RUTCOR
RESEARCH
REPORT

OPTIMAL MANAGEMENT OF OIL LEASE INVENTORY: OPTION VALUE AND NEW INFORMATION

Michael H. Rothkopf[a] Radford Schantz[b]
Lee Upton, III[c]

RRR 22-2006, SEPTEMBER 2006

RUTCOR
Rutgers Center for
Operations Research
Rutgers University
640 Bartholomew Road
Piscataway, New Jersey
08854-8003
Telephone: 732-445-3804
Telefax: 732-445-5472
Email: rrr@rutcor.rutgers.edu
http://rutcor.rutgers.edu/~rrr

[a] RUTCOR, Rutgers University, 640 Bartholomew Road, Piscataway, NJ 08854-8003
[b] U.S. Interior Department, Herndon, VA 22090
[c] Innovation and Information Consultants, Inc., Concord, MA 01742

RUTCOR RESEARCH REPORT

RRR 22-2006, SEPTEMBER 2006

OPTIMAL MANAGEMENT OF OIL LEASE INVENTORY: OPTION VALUE AND NEW INFORMATION

Michael H. Rothkopf Radford Schantz Lee Upton, III

Abstract. Companies buy and hold many Gulf of Mexico oil leases that expire unexplored. Indeed, they spend more on deep-water leases they never drill than they do on ones they do drill. Hence, rational explanation of their lease purchasing depends upon option value. We propose a simple new model based upon changes in the likelihood that a tract contains oil rather than upon changes in the price of oil. In it, it is rational to buy more leases than will be drilled, new information about the potential that a tract has producible oil arrives randomly, and leases are options to explore. The model optimizes the number of leases an individual firm should add to inventory, should hold in inventory, should drill, should farm out, and should allow expire. Further, it shows how optimal inventory management responds to government polices. This helps explain the record of lease sales following the reduction in royalties for deepwater leases provided by the Deepwater Royalty Relief Act of 1995.

Note: Please do not quote or cite this paper without permission of authors. Their opinions do not necessarily represent US government policies.

Acknowledgements: The research was funded by Minerals Management Service contract number 71722, "Effects of Royalty Incentives for Gulf of Mexico Oil and Gas Leases," final report November, 2004.

1 Introduction

This paper focuses on limited-duration oil and gas leases in the U.S. Gulf of Mexico. It presents a simple model of the decision to buy such leases.

By way of background, the U.S. government offers for sale offshore oil leases in the Gulf of Mexico annually. Generally, each year all unleased tracts are offered for lease. Offshore oil leases usually cover a block that is 5 miles square and have a primary term, meaning the period in which the lessee has exclusive mineral rights at the block covered by the lease. The lessee can pursue activities as he chooses in the primary term, paying a small annual rental if there is no production, but the lease can be held for a longer period only if it is producing oil or gas. The terms of the leases vary by location and the terms of each particular sale. Normally, leases in water depth of 0-200 meters have a primary term of 5 years and a royalty of 1/6 of production. Leases in 200-800 meters expire unless an exploration well is drilled within 5 years, and production must begin within 8 years from issuance to hold the lease; the royalty is1/8. Leases in over 800 meters have a primary term of 10 years and a royalty of 1/8. The Minerals Management Service (MMS), a government bureau, sells available leases by means of a competitive, sealed-bid auction in which a cash bonus is the bidding variable.

The original impetus for the modeling reported in this paper was a research project looking into the effects of deepwater royalty relief. The reason for the focus on lease sales in this paper is the fact that, for several lease sales directly following the adoption of the Deepwater Royalty Relief Act of 1995 (DWRRA), there was a massive increase in the number of leases sold, particularly in water depths greater than 200 meters. One impact of royalty relief may have been to increase the benefits of developing deepwater leases; as a result bidders may have been inclined to buy more leases reflecting the high profit potential afforded by royalty relief. Yet, that explanation alone is insufficient: it is very unlikely that all, or even most, of these leases will be explored before they expire. Other factors may have also influenced the number of leases sold during this period including improved information about the quality of leases in certain areas, the availability of new technology, and the number of leases sold in prior sales. Since many leases are held for their full term, not drilled, and instead allowed to expire, any rational explanation of lease purchasing must depend upon option value.

After giving a brief overview of trends over time in the number of leases sold by water depth, we employ an options approach to lease inventory theory to explain the number of leases sold. We use it to examine the role of DWRRA.

