

Why do Firms Hold Oil Stockpiles?*

Charles F. Mason[†]

H.A. True Chair in Petroleum and Natural Gas Economics

Department of Economics & Finance

University of Wyoming,

1000 E. University Ave.

Laramie, WY 82071

February 24, 2022

Abstract

Persistent and significant privately-held stockpiles of crude oil have long been an important empirical regularity in the United States. Such stockpiles would not rationally be held in a traditional Hotelling-style model. How then can the existence of these inventories be explained? In the presence of sufficiently stochastic prices, oil extracting firms have an incentive to hold inventories to smooth production over time. An alternative explanation is related to a speculative motive - firms hold stockpiles intending to cash in on periods of particularly high prices. I argue that empirical evidence supports the former but not the latter explanation.

Keywords: Crude Oil, Inventories, Stochastic Dynamic Optimization

JEL Areas: Q2, D8, L15

*I thank seminar participants at Tulane and UC Davis and participants at the World Congress of Environmental and Resource Economists, the Association of Environmental and Resource Economists conference and the Allied Social Science Association meetings in Chicago for their useful input on earlier version of this paper. I am particularly grateful to Stephen Holland for his thoughtful discussion at the ASSA meetings. The usual disclaimer applies.

[†]phone: +1-307-766-2178; fax: +1-307-766-5090; e-mail: bambuzlr@uwyo.edu

1 Introduction

In late 2007 through the middle of 2008, crude oil prices reached dizzying heights. In reaction to the pattern of crude spot prices during that period, some key players in OPEC and a number of members of the United States (US) Congress, placed the blame on speculators – which indirectly implies a link to inventory holdings. Then in the aftermath of the price collapse in the last quarter of 2008, in early 2009 through 2010, a number of pundits alleged that prices did not recover because of the significant inventory holdings across the globe. Similarly, over the past year, considerable attention has been focused on crude oil inventories. Indeed, private inventory holdings in the US accelerated considerably during this period. Similarly, significant inventory holdings were alleged to play a role in the sluggish recovery of oil prices in 2016-2017; after these stockpiles were reduced oil prices increased markedly during 2018. More recently, large inventory holdings near Cushing, Oklahoma – site of the West Texas Intermediate trading hub – are thought to have played a significant role in the singular occurrence of negative spot prices in April, 2020. These events underscore the significance of crude oil inventories, and of evaluating the motives for holding these stockpiles.

Since official records were first kept in the US, in 1920, private interests have consistently held significant inventories of crude oil. Over the course of the past few decades, these inventories have averaged around 325 million barrels. While these holdings have fluctuated some they have been remarkably persistent over the past 70 years, ranging from just over 215 million barrels to slightly less than 398 million barrels (see Figure 1). What motivates these substantial inventory holdings? One answer is that stockpiles could be held for speculative purposes – betting on abnor-

mally rapid price run-ups. An alternative explanation is that petroleum extracting firms would like to hedge against substantial swings in extraction costs.¹

Neither explanation is compelling in a deterministic setting. In a deterministic world prices would have to rise at the rate of interest to induce firms to hold stockpiles. But if prices increased at the rate of interest, rents would typically rise faster than the interest rate. Firms would then prefer to delay extraction, so that there would be no fodder from which to build inventories. The answer, I believe, must lie in fluctuating prices.

As I noted above, the claims that speculators were manipulating crude prices in early 2008 directs attention to the role of privately held inventories. If speculation was at play, one would expect resource inventory holders to cash in on abnormally high prices. As I discuss below, while there was a negative correlation between spot prices and inventory holdings, prices only explain a paltry amount of the variation in inventories. Indeed, inventories did not change much even when prices increased or decreased dramatically, as during this past year. It seems likely that some other effect played a more important role.

An alternative, and I believe compelling, motivation is related to the concept of production smoothing (Arrow et al., 1958; Blanchard and Fisher, 1989). If oil prices are driven by a random process, perhaps arising from demand shocks, the induced fluctuations in market price will lead to variations in the firm's optimal extraction rate. So long as there is enough variation in production, relative to the overall downward trend in production that must occur for non-renewable resources, and so long as it is costly to expand production, firms will wish to hold inventories to guard against future cost increases. This explanation will hold true no matter what current price is, and no matter

what the current level of resources *in situ*.

In this paper I explore the implication of such motivations. I start by discussing the conceptual underpinnings of the story in section 2, formally demonstrating that a resource extracting firm would generally not acquire stockpiles in a deterministic world. I then analyze a version of the model allowing for stochastic prices in section 3. I turn to an examination of the data in section 4. Here I argue that the variation in spot prices has been sufficient to motivate the acquisition of inventories for almost all months during the past two decades. By contrast, I find that the impact of spot prices upon both levels of and changes in privately-held oil stocks is modest at best. I conclude with a discussion of potential extensions of the model in section 5.

2 Deterministic Prices and the Incentive to Stockpile

Consider a price-taking firm engaged in the extraction of oil. The firm in question has an initial deposit of the resource of size R_0 , from which it may choose to extract. Its rate of extraction is y_t , and its rate of sales, q_t , are selected to maximize the discounted flow of its profits. It will be convenient to adjust the firm's problem slightly, and use net additions to inventories, $w_t = y_t - q_t$, as a control variable in place of sales.

The firm's reserves at instant t are R_t and its inventory holdings are S_t . I assume the firm starts with no inventories. Reserves decumulate with extraction, while inventories accumulate

according to the difference between extraction and sales:

$$\dot{R}_t = -y_t; \quad (1)$$

$$\dot{S}_t = w_t. \quad (2)$$

When it is actively extracting, the firm bears positive operating costs. I assume marginal extraction costs are positive, upward-sloping and weakly convex, with both total costs and marginal costs decreasing in R . A simple example of a cost function that has these features is

$$c(y, R) = A_0 + A_1 y^\eta / R, \quad (3)$$

which is adapted from Pindyck (1980). This function, which combines flow fixed costs with a power function of the rate of extraction that is proportional to the inverse of reserves, has two desirable features: There is a range of falling average extraction costs, and extraction becomes more costly the greater is the ratio of extraction to reserves; both aspects are consistent with anecdotal evidence. In this functional form, $\eta - 1$ can be interpreted as the elasticity of marginal extraction cost with respect to the rate of extraction. The assumption of weakly convex marginal costs implies $\eta \geq 2$. For now, I assume that it is costless to hold inventories; the implications of relaxing this assumption are discussed below.

Denoting the market price of oil at instant t by P_t , the instantaneous rate of profits is

$$\pi_t = P_t[y_t - w_t] - c(y_t, R_t). \quad (4)$$

The goal is to select time paths of y and w so as to maximize the present discounted value of the flow of profits.

