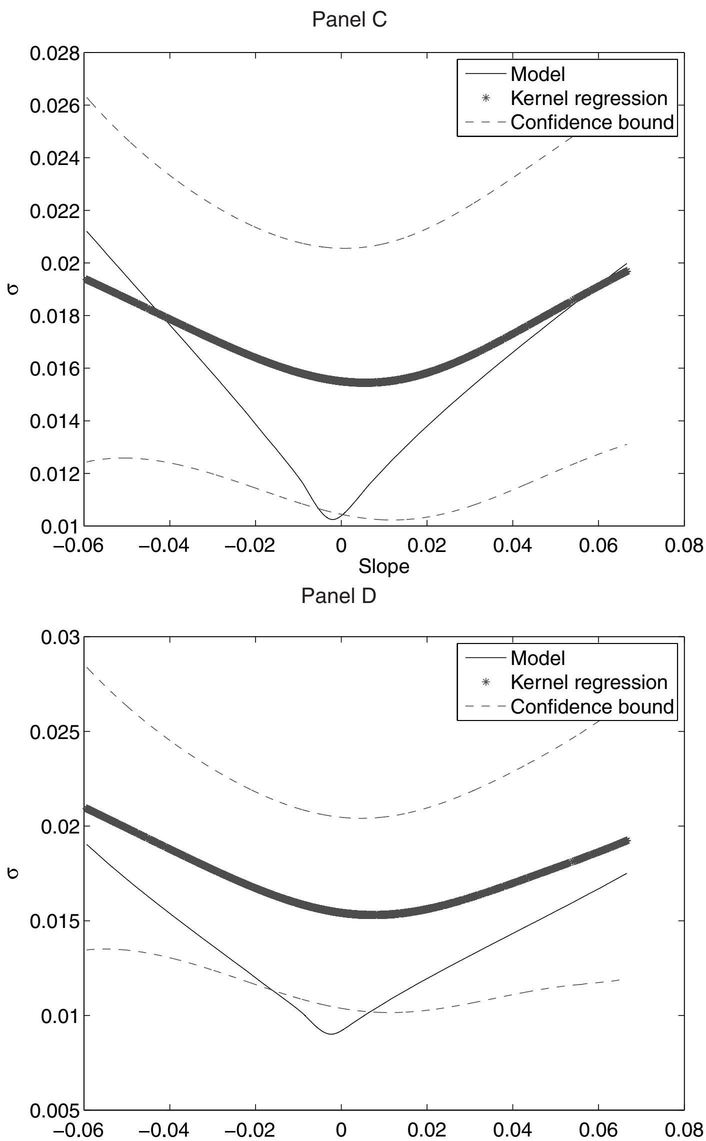

Figure 3. Continued

a sample that ends at $t-1$ with a window of 300 observations. We then use it in the regression

$$\widehat{\sigma}_{GARCH}(t,T) = \alpha_T + \beta_{G,T}\widehat{\sigma}_{GARCH}(t-1,T) + \beta_{1,T}(\widetilde{SL}(t-1))^{+} + \beta_{2,T}(\widetilde{SL}(t-1))^{-} + \varepsilon_T(t). \tag{5}$$

The results are reported in Panel 2 of Table II and show that both $\beta_{1,T}$ and $\beta_{2,T}$ still remain statistically significant.

Table II

GARCH Regressions, Crude Oil Futures

This table reports results for two different regressions. The data are daily returns percentage price changes on NYMEX crude oil (CL) futures from 1985 to 2000. The volatility time series $\sigma_{GARCH}(t, T)$ is obtained by fitting a GARCH(1,1) model to the daily percentage price changes of maturity T. The volatility time series $\widehat{\sigma}_{GARCH}$ is obtained by fitting a GARCH(1,1) model to the time series of daily returns percentage price changes of maturity T up to day $t-1$ and then using it to predict the volatility on day t, $\widehat{\sigma}_{GARCH}(t, T)$. The initial window is set to 300 days. All other variables are defined in the caption to Figure 2.

	1 Month	5 Months	10 Months
$\lvert R(t,T)\rvert = \alpha_T + \beta_{G,T}\sigma_{GARCH}(t,T) + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t)$			
$\beta_{G,T}$	0.5323	0.6145	0.6612
(t-stat)	(8.85)	(9.48)	(12.47)
$\beta_{1,T}$	0.1610	0.0823	0.0595
(t-stat)	(4.32)	(4.01)	(3.81)
$\beta_{2,T}$	−0.1605	−0.0796	−0.0599
(t-stat)	(−2.84)	(−2.92)	(−2.91)
R^2	0.2601	0.2605	0.2495
$\widehat{\sigma}_{GARCH}(t,T) = \alpha_T + \beta_{G,T}\widehat{\sigma}_{GARCH}(t-1,T) + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t)$			
$\beta_{G,T}$	0.9403	0.9748	0.9739
(t-stat)	(68.82)	(155.06)	(161.50)
$\beta_{1,T}$	0.0269	0.0079	0.0064
(t-stat)	(3.06)	(3.34)	(2.90)
$\beta_{2,T}$	−0.0228	−0.0060	−0.0053
(t-stat)	(−3.04)	(−2.23)	(−1.83)
R^2	0.9108	0.9627	0.9586

As a final robustness check we split our sample into pre– and post–Gulf War subsamples. We perform the same analysis as in the case of the full sample on the post–Gulf War subsample. We find the same V-shape in the relationship between the volatility of futures prices and the slope of the term structure of prices.

II. Model

In this section, we present our model for spot prices and derive futures prices.

A. Setup

We consider a continuous-time infinite-horizon economy. We focus on a competitive industry populated by a large number of identical firms using an identical production technology. Firms produce a nonstorable consumption good by means of a production function that exhibits constant returns to scale

$$Q_t = XK_t, \tag{6}$$

where K_t is capital and X is the productivity of capital, which is assumed to be constant. Without loss of generality, we assume below that $X = 1$. Our results can be easily adjusted to accommodate the case in which X is a stochastic process. We also abstract away from production costs.

Firms can adjust their capital stock according to

$$dK_t = (I_t - \delta K_t)\,dt, \tag{7}$$

where I_t is the investment rate and δ is the capital depreciation assumed to be a nonnegative constant. We assume the unit cost of capital is equal to one.

We next assume that investment is irreversible, that is, $I_t \geq 0$, and the rate of investment is bounded. Specifically,

$$I_t \in \left[0, \bar{\imath} K_t^A\right], \tag{8}$$

where K_t^A is the aggregate capital stock in the industry. This constraint implies that the higher the aggregate capital stock in the industry (or the higher the aggregate output rate), the better the investment opportunities faced by an individual firm. One can think of this as a learning-by-doing technology. These investment frictions give rise to nontrivial dynamic properties of futures prices. It is worth noting that one could derive the same functional form of price dynamics by assuming that investment opportunities depend on the firms' own capital stock. However, the above assumption significantly simplifies formal analysis of the model and leads to fewer restrictions on model parameters.

We do not explicitly model entry and exit in equilibrium. Our assumption of nonnegative investment rates effectively implies that there is no exit. We are assuming that all investment is done by existing firms, so there is no entry either. One could equivalently allow for entry into the industry, as long as the total amount of investment by old and new firms satisfies the constraint (8).[4]

Firms sell their output in the spot market at price S_t. We assume that financial markets are complete and the firms' objective is to maximize their market value, which in turn is given by

$$V_0 = \mathrm{E}_0\left[\int_0^\infty e^{-rt}(S_t Q_t - I_t)\,dt\right]. \tag{9}$$

We assume that the expected value is computed under the risk-neutral measure and the risk-free rate r is constant.

The consumers in the economy are represented by the demand curve

$$Q_t = Y_t S_t^{-\frac{1}{\gamma}}, \qquad Q_t \in (0, \infty), \tag{10}$$

where unexpected changes in Y_t represent demand shocks. We assume that under the risk-neutral measure Y_t follows a geometric Brownian motion process

$$\frac{dY_t}{Y_t} = \mu_Y\,dt + \sigma_Y\,dW_t. \tag{11}$$

We also assume that $\gamma > 1$. Results for the case $\gamma \leq 1$ are analogous.

