

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 18, 2009

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: Petroleum Economics, Stochastic Dynamic Optimization

JEL Areas: Q2, D8, L15

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

1 Introduction

Since official records were first kept in the United States (U.S.), 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 per day. 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 abnormally 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.

Reacting to the volatile changes in petroleum prices during the past year, some key players in OPEC and a number of members of the U.S. Congress placed the blame on speculators. One issue left unanswered in this dialog was 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 explain 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)). Suppose that oil prices are driven by a random process that arises 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:

$$P_t - \frac{\partial c}{\partial y} - \lambda_t = 0; \quad (5)$$

$$P_t - \mu_t \begin{cases} > 0 & \Rightarrow w_t = 0 \\ = 0 & \Rightarrow w_t \text{ is indeterminate.} \\ < 0 & \Rightarrow w_t = -\infty \end{cases} \quad (6)$$

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. The firm's optimal time path of w must therefore be based on the first two lines of (6).

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

$$\begin{aligned} d\left[P_t - \frac{\partial c}{\partial y}\right]/dt &= \dot{\lambda} = r\left[P_t - \frac{\partial c}{\partial y}\right] + \frac{\partial c}{\partial R}, \text{ or} \\ \left[\dot{P}_t/P_t - r\right]P_t &= d\left[\frac{\partial c}{\partial y}\right]/dt - r\frac{\partial c}{\partial y} + \frac{\partial c}{\partial R}. \end{aligned} \quad (9)$$

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

$$\begin{aligned} d\left[\frac{\partial c}{\partial y}\right]/dt - r\frac{\partial c}{\partial y} + \frac{\partial c}{\partial R} &= 0 \iff \\ (\partial^2 c/\partial y^2)y - (\partial^2 c/\partial y\partial R)y - r\frac{\partial c}{\partial y} + \frac{\partial c}{\partial R} &= 0. \end{aligned}$$

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:

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

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.²

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):

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

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

$$-P_t + \partial V / \partial S \geq 0. \quad (14)$$

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.³

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 deterministic framework, the next step in the analysis 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.⁴⁴ 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 unit of stockpiled ore 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 ore when the anticipated rate of change in marginal extraction cost just equals the capitalized level of marginal cost:

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

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.⁵

To make further headway, I assume that extraction costs are given by the specific functional form in eq. (3), with $\eta = 2$.⁶ 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. For production smoothing to motivate 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 U.S Energy Information Administration (EIA) website (<http://www.eia.doe.gov/>). There, one can find statistics on spot prices and U.S. 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. In light of my assumption that $\eta = 2$ marginal costs are linear. Consequently, the aggregate results I discuss below map naturally into firm-level implications.

The second issue is the potential endogeneity of one of the key right-side variables, namely price. To the extent that the endogenous variables in the regressions I report below, U.S. 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, the market for oil is international, and producer behavior within the U.S. seems unlikely to exert much impact on global supply. During the period from 1986-2007, the combined market share of U.S. production

in total world output fell monotonically over the sample period from a high of 16.9% in 1986 to a low of 8.4%. It seems clear that individual U.S. producers were too small to influence the world price of crude. Second, Adelman (1995) argues that the Organization of Petroleum Exporting Countries (OPEC) has played a significant role in determining price during the last twenty years, which corresponds to 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 U.S. producer behavior. This point also suggests that U.S. 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, while spot prices, production and stockpiles are reported monthly and annually.⁷ 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 regarded reserves as deterministic, in practice they are not because of new discoveries of oil. These arise from exploratory efforts, which the EIA measures by "exploratory wells." They further distinguish "dry exploratory wells," which presumably impact the probability that future exploration will be successful (and hence impact the current predictions of available oil). Data on both exploratory wells and dry exploratory wells are available at the monthly and annual level. The relation between exploratory efforts - both successful and failed - should have a bearing upon reserves at any level of frequency. Thus, while data are only available to identify a linkage between exploration and reserves at the annual level, that relation ought in principle to also hold true at the monthly level. I therefore start by estimating an empirical model of reserves based upon exploration using annual data, and then employ this regression model to produce synthetic data for reserves at the monthly level. This latter data is then exploited, along with the monthly data on

oil prices and production, to estimate a relation between y and the state variables R and P , at the monthly level.