2 Trends in Leases Sold

We examined the time trend of leases sold since 1983 when MMS shifted to the current area-wide leasing system[1] using the water depth categories of 0-200 meters, 200-800 meters and over 800 meters.

[1]All of our data and statistical analysis use 1983 as the starting point. The shift to area-wide leasing was such a dramatic change to the manner in which MMS conducted its lease sales that we believe it would be inappropriate to compare the attributes of leasing such as number of leases sold or high bid with sales from the earlier lease-nomination period.

Figure 1 shows the number of leases sold between 1983 and 2003 broken out by water depth. As can be seen, during the 1980s and early 1990s the number of leases sold followed a cyclical pattern, possibly reflecting the lease term in effect for each water depth. This is especially apparent for shallow water leases with a five-year term. In the period following the oil price crash in 1986, there was an increase in deepwater leasing activity, followed by a significant decline. This increase in 1987 and 1988 may have been caused by the announcement of several new discoveries in deepwater areas ranging up to 1,600 meters in depth. In 1987, there was also a shift in the MMS leasing program with a reduction in the minimum amount that companies were required to bid on a per acre basis.

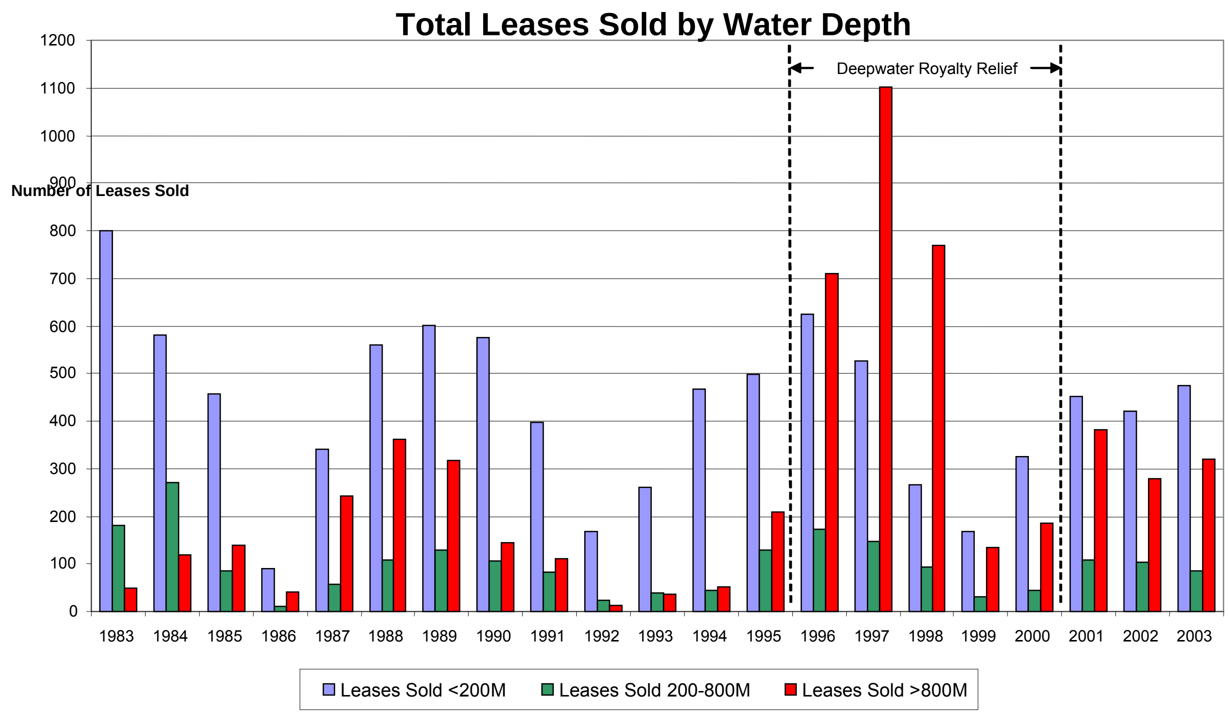

Figure 1.

In 1995, just prior to the passage of DWRRA, there was an increase in the number of deepwater leases sold. Sales of shallow water leases also increased. Then, as the figure shows, after passage of DWRRA, there was a large increase in the number of leases sold in water depths greater than 800 meters. After three years, however, there was a very significant decline in the number of leases sold, including in deepwater areas. Then, after the expiration of DWRRA and imposition of the administrative royalty relief program (which we label "post-DWRRA"), there was another increase in leasing activity. This increase appeared to be somewhat more uniform across different water depths, and perhaps reflected a general recovery in oil and gas prices

commencing in 2000-2001 period. Figure 2 presents data on the outstanding lease inventories by water depth.