The firm's current value Hamiltonian is

$$\mathcal{H} = P_t(y_t - w_t) - c(y_t, R_t) - \lambda_t y_t + \mu_t w_t,$$

where λ_t and μ_t are the current-value shadow prices of reserves and inventories, respectively. Pontryagin's maximum condition gives the necessary conditions for optimization:

$$P_t - \frac{\partial c}{\partial y} - \lambda_t = 0; \tag{5}$$

$$P_t - \mu_t \begin{cases} > 0 & \Rightarrow w_t = -\infty \text{ if } S_t > 0; w_t = 0 \text{ if } S_t = 0 \\ = 0 & \Rightarrow w_t \text{ is indeterminate.} \\ < 0 & \Rightarrow w_t = y_t \end{cases} \tag{6}$$

In principle, it is possible for the firm to liquidate some of its inventories by choosing $w = -\infty$. As such action would radically depress market price it can be ruled out by market clearing. On the other hand, if the firm does not hold inventories then $w \geq 0$ (i.e., there are no inventories to sell from). Moreover, since as a general rule oil firms do not stockpile all their extraction, it seems the first branch is empirically implausible. I therefore proceed assuming the firm's optimal time path of w is based on the middle branch of (6), unless it never pays to acquire inventories.²

In addition to the first-order conditions above, the solution must satisfy the equations of

motion for the shadow values:

$$\dot{\lambda} = r\lambda + \frac{\partial c}{\partial R}; \quad (7)$$

$$\dot{\mu} = r\mu; \quad (8)$$

where r is the interest rate. It is apparent that the solution to the differential equation governing μ is an exponential, with that shadow value growing at the rate r .

If the firm is actively extracting over an interval of time then one may time-differentiate eq. (5). Then combining with eq. (7), one obtains

$$d\left[P_t - \frac{\partial c}{\partial y}\right]/dt = \dot{\lambda} = r\left[P_t - \frac{\partial c}{\partial y}\right] + \frac{\partial c}{\partial R}, \text{ or}$$

$$\left[\dot{P}_t/P_t - r\right]P_t = d\left[\frac{\partial c}{\partial y}\right]/dt - r\frac{\partial c}{\partial y} + \frac{\partial c}{\partial R}. \quad (9)$$

Suppose now that the firm finds it optimal to add to inventories over a period of time. Then the middle branch of eq. (6) applies; time-differentiating and combining with eq. (8), one infers that price would then rise at the rate of interest. The conclusion is that prices must increase at the rate of interest for the firm to be willing to add to inventories. But eq. (9) then implies

$$d\left[\frac{\partial c}{\partial y}\right]/dt - r\frac{\partial c}{\partial y} + \frac{\partial c}{\partial R} = 0 \iff$$

$$(\partial^2 c/\partial y^2)\dot{y} - (\partial^2 c/\partial y\partial R)y - r\frac{\partial c}{\partial y} + \frac{\partial c}{\partial R} = 0.$$

With the particular functional form in eq. (3), this condition reduces to

$$\dot{y}/y = \frac{r}{\eta - 1} - \frac{y}{\eta R}. \quad (10)$$

If this simple relation fails the firm will either sell all or none of its extracted oil. Since as a general rule oil firms do not stockpile all their extraction, it seems the empirically likely outcomes are either stockpiling (if eq. (10) holds) or no stockpiling (if it does not).

Intuitively, if the firm were to hold stockpiles, it would possess two classes of stocks, inventories and *in situ* reserves. These stocks differ in terms of their extraction costs: inventories can be costlessly used (since the extraction costs have already been paid), while reserves in the ground are costly to extract. In this case, the optimal program must use up the lower cost reserves first. However, the only way inventories could exist in the first place is if excess extraction were to occur at some point in time, and so it follows that no inventories would ever be held.

It is worth reiterating that prices are deterministic in this context – i.e., the entire price path is known. What is the implication of relaxing this assumption, allowing for stochastic prices?

3 A Model With Stochastic Prices

Now suppose that the spot price of oil follows a random process, where the fluctuations in price result from demand-side shocks. For concreteness I take this random process to be geometric Brownian motion:³

$$dP_t/P_t = \mu dt + \sigma dz, \quad (11)$$

where dz is an increment from a standard Wiener process. Convergence of the model requires that the trend in prices does not exceed r , the firm's discount rate: $\mu < r$.

The nature of the firm's decision problem is similar to those studied by Pindyck (1980, 1982). At each instant the firm's decision problem is governed by the level of its reserves, its inventories and market price. For expositional simplicity I assume the firm chooses to actively extract over the time horizon in question; allowing for the possibility the firm might wish to cease extraction, or re-activate extraction, can be readily incorporated, though at the cost of some extra complexity.⁴

Let $V(t, R_t, S_t, P_t)$ denote the optimal value function when the firm is currently active at instant t , with *in situ* reserves of R_t , inventories of S_t and market price equal to P_t . The fundamental equation of optimality for a currently active firm is (Kamien and Schwartz, 1991):

$$\mathbf{max}_{y_t, w_t} \left\{ \pi_t e^{-rt} + \partial V / \partial t - y_t \partial V / \partial R + w_t \partial V / \partial S + \mu P_t \partial V / \partial P + (\sigma^2 P_t^2 / 2) \partial^2 V / \partial P^2 \right\} = 0. \quad (12)$$

As in the deterministic case, the optimal extraction rate balances current rents against the shadow price of reserves *in situ*, for a firm that actively extracts at instant t :

$$P_t - \frac{\partial c}{\partial y}(y_t^*, R_t) - \partial V / \partial R = 0, \quad (13)$$

where y_t^* solves the maximization problem in (12). Also as in the deterministic case, the maximand

in (12) is linear in w_t . Thus, optimal adjustments to inventories satisfy

$$P_t - \partial V / \partial S \begin{cases} > 0 & \Rightarrow w_t = -\infty \text{ if } S_t > 0; w_t = 0 \text{ if } S_t = 0 \\ = 0 & \Rightarrow w_t \text{ is indeterminate.} \\ < 0 & \Rightarrow w_t = y_t \end{cases} \quad (14)$$

As above, if P_t exceeds the shadow value of inventories, here measured by $\partial V / \partial S$, the firm is motivated to draw down its inventories as rapidly as feasible. If the shadow value of inventories is larger than current price, all production is allocated to inventories. If $P_t = \partial V / \partial S$, then w_t is indeterminate.

It is instructive to think of the firm as solving a sequence of problems. At each instant t , the firm determines an optimal program, based on the current (and observed) demand shock. This consists of extraction and inventory plans for each future instant that maximize the discounted expected flow of profits, conditional on current demand, where the expectation is with respect to the future stream of prices. This program is subject to the anticipation that reserves will be exhausted at the terminal moment (Pindyck, 1980). Then in the next instant, a new demand shock is observed and the firm re-optimizes.