The data on reserves, in millions of barrels, is reported as of 31 December in each calendar year from 1949 to 2007. Owing to the substantial find in Alaska in the summer of 1968 there is a dramatic jump in reserves for 1970. In light of this singularly important event, I include a dummy variable that takes the value '1' for years prior to 1969 and '0' for all years after 1968.⁸ In the regressions reported below, this dummy variable is termed 'Alaska.' One also expects that lagged reserves are important, in that reserves are a state variable (and hence evolve slowly). To that end, I regress lagged reserves, exploratory wells, dry exploratory wells and the dummy variable 'Alaska' upon current reserves. The results of that regression are reported in Table 1. The key points are that the regression fits the data very well, and that the coefficients on lagged reserves, exploratory wells and dry wells are statistically important.

Based on the results from this regression, I then form the best-fit variable \hat{R} . Because data on spot prices are only available starting in 1986, the role of the Alaska dummy variable is moot. I use the value for reserves as of 31 December 1985 for the value of reserves at the start of the monthly dataset (January 1986), and form the fitted value \hat{R} recursively based on exploratory data. Under standard assumptions, \hat{R} 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.⁹

The next step is to use the fitted value R and the data on spot prices P , along with the induced data on monthly oil production, y , to estimate the effect of reserves and prices upon production. To this end, I regress linear and quadratic terms in R and P , as well as interaction terms, upon y ; this

regression model can be interpreted as a Taylor's series approximation. One particular advantage of this specification is that the induced marginal effect of reserves upon extraction is linear. As such they 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.

Table 2 presents results from the regression analysis of extraction. Here, the units of measurement for production and reserves are millions of barrels per month. All variables aside from P and P^2 are significant at the 5%-level or better; the price variables are marginally insignificant at the 5%-level but significant at the 10%-level.

The results from this regression allow me to estimate the marginal impacts $\frac{\partial y}{\partial R}$, $\frac{\partial y}{\partial P}$ and $\partial^2 y / \partial P^2$; using that information I calculate the implied value of $\underline{\sigma}^2$ from eq. (18). 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 3, including the first three quartiles, mean, 90%, 95% and 99% values, as well as the maximum value.

During the sample period the variance in monthly spot prices is .285544. The relevant unit of time for my purpose is daily spot price. Since the variance of daily spot price during a month is about 21 times the variance of the average monthly spot price, one estimate of the variance in daily spot price as 5.9964.¹⁰ It is noteworthy that the implied daily variation in spot prices exceeds the 99% level of $\underline{\sigma}^2$ at each of the three discount rates; indeed, it is almost as large as the largest value observed at a discount rate of 2.5%. This, I think, is strong evidence supporting the empirical plausibility of production smoothing as a motive for holding oil inventories.¹¹

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 pos-

sibility 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. Indeed, the correlation coefficient between crude stocks and prices is negative (-.241), and so is at least broadly consistent with this story. But closer examination reveals that crude stock levels and crude prices are not closely linked.

Figure 2 shows weekly crude stock and price levels, as ratios of their respective values at the start of 1986 (when weekly data is first available). Over this time period, there are dramatic swings in price: at one point, crude prices fell to below 50% of the initial level. More recently, prices have increased sharply, to over 550% of the initial level. But even with these dramatic swings in price crude stock levels are far less volatile; in fact, they 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.

The correlation is also small if one considers the relation between changes in crude stocks and prices (.0297). As Figure 3 illustrates, there are significant movements in the weekly changes in stocks during the period from 1986 to 2000, despite the fact that crude prices were relatively constant during that period. While the pattern of changes in stocks is less pronounced after 2001, when crude prices started to rise, there is still no clear indication that changes in stocks are more likely to be negative during periods of high prices.¹² On balance, then, there seems to be little evidence to suggest firms are holding stocks so as to make a quick killing when prices rise dramatically.

Finally, one might wonder if speculators held inventories in anticipation of rapidly rising

prices (as opposed to basing their decision on current levels or changes). Plausibly, 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. This web site 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 first week in February, 2009.

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 a strict interpretation, one measure of the net benefit from speculating would be $\ln(P_{t,T}) = T \ln(p_t)$, 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 . Under this interpretation, a regression of changes in inventories

upon the regressors $\ln(P_{t,T})$, $T\ln(p_t)$, and T would yield coefficients k_1, k_2 and k_3 , with k_1 positive and the other two negative; it would also explain much of the variation in stock changes. (A literal interpretation of the coefficients would be $k_2 = -(1+r)$ and $k_3 = c$, where r is the market interest rate and c is the unit cost of holding inventory for a week). A less strict interpretation would regress changes in inventories upon the regressors $\ln(P_{t,T}/p_t)$ and T , where presumably the coefficient on the former would be positive (reflecting the sensitivity of stockpiling decisions to potential gains) and the coefficient on T would be negative, reflecting the opportunity cost of tying up capital while stockpiling. 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,' 'Regression 2' and 'Regression 3,' respectively.