[4] For example, other recent applications in finance using competitive industry equilibrium include Fries, Miller, and Perraudin (1997) and Miao (2005).

B. Equilibrium Investment and Prices

We adopt a standard definition of competitive equilibrium. Firms must choose an investment policy that maximizes their market value (9), taking the spot price of output and the dynamics of the aggregate capital stock in the industry as exogenous. The spot market must clear, that is, the aggregate output and the spot price must be related by (10). Finally, the dynamics of the aggregate capital stock in the industry are given by

$$dK_t^A = \left(I_t^A - \delta K_t^A\right) dt, \tag{12}$$

where I_t^A is the aggregate investment rate.

We guess what the equilibrium investment policy and associated price processes should be, and verify formally that firms' optimality conditions are satisfied and markets clear. The details of the solution are provided in Appendix A.

Intuitively, firms invest only when the net present value of profits generated by an additional unit of capital is positive. As it turns out, the spot price follows a univariate Markov process in equilibrium, and thus firms invest at the maximum possible rate when the spot price is above a certain threshold, and do not invest otherwise. Formally, we prove the following:

PROPOSITION 1: *A competitive equilibrium exists and the equilibrium investment policy is given by*

$$I_t^* = \begin{cases} \bar{i} K_t^A, & S_t \geq S^*, \\ 0, & S_t < S^*. \end{cases} \tag{13}$$

The investment threshold S^* is defined in Appendix A.

See Appendix A for the proof.

To make sure that the equilibrium exists and firm value is finite, we need to impose an additional nontrivial restriction on parameter values:

$$\frac{\sigma_Y^2 \gamma^2}{2} - \gamma \mu^+ - (r + \delta) < 0, \tag{14}$$

where we define $\mu^- = \delta + \mu_Y - \frac{1}{2}\sigma_Y^2$ and $\mu^+ = \bar{i} - \mu^-$. When calibrating the model, we impose the above restriction as a weak inequality and verify that it does not bind at the calibrated parameter values.

The risk-neutral dynamics of the spot price are of the very simple form

$$\frac{dS_t}{S_t} = \left(-\gamma(\bar{i} 1_{[S_t \geq S^*]} - \mu^-) + \frac{\gamma^2 \sigma_Y^2}{2}\right) dt + \gamma \sigma_Y dW_t, \tag{15}$$

where $1_{[\cdot]}$ is an indicator function. When the spot price is above the critical value S^*, it follows a geometric Brownian motion with drift $\gamma\mu^- + \frac{\gamma^2\sigma_Y^2}{2}$. When it is below S^*, the drift changes to $-\gamma(\bar{i} - \mu^-) + \frac{\gamma^2\sigma_Y^2}{2}$. As long as $0 < \mu^- < \bar{i}$,

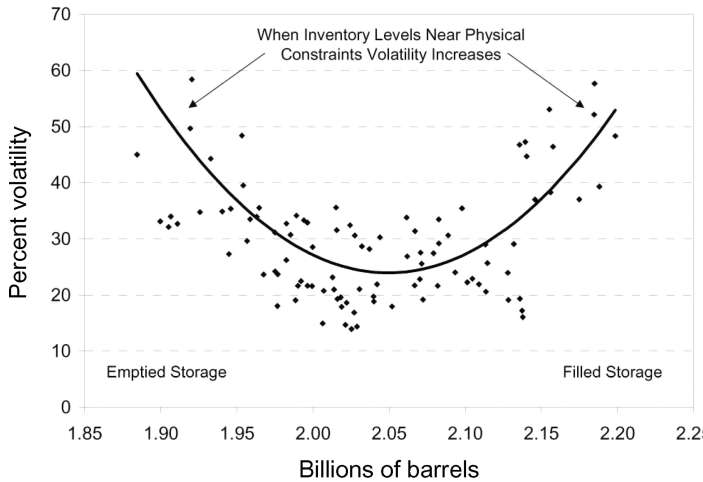

Figure 4. Volatility and supply capacity constraints. The figure gives actual WTI crude oil futures volatility versus expected end-of-March Atlantic Basin inventories. Source: Strongin, Currie, and Fleischmann (1999).

the spot price process has a stationary long-run distribution with density function

$$p(S) = \frac{2\mu^-(\bar{i} - \mu^-)}{\gamma \bar{i} \sigma_Y^2 S} \left(\frac{S}{S^*} \right)^{-\frac{2}{\gamma \sigma_Y^2}(\bar{i} 1_{[S \geq S^*]} - \mu^-)}. \tag{16}$$

The details of the derivation are provided in Appendix B.

Before continuing with estimation of the model, it is worth mapping our general investment constraint model into the features of the oil industry. Oil (Q) is the output produced using physical capital K (e.g., oil rigs, pipes, and tankers). Implicitly we are assuming there is an infinite supply of underground oil, and production is constrained by the existing capital stock K. This supply of capital and consequently of oil output leads to price fluctuations in response to demand shocks. Futures prices (volatility) depend on anticipated future production, which depends on the degree to which investment is constrained.

While our model is admittedly stark, it does capture many of the essential features of the investment and supply constraints in the oil industry. Figure 4 displays the relation between oil supply capacity and volatility as given in Strongin, Currie, and Fleischmann (1999). The essence of this figure is the V-shape relation between the volatility of the spot price and the level of inventories. This pattern is distinct from the one analyzed in our paper: While inventory levels are clearly important for short-run fluctuations in the spot market, their effect on futures prices of longer maturities is much weaker. Thus, a different mechanism must be responsible for the behavior of volatility of longer-maturity futures. However, the logic of time-varying supply elasticity applies to the pattern in Figure 4 as well, wherein instead of production constraints one must recognize natural physical constraints on storage levels.

Furthermore, market analysts seem to concentrate on two key features of the market, namely, the long-term and seasonal demand patterns and the supply features of this industry. Our model clearly focuses on the second of these two issues. In particular, investments in this industry are concentrated in several key facets of production: (i) basic extraction in the form of finding new fields and constructing, installing, and maintaining rigs, and (ii) expansion and improvements in the delivery process. Both of these types of investment take time, face capacity constraints, and constrain supply flexibility in this market—the exact channels that our model focuses on.[5]

C. Futures Prices

The futures contract is a claim on the good that is sold on the spot market at prevailing spot price S_t. The futures price is computed as the conditional expectation of the spot price under the risk-neutral measure:

$$P(t, T) = \mathrm{E}_t[S_{t+T}], \quad \forall T \geq 0, \tag{17}$$

where $P(t, T)$ denotes the price of a futures contract at time t with maturity date $\tau = t + T$. Since no analytical expression exists for the above expectation, we evaluate it numerically using a Markov chain approximation scheme. Figure 5 illustrates futures prices generated by the model.

III. Estimation and Numerical Simulation

In this section we study how well our model can replicate quantitatively the key features of the behavior of futures prices reported in Section I. Since our model is formulated under the risk-neutral probability measure, while the empirical observations are made under the "physical" probability measure, one has to make an explicit assumption about the relation between these two measures, that is, about the risk premium associated with the shock process dW_t. To keep our specification as simple as possible, we assume that the risk premium is constant, that is, the drift of the demand shock Y_t under the "physical" probability measure is equal to $\mu_Y + \lambda$, where λ is an additional parameter of the model. Clearly, one could achieve greater flexibility and better fit of the data by allowing for a time-varying risk premium process. This, however, is

[5] Analysts often describe the "cushion" in this market as spare productive capacity. For example, according to estimates in Strongin, Currie, and Fleischmann (1999) the spare capacity was about 7% to 12% around the fourth quarter of 2000. The report continues to say on page 4 that "These fields will require significant investment and drilling to not only increase production but offset the decline rates. This will take time, . . . , even if rig counts rebound substantially, during the fourth quarter, new supplies would not be available until late second quarter at the earliest." While undoubtedly inventories and stock-outs affect capacity constraints, they seem to be measured in days (around 20 days for full coverage, of which only a few days are full working inventory as the rest represents minimum operating requirements). This is consistent with our view that the final inventories affect the short end of the forward curve but are not likely to affect it along the several-month horizon we investigate.