In the analysis I conducted within the deterministic framework, the next step was to time-differentiate the condition governing the optimal extraction rate. Here, however, the optimal extraction rate will generally be a function of the stochastic variable P , as will the marginal value of reserves. As a result, there is no proper time derivative for either side of eq. (13). The stochastic analog of the time derivative, Ito's differential operator, $\frac{1}{dt}E[d(\bullet)]$, is used in its place (Kamien

and Schwartz, 1991). Applying this operator to eq. (13) yields:

$$\frac{1}{dt}E[d(P)] - \frac{1}{dt}E\left[d\left(\frac{\partial c}{\partial y}\right)\right] = \frac{1}{dt}E[d(\partial V/\partial R)], \quad (15)$$

where I have omitted the time subscript where there will be no confusion.

In the deterministic case, one expects the time rate of change in marginal costs to be smaller than the present value of current marginal cost.⁵ Unlike the deterministic case, however, marginal costs can rise over time in the context of stochastic demand. Despite the overall tendency for production to decline over time, on average, the stochastic nature of extraction can yield an increase in anticipated marginal cost if the variation in extraction is sufficiently large, relative to the slope of marginal costs. This occurs because the optimal extraction rate is subject to a stochastic influence, which in turn means that marginal extraction cost will typically fluctuate. If there is enough variation in the demand shock, this more than compensates for the reductions in extraction that will occur on average.

From the discussion above, if the firm is to be willing to hold inventories then it must be the case that $P_t = \partial V/\partial S$. The analysis leading up to equation (12) in Pindyck (1980) can be applied here to show that $\frac{1}{dt}E[d(\partial V/\partial S)] = r\partial V/\partial S$. It follows that a necessary condition for the firm to be willing to stockpile oil is

$$\frac{1}{dt}E[d(P)] = rP. \quad (16)$$

Intuitively, a firm holding a barrel of stockpiled oil has the option of selling it at instant t or holding it for a brief period, and obtaining a capital gain. The opportunity cost of holding

the inventory is the capitalized value of foregone sales, rP , while the expected capital gain is $\frac{1}{dt}E[d(P)]$. If the latter is not smaller than the former, there will be an incentive to stockpile some ore (Pindyck, 1980, 1982). In light of eqs. (13), (15) and (16), it is apparent that there will be an incentive to stockpile oil when the anticipated rate of change in marginal extraction cost just equals the capitalized level of marginal cost:

$$\frac{1}{dt}E\left[d\left(\frac{\partial c}{\partial y}\right)\right] = r\frac{\partial c}{\partial y}.$$

I show in the appendix that this condition corresponds to

$$-y\left(\frac{\partial^2 c}{\partial y^2}\frac{\partial y}{\partial R} + \frac{\partial^2 c}{\partial y\partial R}\right) + \frac{\sigma^2 P^2}{2}\left[\frac{\partial^2 y}{\partial P^2}\frac{\partial^2 c}{\partial y^2} + \frac{\partial^3 c}{\partial y^3}\left(\frac{\partial y}{\partial P}\right)^2\right] - r\frac{\partial c}{\partial y} = 0. \quad (17)$$

It may seem counter-intuitive that a firm holding both reserves and inventories would be willing to simultaneously extract and stockpile, as *in situ* reserves are higher cost to develop than are stockpiles. Indeed, such simultaneous activities cannot be part of an optimal program under deterministic conditions. But this need not be the case in a stochastic environment. In particular, it can pay the firm to use up its higher cost reserves first, holding the lower cost reserves until a later date when demand is stochastic (Slade, 1988). This is one interpretation of behavior in my model: firms hold onto the lower cost inventory reserves, electing not to sell them until after the higher cost (*in situ*) reserves are exhausted.⁶

To make further headway, I assume that extraction costs are given by the specific functional

form in eq. (3), with $\eta = 2$.⁷ Incorporating this specific form into eq. (17) and simplifying yields

$$\frac{\partial c}{\partial y} \left\{ \frac{y}{R} - \frac{\partial y}{\partial R} + \left(\frac{\sigma^2 P^2}{2y} \right) \frac{\partial^2 y}{\partial P^2} - r \right\} = 0, \text{ or}$$

$$\frac{y}{R} - \frac{\partial y}{\partial R} + \left(\frac{\sigma^2 P^2}{2y} \right) \frac{\partial^2 y}{\partial P^2} - r = 0. \quad (18)$$

Let $\underline{\sigma}^2$ satisfy eq. (18) as an equality. If $\sigma^2 \geq \underline{\sigma}^2$, the anticipated rate of change in marginal extraction costs can equal the capitalized level of marginal costs. In such a scenario the firm has an incentive to acquire and hold stockpiles of oil.

The motive underlying inventory accumulation here is “production smoothing” (Abel, 1985; Arrow et al., 1958; Blanchard and Fisher, 1989). The idea is that when the production cost function is convex, firms can lower the expected discounted flow of costs by using inventories as a buffer, to mitigate abrupt changes in production that are induced by fluctuating demand. In the present case, this motive is offset somewhat by the overall expected downward trend in production associated with a non-renewable resource. Even so, the fundamental wisdom in the literature on inventories can be applied here, given enough variability in demand.

4 Empirical Analysis

The model presented above leads naturally to an empirical investigation. Two types of empirical evidence are germane here. At a somewhat less formal level, if production smoothing motivates inventory holding one would expect greater variation in sales than in extraction. Figure 2 provides visual evidence to this effect. In this figure, I plot monthly observations of sales (the left-most

panel) and extraction (the right-most panel) from January 1981 to December 2020; it is readily apparent that sales rates exhibit greater variability than do extraction rates.

At a more formal level, if production smoothing motivates inventory holding it must be the case that $\sigma^2 \geq \underline{\sigma}^2$. In order to test that condition, one first needs to identify the linkage between optimal extraction and the state variables P and R .

To identify the impact of these state variables upon extraction I utilize data available at the US Energy Information Administration (EIA) website (<http://www.eia.doe.gov/>). There, one can find statistics on spot prices and US crude oil reserves and production. There are three issues that must be confronted. The first issue is that data are only available at the aggregate level, whereas the model above describes motivations to the individual firm. If one assumes that marginal costs are linear, the aggregate results I discuss below map naturally into firm-level implications.⁸

The second issue is the potential endogeneity of one of the key right-side variables, namely price.⁹ To the extent that the endogenous variables in the regressions I report below, US reserves and production, do not influence the world price of crude oil, one can safely ignore the potential endogeneity of price. This seems likely to be the case for at least two reasons. First, US production was a relatively small part of world production during the sample period, and so would seem unlikely to have exerted much impact on global supply. The largest share of world production occurred in 1986; after 1986 the share of US production in total world output fell monotonically until the rise of shale oil in 2009. But even then, most US firms were very small relative to the global market. Second, Adelman (1995) argues that the Organization of Petroleum Exporting Countries (OPEC) has played a significant role in determining price during my sample period. Since OPEC

sets target prices, and associated quotas, based on world market conditions, it seems implausible that they would adjust their actions on the basis of US producer behavior. This point also suggests that US producer behavior is unlikely to exert much influence on the world equilibrium price.