The results from these three regressions are collected into Table 4. Column 2-4 report results from, respectively, Regressions 1, 2 and 3; standard errors are listed in parentheses below the corresponding point estimates. One is struck by the poor performance of each regression. Indeed, the only variable that exerts a statistically significant effect is the difference between future and spot price, as reported in Regression 3; even here the significance is only at the 10% level. None of the three regressions explain any meaningful amount of the variation in inventory adjustments. Overall, one is left with the inescapable conclusion that speculation has very little to do with inventory accumulation.

5 Conclusion

In this paper, I present a model of firm behavior when oil prices are stochastic. 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 U.S. inventories, I find evidence that there was sufficient volatility in crude prices over the period from early 1986 to late 2008 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 mine. In particular, extracted ore 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.




Figure 1: U.S. Petroleum Stocks

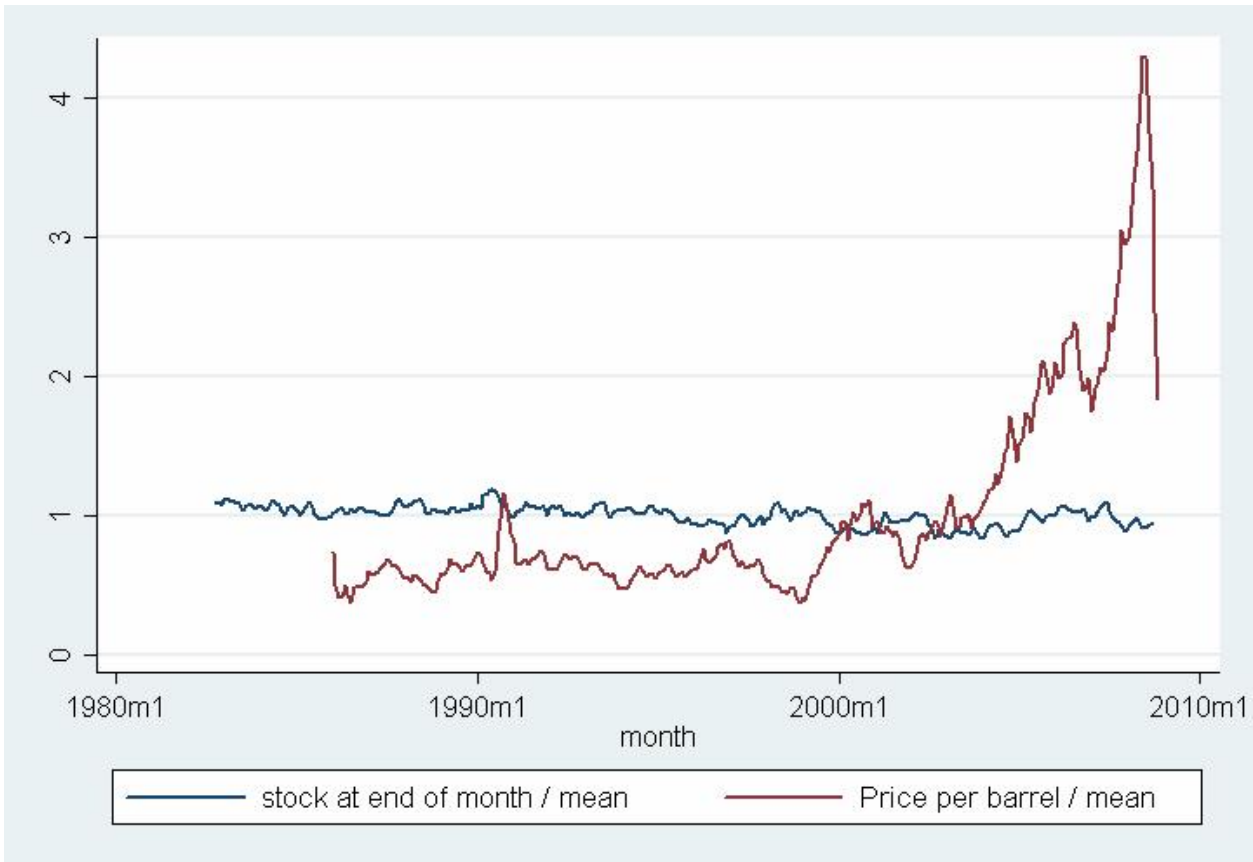


Figure 2: Crude Prices and Stocks




Figure 3: Changes in Stocks and Prices

Table 1: Reserves as a function of exploratory drilling

variable	coefficient	std. err.	t-stat
lagged reserves	.6430	.0934	6.88
Alaska	-1019.4	680.55	-1.50
exploratory wells	-4.642	1.412	-3.29
dry wells	1.159	.3342	3.47
constant	7940.5	2176.4	3.65