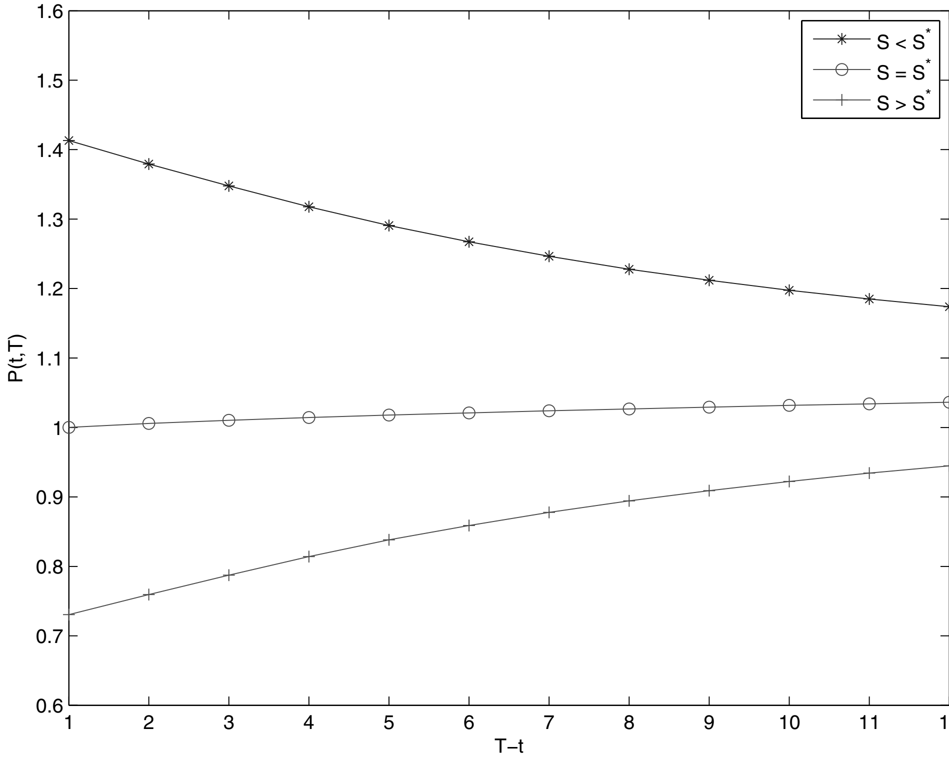

Figure 5. Futures prices, model output. Futures prices are generated by the model using parameter values reported in Table III. The investment threshold, S^*, is normalized to one.

entirely beyond the scope of our paper as our model has implications for the spot price dynamics, but not for the price of risk in the aggregate economy. To have a meaningful discussion of the price of risk process, one would need a full general equilibrium model. We first estimate the model's parameters using a simulated method of moments and then proceed to analyze and discuss some additional implications of the model.

A. Simulated Moments Parameter Estimation

A.1. Estimation Procedure

Our goal is to estimate a vector of structural parameters, $\theta \equiv \{\gamma, \mu_Y, \sigma_Y, \bar{\imath}, r, \delta, \lambda\}$. We do this using a procedure that is similar to those proposed in Lee and Ingram (1991), Duffie and Singleton (1993), Gourieroux and Monfort (1996), and Gourieroux, Monfort, and Renault (1993). Let x_t be the vector-valued process of historical futures prices and output and consider a function of the observed sample $F_T(x_t)$, where T is the sample length. The statistic $F_T(x_t)$ could represent a collection of sample moments or even a more complicated estimator, such as the slope coefficients in a regression of volatility on the term structure as in (3). Assume that as the sample size T increases, $F_T(x_t)$ converges in probability to a limit $M(\theta)$, which is a function of the structural parameters. Since many of the useful population moments cannot be computed analytically, we estimate them using Monte Carlo simulation. In particular, let $m_N(\theta) = \frac{1}{N}\sum_{n=1}^{N} F_T(x^n;\theta)$ represent the estimate of $M(\theta)$ based on N independent model-based statistics, where x^n represents a vector-valued process of simulated futures prices and output of length T based on simulating the model at parameter values, θ.[6] Let $G_N(x,\theta) = m_N(\theta) - F_T(x_t)$ denote the difference between the estimated theoretical mean of the statistic F and its observed (empirical) value. Under appropriate regularity conditions, it can be shown that as the sample size T and the number of simulations N increase to infinity, the GMM estimate of θ,

$$\theta_N = \arg\min_{\theta} J_T = \arg\min_{\theta} G_N(x,\theta)' W_T G_N(x,\theta), \tag{18}$$

will be a consistent estimator of θ. The matrix W_T in the above expression is positive definite and assumed to converge in probability to a deterministic positive definite matrix W. The Simulated Method of Moments approaches in Lee and Ingram (1991) and Duffie and Singleton (1993) focus on one long simulation while Gourieroux and Monfort (1996), and Gourieroux et al. (1993) also discuss an estimator based on multiple simulations. Our approach simulates samples of equal length to that in the data, T, and then averages across N such simulations. Given our nonbalanced panel data this approach allows for easier mapping from the model to the data and easier computation of standard errors, which are based on the distribution emanating from the cross-section of the simulations.

Assume that V is the asymptotic variance-covariance matrix of $F_T(x;\theta)$. Then, if we use the efficient choice of the weighting matrix, $W = V^{-1}$, the estimator θ_N is asymptotically normal with mean θ and covariance matrix $(D'V^{-1}D)^{-1}$, where $D = \nabla_\theta M(\theta)$.

We perform estimation in two stages. During the first stage, we use an identity matrix for the weighting matrix W. During the second stage, the weighting matrix is set equal to the inverse of the estimated covariance matrix: $W = V_N^{-1}$, where V_N is the sample-based covariance matrix of $F_T(x^n;\theta)$. To compute standard errors, we use as an estimate for D, $D_N = \nabla_\theta m_N(\theta)$.

We estimate the vector of seven model parameters, θ, by matching both the unconditional and conditional properties of futures prices. The unconditional properties include mean and volatility of daily percent price changes for futures with maturities equal to 3, 6, and 9 months, as well as the mean, volatility, and 30-day autoregressive coefficient of the slope of the forward curve

[6] Specifically, for any given value of θ, we draw N realizations of the spot price S_t from its long-run steady-state distribution (which itself depends on model parameters and is given by (16)). Then, for each set of initial conditions, we simulate a path of the state variable of the same length as the historical sample and evaluate the function $F(x,\theta)$ for each simulated path of the economy.

of crude oil futures prices. In order to see how far we can push our simple single-factor model, we also fit the relation between the volatility of futures prices and the slope of the term structure. Specifically, the conditional moments include regression coefficients $\beta_{1,T}$ and $\beta_{2,T}$ from equation (4) for T equal to 3, 6, and 9 months.

A.2. Identification

Not all of the model parameters can be independently identified from the data we are considering. In this subsection, we discuss the relations between structural parameters and observable properties of our model economy. These should suggest which of the structural parameters can be identified and what dimensions of the empirical data are likely to be most useful for estimation.

First, we calibrate the risk-free rate. The risk-free rate is determined by many factors outside of the oil industry and consequently it would not be prudent to estimate it solely based on oil price data. Also, it is clear by inspection that the risk-free rate is not identified by our model. It does not affect any of the moments we consider in our estimation and only appears in the constraint on model parameters in equation (14). Therefore, at best, futures price data can only impose a lower bound on the level of the risk-free rate, as implied by (14). Given all of the above considerations, we set the risk-free rate at 2%.[7] Next, consider a simple re-normalization of the structural parameters. Futures prices in our model depend solely on the risk-neutral dynamics of the log of the spot price, which evolves according to

$$d \log S_t = -\left[\gamma \bar{i} 1_{[S \geq S^*]} - \gamma \mu^-\right] dt + \gamma \sigma_Y d W_t. \tag{19}$$

Since we normalize the productivity parameter in (6) to one, only relative prices are informative and therefore we can ignore the dependence of S^* on structural parameters. Thus, the risk-neutral dynamics of futures prices are determined by only three combinations of five structural parameters: $\gamma\mu^-$, $\gamma\bar{i}$, and $\gamma\sigma_Y$. Therefore, we cannot identify all the model parameters separately from futures data alone.