The third issue is that of data frequency: Reserves are reported annually (as of the end of December), while spot prices, production and stockpiles are reported monthly and annually.¹⁰ To match the data on reserves with the monthly data on all other variables of interest, I use a strategy in the spirit of Chow and Lin (1971) and Santos-Silva and Cardoso (2001). I first note that, while the theoretical model I presented above assumed zero net reserve adjustments, $\dot{R} + y = 0$, in practice these adjustments are not identically equal to zero. This is because reserves are regularly adjusted as firms' information concerning their deposits is improved, or as new deposits are discovered. Both improved information and new finds are generally the direct result of drilling oil wells. Baker-Hughes reports the number of oil wells that are actively being drilled, for every week since July 1987.¹¹ Aggregation this data allows one to calculate the number of oil wells drilled a) for each month, and b) for each year. Using the annual construct allows me to correlate drilling (as measured by active oil wells) with reserve adjustments (net of production); that correlation is 0.65.¹² Although data are only available to estimate this correlation at the annual level, that relation ought in principle to also hold true at the monthly level.

To construct synthetic observations on the basis of monthly drilling, I interpolate predicted reserves during the other months. Though data on reserves and production is available for many years, data on drilling is only available after 1987. I therefore apply this method for months between December 1987 and December 2020.¹³ I therefore assigning to each month the share of

reserve changes between observations at the end of December in year $t - 1$ to the end of December in year t equal to the fraction of drilling during year t that occurred during the month in question. Under standard assumptions, this construct is an unbiased estimator of the true (but unobserved) monthly levels of reserves. So long as the measurement error implied by this process is uncorrelated with the disturbance in the regression model for monthly extraction levels, this approach will generate unbiased estimates of the marginal effects of interest.¹⁴ In this way, I am able to splice the synthetic data between the observed values for December in year $t - 1$ and t , for every t from 1989 to 1920.

The next step is to regress the fitted value R , real spot prices P and private inventories on observed oil production. The optimal level of production, as described in (13), should be determined by a balancing of rents with marginal value of reserves. The latter could in principle depend on a combination of prices, inventories and reserves. If marginal costs are inversely related to reserves, optimal output will be proportional to the product of price and reserves, as well as the product of reserves with marginal value $\partial V/\partial R$. To this end, I regress linear and quadratic terms in R , as well as interaction terms involving P, R and S ¹⁵ upon y ; this regression model can be interpreted as a Taylor's series approximation of marginal value multiplied by R . In light of the time-series nature of the data, I allow for serial correlation. To enhance comparability of observations from different months, I convert production into values per day (in millions of barrels). Table 1 presents results from the regression analysis of extraction, for three combinations of variables. Regressions 1 and 2 include linear and quadratic effects for P and R , along with interaction effects including various combinations of P, R and S . Many of these variables are significant in regression 1 (OLS).

That noted, the results from regression 2 indicate serial correlation is quite likely—the first-order parameter, ρ , is very large. The significance of various explanatory variables is less compelling in this regression, with only P, PR, PR^2, PRS and PR^2S having statistically significant coefficients. The specification in regressions 1 and 2 allow the state variable S to play an important role in $\partial V/\partial R$. But one might anticipate R and S entering the value function additively, as they are substitute sources of sales. To investigate this possibility, I ran comparable regressions to 1 and 2 after dropping the seven variables involving S ; these results are presented as regressions 3 and 4. The key point here is that dropping these explanatory variables leads to a substantial increase in summed squared errors; indeed, the hypothesis that all coefficients associated with variables including S are jointly zero is easily rejected. The specification in regression 1 also points to the possible lack of stock effects in costs, in that variables and interactions not including R are present. But the results from regression 2 run counter to this interpretation, as the only significant variables involve R . To delve further into this question, regressions 5 and 6 drop the five variables not including R . Here again this revision to the regression equation significantly raises the summed squared errors, and the hypothesis that the coefficients associated with variables including R are jointly zero is easily rejected.

One might object that other factors such as technology, taxes and distribution costs might influence extraction. As the regressions in Table 1 omit these variables, my results would be suspect if these factors were correlated with my regressors. To check on this possibility, I reran regressions 1, 3 and 5 allowing for yearly fixed effects. Results are given in Table 2 (in the interest of brevity I do not report the annual fixed effects). As in the first set of results, many variables

are statistically significant in the regression that includes all variables, and the variations that omit regressors involving S or not involving R are not supported. In addition, the estimated coefficients are generally similar to those reported in Table 1, as are signs and significance. Finally, the annual fixed effects tend to change over time, with positive effects for earlier years and negative effects for more recent years. This suggests the omitted factors identified above might well matter.

The regression results allow me to estimate the marginal impacts $\frac{\partial y}{\partial R}$, $\frac{\partial y}{\partial P}$ and $\partial^2 y / \partial P^2$, which can then be used to calculate the implied value of $\underline{\sigma}^2$ from eq. (18). I use the parameter estimates from regression 1 allowing for fixed effects, though the broad pattern I describe below holds for other specifications. There are induced values for each month in the sample, so rather than list all these values I present a variety of statistics in Table 4, including the first three quartiles, mean, 90%, and standard errors. I report these values for three values of the real discount rates: 1%, 2% and 3%.

During the sample period, the variance in monthly real spot prices is .2086.¹⁶ It is noteworthy that the sample variance exceeds the mean and median level of $\underline{\sigma}^2$ at each of the three discount rates, as well as the value at the 75% for the small real discount rate. At the medium real discount rate, the implied values of $\underline{\sigma}^2$ exceed the estimated variance σ^2 for slightly less than 75% of the observations. While this evidence is not overwhelming, I think it solidly supports the empirical plausibility of production smoothing as a motive for holding oil inventories.¹⁷

Of course, observing that variations in price are sufficient to motivate production smoothing does not imply there are no other potential explanations for inventory holding. One obvious possibility is that firms hold inventories in order to cash in on unanticipated price increases, whether

they extract more or not in the face of such price increases. Such an explanation has much in common with the idea that wild gyrations in crude prices are related to (and perhaps even caused by) speculation. If such an explanation were correct, one would expect to see sharp increases in crude prices leading to clear reductions in inventories.