number of observations = 58

R-squared = .9097

Table 2: Extraction as a function of price and reserves

variable	coefficient	std. err.	t-stat
P	1269.8	237.93	5.34
R	2.5411	.3309	7.68
P^2	-.4956	.2592	-1.91
PR	-.1080	.0203	-5.32
R^2	-5.22e-05	7.01e-06	-7.45
P^2R	2.19e-05	1.15e-05	1.91
PR^2	2.29e-06	4.32e-07	5.30
constant	-30618.3	3905.21	-7.84

number of observations = 273

R-squared = .7586

Table 3: Implied lower bounds on variance in price

statistic	$\frac{\partial y}{\partial R}$	y/P^2	$\partial^2 y / \partial P^2$	$\underline{\sigma}^2(.01)$	$\underline{\sigma}^2(.025)$	$\underline{\sigma}^2(.05)$
25%	0.08	0.18	0.02	0.68	0.77	0.96
mean	0.08	0.46	0.03	1.54	1.76	2.14
median	0.09	0.46	0.03	1.54	1.78	2.17
75%	0.09	0.65	0.03	2.25	2.57	3.12
90%	0.10	0.92	0.04	3.04	3.47	4.17
95%	0.10	1.12	0.04	3.41	3.89	4.67
99%	0.10	1.61	0.13	4.34	4.94	5.94
maximum	0.10	2.00	0.25	5.42	6.17	7.43

note: $\underline{\sigma}^2(r)$ listed for $r = .01$, $r = .025$ and $r = .05$

Table 4: Analysis of Contango

variable	Regression 1	Regression 2	Regression 3
$\ln(P_{t,T})$	-507.47 (1037.9)	—	—
$T \ln(p_t)$	85.619 (169.17)	—	—
$\ln(P_{t,T}/p_t)$	—	5981. (4518.6)	—
$P_{t,T} - p_t$	—	—	248.36* (142.35)
T	-269.46 (559.79)	5.6496 (93.293)	4.6642 (93.24)
constant	1614.0 (559.79)	11.8372 (574.37)	4.7097 (573.31)
R^2	.0002	.0015	.0025

number of observations = 1205

*: statistically significant at the 10% level

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 U.S. oil industry, however: average daily input into U.S. 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 U.S. refiners for almost 20 days.

² See Brennan and Schwartz (1985), Dixit and Pindyck (1993) and Mason (2001) for analysis of such a model.

³ Shadow values in excess of current price would induce the firm to discontinuously reduce its inventory stock, which would imply selling at an infinite rate. As discussed above, this cannot occur because of market clearing.

⁴ 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.

⁵ 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.

⁶ 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.

⁷ In fact, spot prices are reported on a daily, weekly, monthly and annual basis. Data on stock-

piles 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.

⁸ While the deposit was discovered in 1968 it is arguably true that the full extent of the find was not known for several months. As such, it would not impact the reported level of reserves until after 31 December 1968 and hence would not alter reported reserves for 1968.

⁹ An alternative approach is to use annual data to estimate a relation between extraction and prices and reserves. 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 2, in the sense that estimated coefficients have the same signs and are generally the same order of magnitude. However, production data are only available since 1986, and so there are only 22 annual observations. With such a small data set there are very few degrees of freedom, and none of the coefficients are statistically significant.

¹⁰ The variance of x_a , the average of N observations x_i with common variance σ^2 , is $\sigma_a^2 = \sigma^2/N$. In a typical year there are 254 trading days (365 days, less 104 weekend days, less 7 holidays), which corresponds to just over 21 trading days per month. With $\sigma_a^2 = .285544$ one may estimate σ^2 by $21 \times .285544 = 5.9964$.

¹¹ 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 $\Delta y/y + y/2R$ 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.

¹² 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. The important point here is that there is no evidence of significantly smaller values for changes in stocks as prices increased. It does appear that changes in crude stocks were less volatile after 2000, but again this does not indicate that agents were more likely to speculate on price changes as spot prices increased.