We obtain an additional identifying condition from the oil consumption data. Cooper (2003) reports that individual growth rates vary for the 23 countries in his sample, typically falling between -3 to 3%. For the United States, the reported growth rate averages -0.7%. As documented in Cooper (2003), world crude oil consumption increased by 46% per capita from 1971 to 2000, implying an average growth rate of approximately 1.25%, which we attempt to equate with the expected growth rate of oil consumption, g_C, implied by the model

$$g_C = \bar{i}\,\Pr(S \geq S^*) - \delta = \lambda + \mu_Y - \frac{1}{2}\sigma_Y^2, \tag{20}$$

where $\Pr(S \geq S^*) = \int_{-\infty}^{S^*} p^+(S) dS = (\bar{i})^{-1}\mu^-$ is the unconditional probability that S is below the investment trigger.

[7] Our results regarding the V-shape response in prices are not affected by this choice.

Table III
Parameter Estimates, Crude Oil Futures

This table reports our parameter values. We use a two-step SMM procedure to estimate a vector of seven structural parameters $\theta \equiv \{\gamma, \mu_Y, \sigma_Y, \bar{\iota}, r, \delta, \lambda\}$. Since only five model parameters can be independently identified from the data (see Section III.A.2), we fix δ and r and estimate the remaining five parameters. We match the unconditional properties of crude oil futures prices, specifically, the historic daily return on fully collateralized 3-month futures position; the unconditional volatility of daily percent price changes for futures of various maturities; as well as the mean, volatility, and 30-day autoregressive coefficient of the slope of the forward curve. In addition, we match the expected growth rate of crude oil consumption. The standard errors are reported where applicable.

Parameter	Value	Std. Error
γ	3.1484	0.2396
$\bar{\iota}$	0.2362	0.0531
r	0.0200	NA
μ_Y	0.01438	0.0048
σ_Y	0.1043	0.0216
δ	0.1200	NA
λ	2.47×10^{-4}	1.4×10^{-4}

Finally, to estimate the risk premium λ, we use average historical daily returns on fully collateralized futures positions (we use three-month contracts). We are thus left with five independent identifying restrictions on six structural parameters. Following Gomes (2001), we fix the depreciation rate of capital at $\delta = 0.12$ per year and do not infer it from futures prices and thus estimate the remaining five parameters.

A.3. Parameter Estimates

Our estimated parameter values and the corresponding standard errors are summarized in Table III. The first parameter value in the table, $\gamma = 3.15$, implies that the price elasticity of demand in our model is -0.32. Cooper (2003) reports estimates of short-run and long-run demand elasticity for a partial adjustment demand equation based on U.S. data of -0.06 and -0.45, respectively. In our model, there is no distinction between short-run and long-run demand, as demand adjustments are assumed to be instantaneous. Our estimate falls half-way between the two numbers reported in Cooper (2003) for the U.S. and is close to the average of the long-run elasticity estimates reported for all 23 countries considered in that study, which is -0.2.

Our second parameter is $\bar{\iota}$, the maximum investment rate in the model. This variable parameterizes the investment technology used by the firms. In order to make relative empirical comparison we compare it to the growth rate of the number of operating oil wells between 1999 and 2000 for several leading crude producers worldwide.[8] During this period the number of operating wells

[8] These data are publicly available from www.WorldOil.com.

increased by 4.7% in the whole Middle East region, by 9.3% in Russia, by 22.3% in Venezuela, and by 10.3% in Norway.[9] The corresponding growth rate implied by our model is $\bar{i} - \delta = 0.12$. These numbers show that the *upper bound* of $\bar{i} = 24\%$ would allow for a plausible range of realized annual investment rates.

The average growth rate of demand is close to zero. For comparison, the average annualized change in futures prices is approximately 2.8% in the data, which falls within the 95% confidence interval of the model's prediction. The volatility of demand shocks is not directly observable. The estimated value of σ_Y together with the demand elasticity parameter γ^{-1} imply annualized volatility of the spot price of approximately 33%, which is close to the observed price volatility of short-maturity futures contracts.

Finally, the market price of risk in our model is equal to 2.47×10^{-4}. This would imply that excess expected returns on a fully collateralized futures strategy should be close to zero. This is consistent with empirical data. While an assumption of constant risk premium is clearly restrictive, it is made for simplicity: Nothing in our model prevents one from assuming a time-varying price of oil price risk. However, such an assumption would be exogenous to the model, and hence would not add to our understanding of the underlying economics of the problem.

B. Results and Discussion

B.1. Quantitative Results

We first illustrate the fit of the model by plotting the term structure of unconditional futures price volatility (to facilitate comparison with empirical data, we express our results as daily values, defined as annual values scaled down by $\sqrt{252}$). We choose model parameters, as summarized in Table III, to match the behavior of crude oil futures. Figure 6 compares the volatility of prices implied by our choice of parameters to the empirical estimates. Our model seems capable of reproducing the slow-decaying pattern of futures price volatility. This feature of the data presents a challenge to simple storage models, as discussed in Routledge et al. (2000). To see why it may not be easy to reproduce the slow-decaying pattern of unconditional volatilities in a simple single-factor model, consider a reduced-form model in which the logarithm of the spot price process follows a continuous-time AR(1) process (Ornstein-Uhlenbeck process). Specifically, assume that the spot price is given by

$$S_t = e^{y_t} \tag{21}$$

and under the risk-neutral probability measure y_t follows

[9] The total worldwide growth in the number of operating oil wells has been around 0.6% for the same period. This number is largely driven by the negative 2% growth in the United States, where most of these wells produce relatively small volumes of oil, often on an intermittent and marginally economic basis (commonly called "stripper" wells). The number of producing stripper wells changes depending on how many wells enter the ranks (by declining in production) and leave the ranks (by increasing production or being plugged and abandoned) of stripper wells each year. The United States' stripper oil well population has been gradually declining over the past decade.

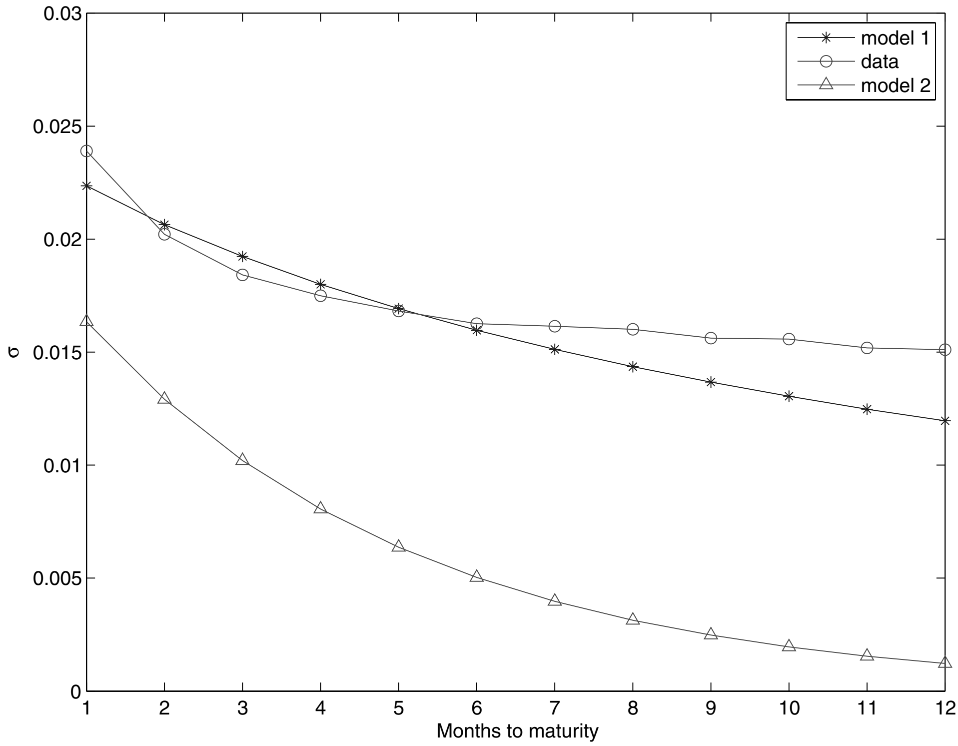

Figure 6. Unconditional volatility of futures prices, model output. The unconditional standard deviation of daily percentage changes in futures prices is constructed based on the output from the model (model 1). The model is fitted to the data on crude oil futures (CL) using a two-step simulated method of moments described in Section III.A.1. Parameter values are reported in Table III. Model 2 corresponds to the unconditional standard deviation of percentage changes in futures prices implied by equation (23). We report averages across 2,000 simulations.

$$dy_t = \theta_y(\bar{y} - y_t)\,dt + \sigma_y\,dW_t, \tag{22}$$

where θ_y is the mean-reversion coefficient and $\bar{y}$ is the long-run mean of the state variable. According to this simple model, the unconditional volatility of futures price changes is an exponential function of maturity τ:

$$\sigma^2(\tau) = \sigma_y^2 e^{-2\theta\tau}. \tag{23}$$

To compare the term structure of unconditional volatility implied by this model to the one generated by our model, we calibrate parameters θ_y and σ_y so that the simple model exhibits the same volatility of the spot price and the same 30-day autocorrelation of the basis as our model. Figure 6 shows that, as expected, unconditional volatility implied by the simple model above decays too fast relative to our model and data.