Figure 3 shows weekly crude oil prices (using the West Texas Intermediate spot price), measured on the right-most y-axis, and changes in crude oil inventories, measured on the left-most y-axis, for every week from the start of 1986 (when weekly data is first available) to the end of 2020. While significant movements in the weekly changes in stocks occurred throughout this period, crude prices were relatively constant from 1986 to 2000. And while the pattern of changes in stocks is less pronounced after 2001, when crude prices started to vary substantially, there is still no clear indication that changes in stocks are more likely to be negative (respectively, positive) during periods of high (respectively, low) prices.¹⁸ And while it does appear that changes in crude stocks were less volatile after 2000, this does not indicate that agents were more likely to speculate on price changes as spot prices increased. On balance, then, there seems to be little evidence to suggest firms are holding stocks so as to make a killing when prices rise dramatically.

Perhaps speculators held inventories in anticipation of rapidly rising prices, as opposed to basing their decision on current price. If so, it seems plausible that such agents would take their cues from existing futures markets. When futures prices were well in excess of current spot prices, a situation referred to as contango, there would be a motive to buy and hold inventories. To get at this hypothesis, I collected futures data from the EIA webpage, which lists data from four futures contracts. The first of these, “contract 1,” lists futures prices for delivery in the following month.

As this delivery could be within a week or so of the trading date, these futures prices can be very close to current spot prices, particularly as the end of the month approaches. “Contract 2” lists futures prices for delivery in the month after contract 1; “contract 3” is for delivery in the month after contract 2, and “contract 4” is for delivery in the month after contract 3. following month. Since the data from contract 1 seem less likely to produce conditions favorable to speculation, especially at the end of the month, I based the analysis reported below on contract 2 data. Weekly data are available for spot prices, Futures 2 contract prices, and inventory levels from the first week in January, 1986 to the last week in 2020.

While the presence of contango suggests potential benefits from speculation, one also needs to take opportunity costs into account. Irrespective of the presence or absence of holding costs, the ‘buy and hold’ strategy ties up capital resources for a period of time; how long depends on how long the speculator must wait before selling. Accordingly, for each date I calculated the number of weeks until the start of the month in which the contract was to be exercised; this variable is termed “week” in the results reported below. A literal interpretation would set the opportunity cost of tying up capital would be equal to the present value of \$1 received in the future week in question. Under such an interpretation, the net benefit from speculating would be $P_{t,T} - e^{rT} p_t$, where $P_{t,T}$ is the price of a future contract at time t for delivery in T periods and p_t is the spot price in period t . Thus, one could measure of the net benefit from speculating would be $\ln(P_{t,T}/p_t) - rT$. A regression of changes in inventories upon the regressors $\ln(P_{t,T}/p_t)$, and T would then yield coefficients k_1 and k_2 , with k_1 positive (reflecting the sensitivity of stockpiling decisions to potential gains) and k_2 negative (reflecting the opportunity cost of tying up capital while stockpiling). To the extent

this model correctly captures agents' motives, it should also explain much of the variation in stock changes. Alternatively, one might replace the log-ratio of futures to spot price with the difference between future and spot price; the coefficients would take similar interpretations. In the results reported below, I refer to these regressions as 'Regression 1' and 'Regression 2,' respectively.

The results from these three regressions are collected into Table 4. Column 2 reports results from Regression 1, while column 3 reports results from Regression 2; standard errors are listed in parentheses below the corresponding point estimates. One is struck by the relatively poor performance of each regression. While the coefficients have the anticipated signs, and the coefficient on anticipated benefits is statistically important, neither regression explains a meaningful amount of the variation in inventory adjustments. Overall, these results indicate that speculation had little to do with inventory accumulation during the sample period.

5 Conclusion

In this paper, I present a model of firm behavior when oil prices are stochastic. The motive that I focus on is production smoothing, the notion that firms may wish to hold buffer stocks to even out pressures on production. This argument requires that extraction costs are convex (*i.e.*, that marginal costs are upward-sloping).¹⁹ In this framework, the firm has an incentive to hold inventories if prices are sufficiently volatile. Using data on monthly crude prices and privately-held US inventories, I find evidence that there was sufficient volatility in crude prices over the period from January 1986 through December 2009 to motivate inventory holding. By contrast, the evidence that firms held inventories to speculate on price movements does not seem very strong. I believe

the conclusion is that inventories are more likely to be motivated by attempts to smooth marginal production costs than by speculative motives.

My model assumes that the entire cost of production is born at the deposit. In particular, it assumes that extracted oil can instantly and costlessly be delivered to market, and storage of inventories is costless. These assumptions may be legitimately questioned as unrealistic. Shipping costs for crude oil can be a significant share of delivered price, and there is often an important lag between extraction and sale. However, my central findings seem likely to be robust to each of these potential extensions.

Adding distribution costs to the model above has no major effects upon my central results. While such an alteration lowers the expected gains from holding inventories, it has an equivalent effect on current profits. Correspondingly, the key comparison is between the capitalized value of “distribution rents” (price less marginal distribution cost) and the expected rate of change in distribution rents. If the unit cost of distribution is taken as constant, then my model may be applied by interpreting price as distribution rent. This suggests smaller initial sales (and higher initial price) in conjunction with slower growth of prices over time. Such an alteration reduces the value of inventories, but not the finding that sufficient variation in prices will induce firms to hold stockpiles.

Adding storage costs to the model also leaves the central result unchanged. While the presence of storage costs would make it less desirable to hold inventories, there can still be a motive with sufficiently variable demand. The results in Table 3 suggest that demand is often considerably more variable than required to motivate the holding of inventories. Thus, it seems

plausible that the results reported above are robust to storage costs.

It seems most plausible that there is a lag between extraction and sales, as crude oil must be refined prior to delivery of the final good. An extension of the analysis to allow for such lags can be constructed by distinguishing between the date of sales and the date of extraction. Abel (1985) showed that competitive firms would generally have an incentive to hold inventories in the context of lags between production and sales, to facilitate speculation. His results would seem applicable here as well. Indeed, Blanchard and Fisher (1989) suggest that this motive may be at least as important as the production smoothing motive in explaining inventories of most commodities.

6 APPENDIX

To evaluate the anticipated rate of change in marginal extraction costs, $\frac{1}{dt}E[d(\frac{\partial c}{\partial y})]$, I first note that the optimal extraction rate is an implicit function of R, S and P . Applying Ito's Lemma yields

$$\frac{1}{dt}E[d(\frac{\partial c}{\partial y})] = \frac{\partial^2 c}{\partial y^2} \frac{1}{dt}E[d(y^*)] - \frac{\partial^2 c}{\partial y \partial R} y + \frac{1}{2} \frac{\partial^3 c}{\partial y^3} \frac{1}{dt}E[d(y^{*2})]. \quad (19)$$

As y^* is a function of P, S and R , It's Lemma implies

$$\frac{1}{dt}E[d(y^*)] = \frac{\partial y}{\partial S} w - \frac{\partial y}{\partial R} y^* + \frac{1}{2} \frac{\partial^2 y}{\partial P^2} \sigma^2 P^2, \quad (20)$$

$$\frac{1}{dt}E[d(y^{*2})] = (\frac{\partial y}{\partial P})^2 \sigma^2 P^2. \quad (21)$$

If the firm is to be willing to acquire and hold inventories it must be the case that $\partial V/\partial S = P$. Since market price is plainly independent of the firm's reserves one has $\partial^2 V/\partial R\partial S = 0$, in which case $\partial y/\partial S = 0$ and the first term on the right of eq. (20) vanishes. Substituting eqs. (20) and (21) into eq. (19) then yields eq. (17) in the text.

References

- Abel, A. (1985). Inventories, stockouts, and production smoothing, *Review of Economic Studies* **52**: 283–194.
- Adelman, M. A. (1995). *Genie Out of the Bottle: World Oil Since 1970*, MIT Press, Cambridge, MA.
- Arrow, K. J., Karlin, S. and Scarf, H. (1958). *Studies in the mathematical theory of inventory and production*, Stanford University Press, Palo Alto CA.
- Blanchard, O. J. and Fisher, S. (1989). *Lectures on Macroeconomics*, MIT Press, Cambridge, MA.
- Brennan, M. J. and Schwartz, E. J. (1985). Evaluating natural resource investments, *Journal of Business* **58**: 135–158.
- Chow, G. C. and Lin, A.-I. (1971). Best linear unbiased interpolation, distribution, and extrapolation of time series by related series, *The Review of Economics and Statistics* **53**: 372–375.
- Dixit, A. and Pindyck, R. (1993). *Investment under uncertainty*, Princeton University Press, Princeton, NJ.

- Kamien, M. and Schwartz, N. (1991). *Dynamic Optimization: The Calculus of Variations and Optimal Control in Economics and Management*, North Holland, Amsterdam.
- Mason, C. F. (2001). Nonrenewable resources with switching costs, *Journal of Environmental Economics and Management* **42**: 65–81.
- Pindyck, R. S. (1980). Uncertainty and exhaustible resource markets, *Journal of Political Economy* **88**: 1203–1225.
- Pindyck, R. S. (1982). Uncertainty and the behavior of the firm, *American Economic Review* **72**: 415–427.
- Santos-Silva, J. and Cardoso, F. N. (2001). The Chow-Lin method using dynamic models, *Economic Modelling* **18**: 262–280.
- Slade, M. (1988). Grade selection under uncertainty: least cost last and other anomalies, *Journal of Environmental Economics and Management* **15**: 189–205.

Notes

¹ A third explanation is that inventories might be held to insure against running out of the key resource (the so-called “stock-out” motive). It is hard to believe this motive played a major role in the US oil industry, however: average daily input into US refiners during the same period was just over 14 million barrels per day, and never exceeded 16.5 million barrels per day. As such, the stockpile of crude oil would have supplied all US refiners for almost 20 days. This point notwithstanding, the model I discuss below can be adapted to allow for a stock-out motive by including delivery constraints. I discuss this extension in the conclusion.

² If $P_t < \mu_t$ the firm would be inclined to sell all extracted oil along with any accumulated inventories. If any inventories were held the firm's sales rate would then be infinite, which as I note in the text would violate market clearing. But if the firm has never acquired any inventories there is nothing to prevent $P_t < \mu_t$. In fact, this is the most likely outcome in the deterministic framework.

³ While I assume geometric Brownian motion for analytic convenience, a number of previous authors have made similar assumptions (Brennan and Schwartz, 1985; Dixit and Pindyck, 1993; Mason, 2001; Pindyck, 1980; Slade, 1988).

⁴ See Brennan and Schwartz (1985), Dixit and Pindyck (1993) and Mason (2001) for analysis of such a model.

⁵ As rents rise at the rate of interest, while price generally rises less rapidly, it follows that marginal costs must also rise at less than the rate of interest.

⁶ For example, in December of 2008 Royal Dutch - Shell PLC anchored a supertanker full of crude oil off the British coast in anticipation of higher prices for future delivery.

⁷ With these assumptions $\partial^3 c / \partial y^3 = 0$. As $\partial^3 c / \partial y^3$ exerts a positive influence on the expression in eq. (17), one can argue that this assumption generates the least compelling case for holding inventories.

⁸ If marginal costs are linear, aggregate level results map naturally into results at the level of the individual field-reservoir. If one is willing to draw an analogy between individual field-reservoirs and firms these results are directly relevant to the model discussed in section 3 above.

⁹ The regression model I pose below also uses stocks and reserves as explanatory variables. While these variables are also endogenous, in that their paths are steered by the production and sales decisions, I take the view that it is the preceding observation on stocks and reserves that influences current production choices. As such, the regression model uses predetermined endogenous variables as explanatory variables.

¹⁰ Spot prices are also reported on a daily basis.

¹¹ These data can be accessed at <https://rigcount.bakerhughes.com/static-files/6ccf823f-7606-4d18-9c7b-c2c2e8ee50ef>, under the tab labeled “US Oil & Gas Split”.

¹² Alternatively, one could regress reserves reported for year t on reserves for year $t - 1$, production during year t and drilling during year t ; such a regression has an R^2 value of 0.95, and the coefficient on drilling is positive and statistically significant. This alternative method confirms the

economic importance of the Baker-Hughes rig count in predicting oil reserves, and so corroborates the method I describe for constructing synthetic observations on reserves for the eleven months not reported in the annual data.

¹³ I note that there was an idiosyncratic (positive) change between 1969 and 1970, reflecting the addition of Alaskan reserves. As these reserves were firmly in place before the period in which the Baker-Hughes data is available, this idiosyncratic shift has no impact on my analysis.

¹⁴ However, the approach will impact the standard errors. The results reported below are based on robust standard errors, and so correct for this possibility. An alternative approach to the one I use here is to estimate a relation between extraction and prices and reserves using annual data. The disadvantage of using annual data is the corresponding reduction in number of observations. A regression using annual data generates similar results to those reported in Table 1, in the sense that estimated coefficients have the same signs and are generally the same order of magnitude. However, because production data are only available after 1985, only 25 annual observations are available. With such a small data set there are very few degrees of freedom, and none of the coefficients are statistically significant.

¹⁵ The EIA website reports various measures of inventories. Data on stockpiles are available including or excluding the U. S. strategic petroleum reserve (SPR). As the SPR is both publicly held, and hence motivated by political – as opposed to economic considerations – it seems clear that the data excluding the SPR is preferable for my purposes. Netting out SPR holdings yields a measure of private inventory holdings. One might legitimately object to this measure of private

inventories on the grounds it includes stocks held at refineries or oil in pipelines. The first of these which are really more representative of raw materials in the production process, while the second reflect oil in transit; neither of these seems representative of the sort of buffer stock my model envisions. Accordingly, I use data on stocks held at ‘tank farms’, which do seem more representative of the sort of inventories envisioned by the model.