The main qualitative distinction between the properties of our model and those of basic storage models is in the conditional behavior of futures volatility. As we demonstrate in Section I, the empirical relation between the volatility of

Table IV

Conditional Volatility, Model Output

Conditional variance of crude oil futures prices is compared with the output of the model under two different econometric specifications. See the caption to Figure 2 for a detailed description of the underlying variables.

	1 Month		5 Months		10 Months	
	Data	Model	Data	Model	Data	Model
	$\lvert R(t,T)\rvert = \alpha_T + \beta_T \widetilde{SL}(t-1) + \varepsilon_T(t)$					
β_T	−0.0743	−0.0030	−0.0462	−0.0100	−0.0419	−0.012
	$\lvert R(t,T)\rvert = \alpha_T + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t)$					
$\beta_{1,T}$	0.3191	0.0854	0.1791	0.1482	0.1308	0.1246
$\beta_{2,T}$	−0.3505	−0.0964	−0.2039	−0.1822	−0.1632	−0.1651

futures prices and the slope of the term structure of prices is nonmonotone and has a pronounced V-shape. Intuitively, we would expect our model to exhibit this pattern. When the spot price S_t is far away from the investment trigger S^*, one of the investment constraints is binding and can be expected to remain binding for some time. If the capital stock K_t is much higher than its optimal level, given the current level of demand, firms find it optimal to postpone investment and the irreversibility constraint binds. On the other hand, when K_t is much lower than the optimal level, firms invest at the maximum possible rate and the investment rate constraint binds. In either case, the supply of the commodity is relatively inelastic and futures prices are relatively volatile. The further S_t travels away from the investment trigger, the larger the effect on volatility of long-maturity futures. At the same time, precisely when S_t is relatively far away from the investment trigger S^*, the absolute value of the slope of the term structure of futures prices is large, as illustrated in Figure 5. This is to be expected. All prices in our model are driven by a single mean-reverting stationary spot price, and since futures prices of longer-maturity contracts are less sensitive to the current value of the spot price, the slope of the forward curve tends to be large when the spot price is far away from its long-run average value. The latter, in turn, is not far from S^*, given that S_t reverts to S^*. Thus, our model predicts that the volatility of futures prices should exhibit a V-shape as a function of the slope of the term structure of futures prices.

It should be clear from the above discussion that the critical feature of the model is not the precise definition of the production function, but rather the variable elasticity property of the supply side of the economy. The V-shape pattern in volatilities is due to the fact that supply can adjust relatively easily in response to demand shocks when output is close to the optimal level, but supply is relatively inelastic when the output level is far from the optimum.

We now report the quantitative properties of the model. All data moments used to estimate the model are reported in the first column of Table IV. The expected growth rate of oil consumption implied by the model is equal to 0.9% and is close to an average worldwide growth rate of 1.25%. The long-run average of the slope of the forward curve, $\ln[\frac{P(t-1,6)}{P(t-1,3)}]$, is 0.0065 in the model, compared to the empirical value of −0.0125. Both values are statistically indistinguishable from zero. The long-run standard deviation of the slope in the model, which equals 0.0285, is almost identical to the empirical value of 0.0287. The 30-day autocorrelation coefficient of the slope implied by the model is equal to 0.83, as compared to the value of 0.77 in the data. Overall, our model fits the basic behavior of the slope of the forward curve quite well. Table IV shows the estimates of linear and piecewise linear specifications of conditional variance of futures price changes (4) implied by the model for 1-, 5-, and 10-month futures. The coefficients of the linear regressions are negative and close in magnitude to their empirical counterparts. Such a negative relation between conditional volatility of futures prices and the basis would typically be interpreted as supportive of simple storage models. Note, however, that our model without storage can reproduce the same kind of relation. Our model, however, has a further important implication: The linear model is badly misspecified, since the theoretically predicted relation is nonmonotone. Our piece-wise linear specification produces coefficients $\beta_{1,T}$ and $\beta_{2,T}$ that agree well with their empirical counterparts for longer maturities (3 to 12 months), but the fit worsens for shorter maturities (1 and 2 months). Given the extremely streamlined nature of our model (e.g., the slope of the forward curve is a sufficient statistic for conditional volatility), this should not be surprising. In order to capture the properties of the short end of the term structure, one must take into account storage, which we do not allow in our model. The entire distribution of regression coefficients across maturities of the futures contracts is shown in Figure 7. Finally, Figure 3 helps visualize the V-shape pattern.

B.2. Sensitivity Analysis

In order to understand the sensitivity of our results to the baseline parameters summarized in Table III, we compute elasticities of basic statistics of the model output with respect to these parameters. Each elasticity is calculated by simulating the model twice with a value of the parameter of interest 10% of one standard deviation below (above) its baseline value. Next, the change in the moment is calculated as the difference between the results from the two simulations. This difference is then divided by the change in the underlying structural parameter between the two simulations. Finally, the result is then multiplied by the ratio of the baseline structural parameter to the baseline moment. The elasticities are reported in Table V.

An increase in the demand volatility, σ_Y, or in the elasticity of the inverse demand curve, γ, leads to an increase in the volatility of the spot price, which equals $\gamma^2\sigma_Y^2$. As one would expect, the volatility of futures prices of various maturities increases as well. Qualitatively, both of the parameters σ_Y and γ affect the level of the unconditional volatility curve plotted in Figure 6. However, the demand volatility has a strong positive effect on the expected growth rate of oil consumption since it increases the long-run growth rate of the level of the demand curve, Y_t. The parameter γ has no such effect.

The constraint on the investment rate $\bar{\imath}$ has no effect on the volatility of the spot price. However, it affects the volatility of futures prices. A higher value

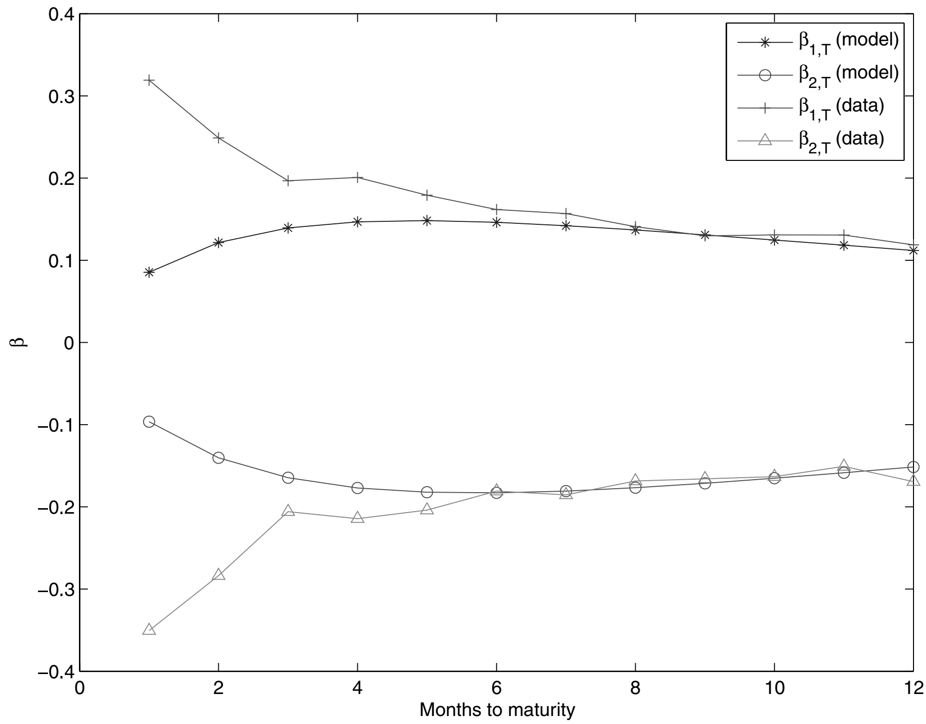

Figure 7. Conditional volatility of futures prices, model output. For each time to maturity, T, we simulate a time series of daily futures prices using parameter values reported in Table III. We then compute daily percentage changes in futures prices, denoted by $R(t, T)$. The instantaneous volatility of futures prices is related to the beginning-of-period slope of the forward curve, defined as