¹⁶ Assuming that prices evolve according to geometric Brownian motion implies that prices are log-normally distributed, *i.e.*, the natural log of prices is normally distributed. During the sample period, the mean and variance of the natural log of real monthly spot prices are 2.894 and .2086, respectively.

¹⁷ As indicated in eq. (10), firms could be motivated to hold inventories even in the absence of stochastic demand so long as the percentage change in production and the ration of production to reserves matched up just right. To investigate this possibility, I formed the discrete time approximation to percentage change in production, $\Delta y_t/y_t = (y_{t+1} - y_t)/y_t$ and the ration of production to reserves, y_t/R_t , for each month t in the sample period. Assuming $\eta = 2$, the construct $\Delta/y + y/2R$ would equal the interest rate. For the sample period, the average value of this construct is 4.2273, implying an interest rate of over 422%. Alternatively, one could run a linear regression of $\Delta y/y$ on $y/2R$; such a regression should yield an intercept of $r/\eta - 1$ and a slope coefficient of $-1/\eta$. Performing such a regression yields estimated coefficients -.0742 for $r/\eta - 1$ and .00822 for $-1/\eta$. Plainly, these results do not lend much support to the notion that firms would have been motivated to hold inventories absent stochastic demand.

¹⁸ In fact, the average change in stocks was negative prior to 2000 (-216,726.2 barrels) and positive after (179,228.6 barrels). This difference in average values, while intriguing, is not statistically significant: the corresponding standard deviations were two orders of magnitude larger. Thus, there is no evidence of significantly smaller values for changes in stocks as prices increased. A similar pattern emerges if one compares levels of inventories against price. During the sample period there are dramatic swings in price, from below 50% of the initial level to over 550% of the initial level. But even with these dramatic swings in price crude stock levels are always within 20% of the initial value. More to the point, there seems to be little evidence that stocks are drawn down during times of particularly large prices, nor are stocks built up during periods of low prices.

¹⁹ A second motive, which I do not explicitly consider, is the “stock out” motive: that firms hold inventories so as to be able to increase deliveries to market in the presence of capacity constraints on production.