$$SL(t) = \ln\left[\frac{P(t-1,6)}{P(t-1,3)}\right],$$

according to the specification

$$|R(t,T)| = \alpha_T + \beta_{1,T}(\widetilde{SL}(t-1))^+ + \beta_{2,T}(\widetilde{SL}(t-1))^- + \varepsilon_T(t),$$

where $\widetilde{SL}(t)$ is the demeaned slope and $(X)^{\pm}$ denotes the positive (negative) part of X. This procedure is repeated 2,000 times and we report the average across simulations. The figure shows $\beta_{1,T}$ and $\beta_{2,T}$ for different times to maturity.

of $\bar{\imath}$ allows capital stock to adjust more rapidly in response to positive demand shocks, thus reducing the impact of demand shocks on the future value of the spot price and therefore lowering the volatility of futures prices. We thus see that $\bar{\imath}$ effectively controls the slope of the term structure of volatility, with higher values of $\bar{\imath}$ implying a steeper term structure. Further, $\bar{\imath}$ has no effect on the expected growth rate of oil consumption, in agreement with equation (20).

An increase in the unconditional mean of the demand shock, μ_Y, has little effect on the level of futures price volatility. This is not surprising given the role

Table V

Sensitivity of Model Moments to Parameters

This table presents elasticities of model moments with respect to the model parameters. The baseline parameters are given in Table III. Each elasticity is calculated by simulating the model twice: once with a value of the parameter of interest 10% of one standard deviation below its baseline value, and once with a value 19% of one standard deviation above its baseline value. Then the change in the moment is calculated as the difference between the results from the two simulations. This difference is then divided by the change in the underlying structural parameter between the two simulations. The result is then multiplied by the absolute value of the ratio of the baseline structural parameter to the baseline moment.

	Data	Baseline Moments	γ	$\bar{i}$	μ_Y	σ_Y	$\lambda(\times 10^{-2})$
Average Slope of the Forward Curve	−0.0124	0.0065	1.8401	−1.2862	0.1814	2.3471	−0.5004
Volatility of the Slope of the Forward Curve	0.0287	0.0285	0.9704	0.5516	−0.0377	0.7005	0.4189
30-days AR(1) Coefficient of Forward Curve Slope	0.7653	0.8337	−0.0041	−0.1221	−0.0038	0.2112	0.1177
Long Term Crude Oil Consumption Growth	0.0092	0.0092	0.0001	−0.0003	1.5652	−1.1843	2.6939
Unconditional Variance 3-month CL Futures	0.0184	0.0186	0.7444	−0.5535	0.0036	1.2367	0.0646
Unconditional Variance 6-month CL Futures	0.0163	0.0149	0.6807	−0.9276	0.0104	1.4487	0.0625
Unconditional Variance 9-month CL Futures	0.0156	0.0124	0.6176	−1.2162	0.0173	1.5690	0.0706
$\beta_{1,T}$ 3-month CL Futures	0.1968	0.1393	−0.0367	1.2496	−0.1222	−0.1146	−0.0488
$\beta_{1,T}$ 6-month CL Futures	0.1616	0.1462	−0.0634	1.0368	−0.1475	0.4632	−0.1799
$\beta_{1,T}$ 9-month CL Futures	0.1298	0.1309	−0.0795	0.8254	−0.1634	0.9137	−0.0662
$\beta_{2,T}$ 3-month CL Futures	−0.2059	−0.1646	0.0745	−1.7212	0.1753	−0.0084	0.3828
$\beta_{2,T}$ 6-month CL Futures	−0.1808	−0.1829	0.0897	−2.6308	0.2155	0.5637	0.4722
$\beta_{2,T}$ 9-month CL Futures	−0.1659	−0.1713	0.0983	−3.5200	0.2551	1.0333	0.5976

μ_Y plays in the evolution of the spot price S_t. An increase in μ_Y raises the drift of S_t uniformly. The impact of this on the volatility of futures prices is ambiguous and depends on the relative magnitude of the drift of S_t above, μ^-, and below, $\mu^+ \equiv \bar{\imath} - \mu^-$, the investment threshold S^*. By symmetry considerations, if $\mu^+ = \mu^-$, an infinitesimal change in μ_Y has no impact on the volatility of futures prices. Under the calibrated parameter values, $\mu^- = 0.1289$ and $\mu^+ = 0.1083$ and futures volatility is not very sensitive to μ_Y. The same is true for the risk premium, λ. Both μ_Y and λ have a strong positive effect on g_C, in agreement with equation (20).

In general, the effect of model parameters on the slope of the forward curve is difficult to interpret intuitively and depends on the chosen parameter values. However, the fact that the moments of the slope have different sensitivities to various model parameters makes them useful in estimating these parameters.

IV. Conclusions

This paper contributes along two dimensions. First, we show that volatility of futures prices has a V-shape relationship with respect to the slope of the term structure of futures prices. Second, we show that such volatility patterns arise naturally in models that emphasize investment constraints and, consequently, time-varying supply elasticity as a key mechanism for price dynamics. Our empirical findings seem beyond the scope of simple storage models, which are currently the main focus of the literature, and point towards investigating alternative economic mechanisms, such as the one analyzed in this paper. Adding finite storage capacity to this economy would likely affect the very short end of the forward curve. It is likely, then, that the volatility of spot prices would be related to the level of inventories. However, adding storage is not likely to have a material effect on the long end of the forward curve, since in practice the total storage capacity is rather small relative to annual consumption. What is difficult to predict, however, is how storage would interact with production constraints at intermediate maturities. This requires formal modeling and future work will entail a model that nests both storage and investment in an attempt to isolate their quantitative effects.

Appendix A: Proof of Proposition 1

We conjecture that the equilibrium investment policy I_t^* is given by (13). Market clearing in the spot market then implies that the spot price process S_t satisfies (15).

A competitive firm chooses an investment policy I_t to maximize the firm value, that is, the present value of future output net of investment costs:

$$\max_{I_t} \mathrm{E}_0 \left[\int_0^\infty e^{-rt} \left(K_t S_t - I_t \right) dt \right], \tag{A1}$$

subject to the capital accumulation rule

$$dK_t = (I_t - \delta K_t)dt, \tag{A2}$$

$$I_t \geq 0, \tag{A3}$$

$$I_t \leq \bar{i} K_t^A. \tag{A4}$$

From (A2), we obtain

$$K_t = \int_0^t e^{-\delta(t-s)} I_s ds + K_0 e^{-\delta t}. \tag{A5}$$

Using this expression for the capital stock, and relaxing the constraint (A4), we re-write the above optimization problems as

$$\max \mathrm{E}_0 \left[\int_0^\infty e^{-rt}(V_t - 1) I_t dt \right] + K_0 V_0, \tag{A6}$$

subject to (A3, A4), where

$$V_t = \mathrm{E}_t \left[\int_t^\infty e^{-(r+\delta)(u-t)} S_u du \right]. \tag{A7}$$

Assuming that the process V_t is well defined (we prove that next), the optimal solution of the firm's problem is

$$I_t^* = \begin{cases} \bar{i} K_t^A, & V_t - 1 > 0, \\ \left[0, \bar{i} K_t^A\right], & V_t - 1 = 0, \\ 0, & V_t - 1 < 0. \end{cases} \tag{A8}$$

The above policy will coincide with (13) as long as $V_t = 1$ whenever $S_t = S^*$. Let $V(S, S^*)$ denote the value of V_t when $S_t = S$ and the optimal investment threshold is S^*. Note that the conjectured form of the equilibrium spot price process implies that

$$V_t = V(S_t, S^*) = S^* V\left(\frac{S_t}{S^*}, 1\right). \tag{A9}$$

Thus, the equilibrium value of S^* can be found as $S^* = V(1, 1)^{-1}$.

We now characterize $V(S; 1)$, and show that V_t is finite. Let $S^* = 1$ and define $\omega_t = -(1/\gamma)\ln S_t$. Then

$$d\omega_t = \left(\bar{i} 1_{[\omega_t \leq 0]} - \mu^-\right) dt - \sigma_Y dW_t. \tag{A10}$$

Let B be an arbitrary positive number and define a stopping time $\tau_B = \inf\{t : \omega_t \leq -B\}$. Let

$$F_t^B = \mathrm{E}_t \left[\int_t^\infty 1_{[s \leq \tau_B]} e^{-(r+\delta)(s-t) - \gamma \omega_s} ds \right]. \tag{A11}$$

We look for $F_t^B = F^B(\omega_t)$. We start by heuristically characterizing $F^B(\omega)$ as a unique solution of the Feynman-Kac equation

$$\frac{\sigma_Y^2}{2}\frac{d^2F^B(\omega)}{d\omega^2} + \left[\bar{i}1_{[\omega\le 0]} - \mu^-\right]\frac{dF^B(\omega)}{d\omega} - (r+\delta)F^B(\omega) + e^{-\gamma\omega} = 0 \quad \text{(A12)}$$

with the boundary condition

$$F^B(-B) = 0. \quad \text{(A13)}$$

We look for a solution of (A12) of the form

$$F^B(\omega) = \begin{cases} Ae^{\varkappa_-\omega} + \dfrac{e^{-\gamma\omega}}{r - \gamma\mu^- - \dfrac{\gamma^2}{2}\sigma_Y^2}, & \omega \ge 0, \\ C_1e^{\overline{\varkappa}_+\omega} + C_2e^{\overline{\varkappa}_-\omega} + \dfrac{e^{-\gamma\omega}}{r + \gamma\mu^+ - \dfrac{\gamma^2}{2}\sigma_Y^2}, & \omega < 0. \end{cases} \quad \text{(A14)}$$

Substituting these solutions into ODE (A12) yields quadratic equations on $\varkappa_\pm$ and $\overline{\varkappa}_\pm$

$$\frac{\sigma_Y^2}{2}\varkappa_\pm^2 - \mu^-\varkappa_\pm - (r+\delta) = 0, \quad \text{(A15)}$$

$$\frac{\sigma_Y^2}{2}\overline{\varkappa}_\pm^2 + \mu^+\overline{\varkappa}_\pm - (r+\delta) = 0. \quad \text{(A16)}$$

The terms $\varkappa_\pm$ and $\overline{\varkappa}_\pm$ denote, respectively, positive and negative roots of the above equations. Define

$$M = \frac{\gamma\bar{i}}{\left(r + \delta - \gamma\mu^- - \dfrac{\gamma^2}{2}\sigma_Y^2\right)\left(r + \delta + \gamma\mu^+ - \dfrac{\gamma^2}{2}\sigma_Y^2\right)}. \quad \text{(A17)}$$

Using the boundary condition (A13) and imposing continuity of the function F^B and its first derivative across zero (to verify that the solution of the differential equation (A12) characterizes the expected value F_t^B, we only need the first derivative to be continuous at zero), we obtain the following system of equations on coefficients C_1, C_2, and A:

$$\begin{cases} C_1 + C_2 = A + M, \\ \overline{\varkappa}_+C_1 + \overline{\varkappa}_-C_2 = \left(\dfrac{2\bar{i}}{\sigma_Y^2} + \varkappa_-\right)A + 2\dfrac{r + \delta + \gamma\mu^+ - \gamma^2\sigma_Y^2}{\gamma\sigma_Y^2}M, \\ C_1e^{-\overline{\varkappa}_+B} + C_2e^{-\overline{\varkappa}_-B} = \dfrac{r + \delta - \gamma\mu^- - \dfrac{\gamma^2}{2}\sigma_Y^2}{\gamma\bar{i}}Me^{\gamma B}. \end{cases} \quad \text{(A18)}$$

Solving this system yields the following expression for C_2:

$$C_2 = M \frac{\dfrac{r+\delta-\gamma\mu^- - \dfrac{\gamma^2}{2}\sigma_Y^2}{\gamma\bar{i}}\left(\overline{\varkappa}_+ - \dfrac{2\bar{i}}{\sigma_Y^2} - \varkappa_-\right)e^{(\overline{\varkappa}_- + \gamma)B} - \left(2\dfrac{r+\delta+\gamma\mu^+ - \gamma^2\sigma_Y^2}{\gamma\sigma_Y^2} - \varkappa_-\right)e^{(\overline{\varkappa}_- - \overline{\varkappa}_+)B}}{\overline{\varkappa}_+ - \dfrac{2\bar{i}}{\sigma_Y^2} - \varkappa_- - \left(\overline{\varkappa}_- - \dfrac{2\bar{i}}{\sigma_Y^2} - \varkappa_-\right)e^{(\overline{\varkappa}_- - \overline{\varkappa}_+)B}}. \tag{A19}$$

Next, we show that, indeed, $F_t^B = F^B(\omega_t)$. To see this, note that the process $X_t = e^{-(r+\delta)t}F^B(\omega_t) + \int_0^t e^{-(r+\delta)s-\gamma\omega_s}\,ds$ is a local martingale. This follows from the fact that, by Ito's lemma, the drift of the process is equal to zero (due to (A12)). Next, since the diffusion coefficient of X_t, $\sigma\, dF^B(\omega)/d\omega$ is bounded on the domain $\omega \geq -B$, the stopped process $X_{t\wedge\tau_B}$ is a martingale. Thus,

$$X_0 = F^B(\omega_0) = \mathrm{E}_0\,[X_T] = \mathrm{E}_0\left[1_{[T\leq\tau_B]}e^{-(r+\delta)T}F^B(\omega_T) + \int_0^T 1_{[s\leq\tau_B]}e^{-(r+\delta)s-\gamma\omega_s}\,ds\right].$$

Since the function $F^B(\omega)$ is bounded on the domain $\{\omega_t \geq -B\}$, we know that $\lim_{T\to\infty}\mathrm{E}_0[1_{[T\leq\tau_B]}e^{-(r+\delta)T}F^B(\omega_T)] = 0$. Thus, as we take $T\to\infty$, by monotone convergence theorem, $F_B(\omega_0) = F_0^B$.

Having found F_t^B, we now take the limit of $B\to\infty$. The $\lim_{B\to\infty}F_t^B$ is well defined. Because $\frac{\sigma_Y^2\gamma^2}{2} - \gamma\mu^+ - (r+\delta) < 0$, it follows that $\overline{\varkappa}_- + \gamma < 0$ and therefore $\lim_{B\to\infty}C_2 = 0$. Also, constants C_1 and A converge to finite limits as $B\to\infty$.

To show that $\lim_{B\to\infty}F_t^B = V_t$, we use the monotone convergence theorem, combined with the observation that $\lim_{B\to\infty}\tau_B = \infty$. The latter follows from the fact that $\omega_t \leq \omega_0 + \mu^- t + \sigma(W_t - W_0)$, which is an arithmetic Brownian motion and for which the corresponding stopping time converges to infinity as $B\to\infty$.