Figure 1: US Petroleum Stocks, 1981 – 2020



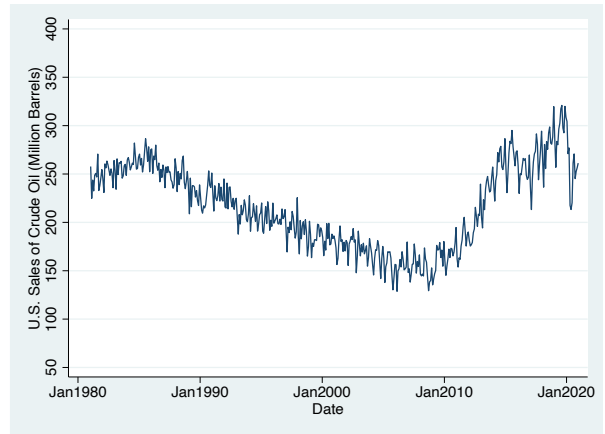


Figure 2: US Petroleum Production and Sales, 1981 – 2020



(a) Production



(b) Sales

Figure 3: Changes in Stocks and Prices

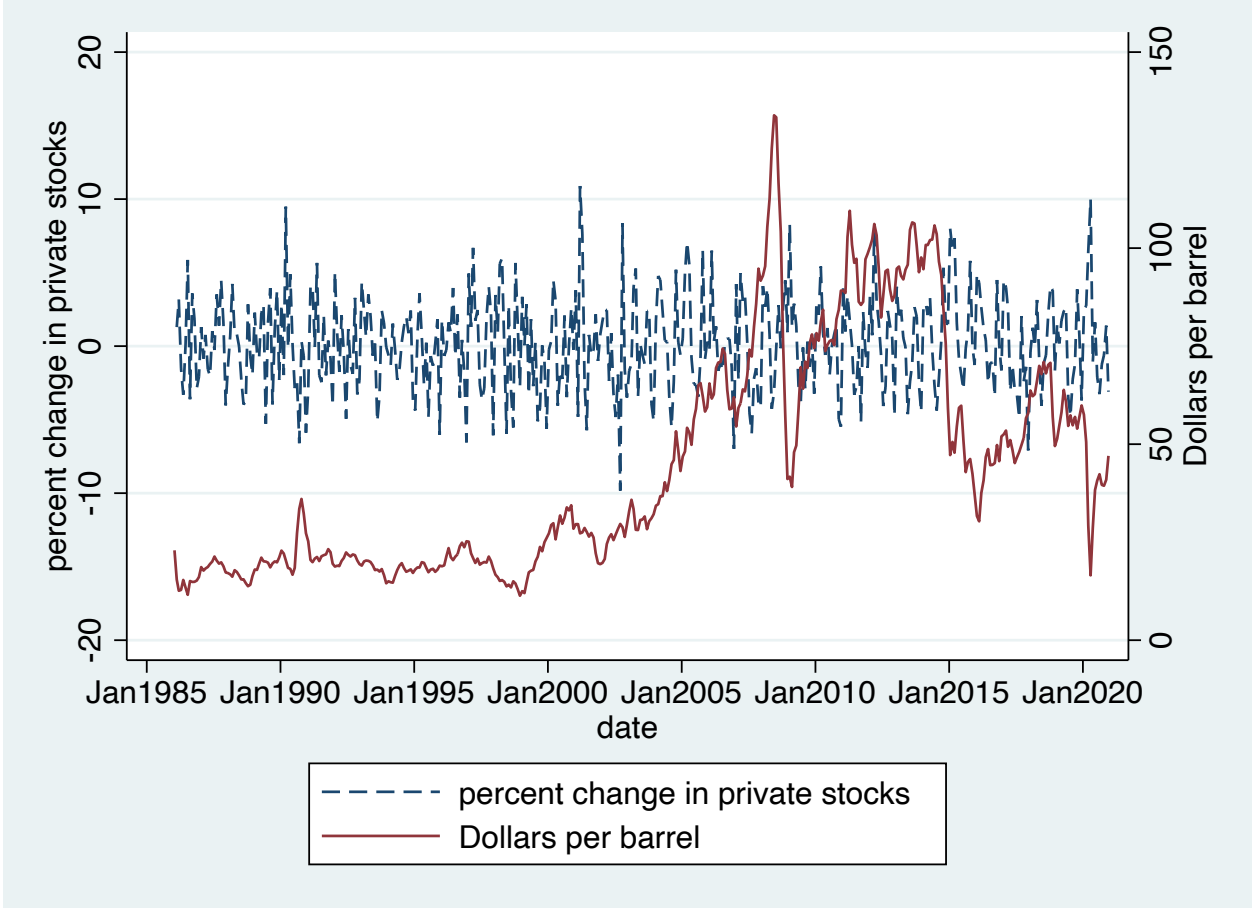


Table 1: Extraction as a function of price, reserves and inventories

variable	Regression					
	(1)	(2)	(3)	(4)	(5)	(6)
<i>R</i>	.0036*** (.00081)	.000062 (.00093)	.0053*** (.00072)	.00031 (.00085)	.00028 (.00023)	-.00054 (.000492)
<i>R</i> ²	-6.56e-08*** (1.54e-08)	4.77e-09 (1.83e-08)	-9.67e-08*** (1.38e-08)	1.7e-09 (1.69e-08)	1.86e-09 (4.52e-09)	1.60e-08 (1.02e-08)
<i>P</i>	10.32*** (3.116)	-4.456* (2.551)	4.464*** (.5827)	.7432 .5814	—	—
<i>P</i> ²	-.0918*** (.0252)	-.01590 (.0201)	-.0205 (.0052)	-.00938** (.00414)	—	—
<i>PR</i>	-.00067*** (.000213)	.00035** (.00018)	-.00036*** (.00004)	-.000058 (.000045)	-.000052*** (7.98e-06)	-9.41e-06 (6.74e-06)
<i>PR</i> ²	8.17e-09** (3.56e-09)	-7.01e-09** (3.0e-09)	6.8e-09*** (7.83e-10)	1.01e-09 (8.45e-10)	1.72e-09*** (3.0e-10)	8.44e-11 (2.40e-10)
<i>P</i> ² <i>R</i>	4.32e-06*** (1.18e-06)	8.31e-07 (9.4e-07)	1.0e-06*** (2.44e-07)	4.64e-07** (1.95e-07)	1.38e-0*** (5.11e-08)	5.50e-08 (3.95e-08)
<i>PS</i>	-.000061** (.000021)	.000028 (.000018)	—	—	—	—
<i>PS</i> ²	1.13e-10 (3.42e-11)	-1.9e-11 (2.69e-11)	—	—	—	—
<i>P</i> ² <i>S</i>	4.22*** (1.3e-07)	6.45e-08 (1.09e-07)	—	—	—	—
<i>PRS</i>	3.31*** (1.24e-09)	-2.06e-09* (1.09e-09)	—	—	2.65e-10*** (5.13e-11)	1.88e-11 (4.51e-11)
<i>PRS</i> ²	-5.41e-15*** (1.55e-15)	6.77e-16 (1.22e-15)	—	—	-1.7e-16* (9.19e-17)	-8.67e-17 (7.02e-17)
<i>P</i> ² <i>RS</i>	1.21e-15 (6.09e-12)	-3.36e-12 (5.08e-12)	—	—	-5.0e-13* (2.88e-13)	-1.87e-13 (2.29e-13)
<i>PR</i> ² <i>S</i>	-1.9e-14 (1.65e-14)	3.95e-14** (1.55e-14)	—	—	-7.96e-15*** (1.49e-15)	1.30e-15 (1.31e-15)
constant	-40.49*** (10.27)	2.767 (11.65)	-62.45*** (11.59)	-.8895 (10.48)	-.1283 (2.889)	10.61* (5.915)
ρ	—	.926	—	.920	—	.924
SSE	17.44	4.869	20.36	5.327	21.787	5.088
R-squared	.948	.685	.940	.674	.936	.676

Note: standard errors in parentheses
 dependent variable: production, million barrels per day
 number of observations = 396
 *: significant at 10% level
 **: significant at 5% level
 ***: significant at 1% level

Table 2: Fixed effects regression of extraction

variable	Regression		
	(1)	(2)	(3)
<i>R</i>	.00041*** (.00062)	.00039*** (.000058)	.00038*** (.000044)
<i>R</i> ²	-5.27E-09** (2.43e-09)	-4.36e-09* (2.34e-09)	-4.06** (1.65e-09)
<i>P</i>	3.862 (2.481)	.4060* (.2473)	—
<i>P</i> ²	-0.06388*** (.0179)	-.00124 (.00312)	—
<i>PR</i>	-0.00023 (.00017)	-.000041** (.000018)	-.000016** (6.67e-06)
<i>PR</i> ²	2.21E-09 (3.04e-09)	9.43e-10*** (3.27e-10)	2.99e-10 (2.31e-10)
<i>P</i> ² <i>R</i>	3.02E-06*** (8.4e-07)	7.86e-08 (1.47e-07)	4.33e-08 (3.85e-08)
<i>PS</i>	-.000026 (.000017)	—	—
<i>PS</i> ²	4.17E-11* (2.53e-11)	—	—
<i>SP</i> ²	3.26E-07*** (9.14e-08)	—	—
<i>SPR</i>	1.35E-09 (1.03)	—	7.22e-11 (4.62e-11)
<i>S</i> ² <i>PR</i>	-2.05E-15* (1.15e-15)	—	-1.53e-16** (7.26e-17)
<i>SP</i> ² <i>R</i>	-1.53E-11*** (4.27e-12)	—	-1.45e-13 (2.23e-13)
<i>SPR</i> ²	-5.23E-15 (1.54e-14)	—	-1.08e-16 (1.28e-15)
SSE	5.907	6.510	6.322

Note: standard errors in parentheses
 dependent variable: production, million barrels per day
 number of observations = 396
 *: significant at 10% level
 **: significant at 5% level
 ***: significant at 1% level

Table 3: Implied lower bounds on variance in price

statistic	$\partial y / \partial R$	y/R	$(\partial^2 y / \partial P^2) \div (y/P^2)$	$\underline{\sigma}^2(.01)$	$\underline{\sigma}^2(.02)$	$\underline{\sigma}^2(.03)$
25%	0.000257	0.000195	0.0176	0.0297	0.0600	0.0918
mean	0.000275	0.000256	0.0527	0.0695	0.1417	0.2141
median	0.000281	0.000233	0.0261	0.0642	0.1295	0.1927
75%	0.000293	0.000333	0.0535	0.0922	0.1874	0.2825
90%	0.000300	0.000378	0.1111	0.1248	0.2577	0.3911
s.d.	0.000023	0.000117	0.0753	0.0525	0.1086	0.1649

Note: $\underline{\sigma}^2(r)$ listed for annual discount rates: $r = .01$, $r = .02$ and $r = .03$
variance of monthly real spot price during sample period = .2086

Table 4: Analysis of Contango

variable	Regression 1	Regression 2
$\ln(P_{t,T}/p_t)$	8987.3** (4163.5)	—
$P_{t,T} - p_t$	—	296.71** (111.42)
T	-136.4 (85.790)	-132.29 (85.738)
constant	846.06* (524.29)	779.66 (573.31)
R^2	.0052	.0064

Note: left-hand side variable equals weekly changes in inventories. Number of observations = 1767.