Thus, $V(S,1)$ is well defined, it is a function of the state variable ω, $F(\omega) = V(e^{-\gamma\omega}, 1)$, which satisfies the equation

$$\frac{\sigma_Y^2}{2}\frac{d^2F(\omega)}{d\omega^2} + \left[\bar{i}1_{[\omega\leq\omega^*]} - \mu^-\right]\frac{dF(\omega)}{d\omega} - (r+\delta)F(\omega) + e^{-\gamma\omega} = 0 \tag{A20}$$

and is given explicitly by

$$F(\omega) = \begin{cases} Ae^{\varkappa_-\omega} + \dfrac{e^{-\gamma\omega}}{r-\gamma\mu^- - \dfrac{\gamma^2}{2}\sigma_Y^2}, & \omega \geq 0, \\[2ex] Ce^{\overline{\varkappa}_+\omega} + \dfrac{e^{-\gamma\omega}}{r+\gamma\mu^+ - \dfrac{\gamma^2}{2}\sigma_Y^2}, & \omega < 0, \end{cases} \tag{A21}$$

where

$$A = \frac{2\dfrac{r+\delta+\gamma\mu^+ - \gamma^2\sigma_Y^2}{\gamma\sigma_Y^2} - \overline{\varkappa}_+}{\overline{\varkappa}_+ - \dfrac{2\bar{i}}{\sigma_Y^2} - \varkappa_-} M,$$

$$C = \left(1 + \frac{2\dfrac{r+\delta+\gamma\mu^+ - \gamma^2\sigma_Y^2}{\gamma\sigma_Y^2} - \overline{\varkappa}_+}{\overline{\varkappa}_+ - \dfrac{2\bar{i}}{\sigma_Y^2} - \varkappa_-}\right) M.$$

This completes our proof.

Appendix B: Stationary Long-Run Distribution of S_t

Define $\omega_t = -(1/\gamma)\ln S_t$. Then

$$d\omega_t = \left(\bar{i}1_{[\omega_t \le \omega^*]} - \mu^-\right)dt - \sigma_Y dW_t, \quad \omega^* \equiv -\frac{1}{\gamma}\ln S^*. \tag{B1}$$

It is enough to calculate the stationary long-run distribution of ω,

$$p(\omega) = \begin{cases} p^+(\omega)\ \omega \le \omega^* \\ p^-(\omega)\ \omega > \omega^* \end{cases}. \tag{B2}$$

This distribution exists if $0 < \mu^- < \bar{i}$ and it satisfies the forward Kolmogorov ODE

$$\frac{d^2p(\omega)}{d\omega^2} - \frac{2\left[\mu^+1_{[\omega\le\omega^*]} - \mu^-1_{[\omega>\omega^*]}\right]}{\sigma_Y^2}\frac{dp(\omega)}{d\omega} = 0. \tag{B3}$$

The distribution $p(\omega)$ also satisfies the normalization condition

$$\int_{-\infty}^{\omega^*} p^+(\omega)d\omega + \int_{\omega^*}^{\infty} p^-(\omega)\,d\omega = 1. \tag{B4}$$

Condition (B4) eliminates a constant as a trivial solution of the ODE (B3). We solve the ODE (B3) separately for $p^+(\omega)$ and $p^-(\omega)$:

$$\begin{aligned} p^+(\omega) &= A^+e^{\frac{2\mu^+}{\sigma_Y^2}\omega}, \\ p^-(\omega) &= A^-e^{-\frac{2\mu^-}{\sigma_Y^2}\omega}, \end{aligned} \tag{B5}$$

and paste the solutions together so that $p(\omega)$ is continuous at ω^*. Thus, $A^{\pm}$ can be found from the boundary condition at ω^* and the normalization condition:

$$A^+e^{\frac{2\mu^+}{\sigma_Y^2}\omega^*} = A^-e^{-\frac{2\mu^-}{\sigma_Y^2}\omega^*}, \tag{B6}$$

$$\frac{\sigma_Y^2}{2\mu^+}A^+e^{\frac{2\mu^+}{\sigma_Y^2}\omega^*} + \frac{\sigma_Y^2}{2\mu^-}A^-e^{-\frac{2\mu^-}{\sigma_Y^2}\omega^*} = 1. \tag{B7}$$

Solving equations (B6) and (B7) for A^+ and A^-, and changing variables according to $\omega = -\frac{1}{\gamma}\ln S$, we obtain (16).

REFERENCES

Atkeson, Christopher G., Andrew W. Moore, and Stefan Schaal, 1997, Locally weighted learning, *Artificial Intelligence Review* 11, 11–73.

Brennan, Michael J., 1991, The price of convenience and the valuation of commodity contingent claims, in D. Lund and B. Oksendal, eds.: *Stochastic Models and Option Models* (Elsevier Science Publishers, Amsterdam).

Brennan, Michael J., and Eduardo S. Schwartz, 1985, Evaluating natural resource investments, *Journal of Business* 58, 135–157.

Carlson, Murray, Zeigham Khoker, and Sheridan Titman, 2006, Equilibrium exhaustible resource price dynamics, NBER working paper.

Casassus, Jaime, Pierre Collin-Dufresne, and Bryan R. Routledge, 2004, Equilibrium commodity prices with irreversible investment and non-linear technologies, NBER Working paper.

Cooper, John C. B., 2003, Price elasticity of demand for crude oil: Estimates for 23 countries, *OPEC Review* 27, 1–8.

Deaton, Angus, and Guy Laroque, 1992, On the behavior of commodity prices, *Review of Financial Studies* 59, 1–23.

Deaton, Angus, and Guy Laroque, 1996, Competitive storage and commodity price dynamics, *Journal of Political Economy* 104, 896–923.

Duffie, Darrell, and Stephen Gray, 1995, Volatility in energy prices, in Robert Jameson, ed.: *Managing Energy Price Risk* (Risk Publications, London, U.K.).

Duffie, Darrell, and Kenneth J. Singleton, 1993, Simulated moments estimation of Markov models of asset prices, *Econometrica* 61, 929–952.

Fries, Steven, Marcus Miller, and William Perraudin, 1997, Debt in industry equilibrium, *Review of Financial Studies* 10, 39–67.

Gomes, Joao, 2001, Financing investment, *American Economic Review* 91, 1263–1285.

Gourieroux, Christian, and Alain Monfort, 1996, *Simulation based econometric methods* (Oxford University Press, Oxford, U.K.).

Gourieroux, Christian, Alain Monfort, and E. Renault, 1993, Indirect inference, *Journal of Applied Econometrics* 8, 85–118.

Grenadier, Steven R., 2002, Option exercise games: An application to the equilibrium investment strategies of firms, *Review of Financial Studies* 15, 691–721.

Kaldor, Nicholas, 1939, Speculation and economic stability, *Review Economics Studies* 7, 1–27.

Kogan, Leonid, Dmitry Livdan, and Amir Yaron, 2005, Futures prices in production economy with investment constraints, NBER Working paper.

Lee, Bong-Soo, and Beth F. Ingram, 1991, Simulation estimation of time series models, *Journal of Econometrics* 47, 197–205.

Litzenberger, Robert H., and Nir Rabinowitz, 1995, Backwardation in oil futures markets: Theory and empirical evidence, *Journal of Finance* 50, 1517–1545.

Miao, Jianjun, 2005, Optimal capital structure and industry dynamics, *Journal of Finance* 60, 2621–2659.

Ng, Victor K., and Craig S. Pirrong, 1994, Fundamentals and volatility: Storage, spreads, and the dynamics of metal prices, *Journal of Business* 67, 203–230.

Novy-Marx, Robert, 2007, An equilibrium model of investment under uncertainty, *Review of Financial Studies* 20, 1461–1502.

Routledge, Bryan R., Duane J. Seppi, and Chester S. 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.

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

Strongin, Steven, Jeffrey Currie, and Allison Fleischmann, 1999, *Energy Monthly*, Goldman Sachs Commodities Research, Goldman Sachs & Co, New York.

Telser, Lester G., 1958, Futures trading and the storage of cotton and wheat, *Journal of Political Economy* 66, 133–144.

Williams, Jeffrey C., and Brian D. Wright, 1991, *Storage and Commodity Markets* (Cambridge University Press, Cambridge, UK).

Working, Holbrook, 1948, Theory of the inverse carrying charge in futures markets, *Journal of Farm Economics* 30, 1–28.

Working, Holbrook, 1949, The theory of the price of storage, *American Economic Review* 39, 1254–1262.