Based on our review of the data and the prior literature as well as our knowledge of leasing behavior, we developed various hypotheses of what influenced the number of leases sold. The number of leases sold clearly varies from lease sale to lease sale and fluctuates among various water depth categories. The number of leases sold also seems to change as new government initiatives are introduced: area-wide leasing, reduction in minimum bid per acre, DWRRA, administrative royalty relief.

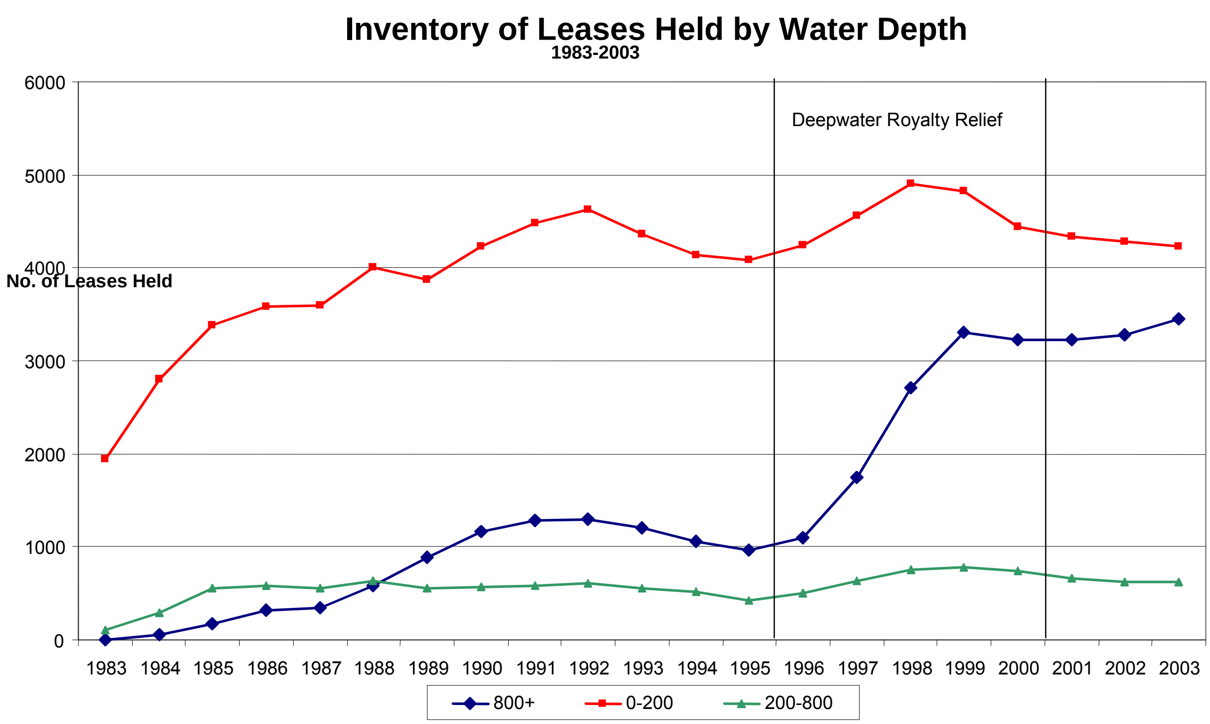

Figure 2.

To a remarkable degree, many leases bought and held in inventory are never drilled. While we have not tried to estimate the precise acquisition and carrying cost of inventory, the amounts of money involved in buying offshore leases that are never drilled are substantial by any measure. These costs are illustrated in the following three figures. These figures illustrate the case of the Central planning area of the Gulf of Mexico; illustrations of the Western area are similar.

For shallow water, figure 3 shows the amount of money spent in each lease sale since 1983 to acquire leases. The cohort of leases purchased is divided into three categories reflecting how those leases were used in the ensuing years: those that were drilled (i.e., exploration, regardless of whether any production ensued); those that were surrendered without having been drilled at all; and those that are still held undrilled, termed "available." It is evident that, since 1987, the majority of the money spent on leases was spent on leases that were never drilled. For the two deepwater classes shown in figures 4 and 5, the pictures lead to a similar conclusion.

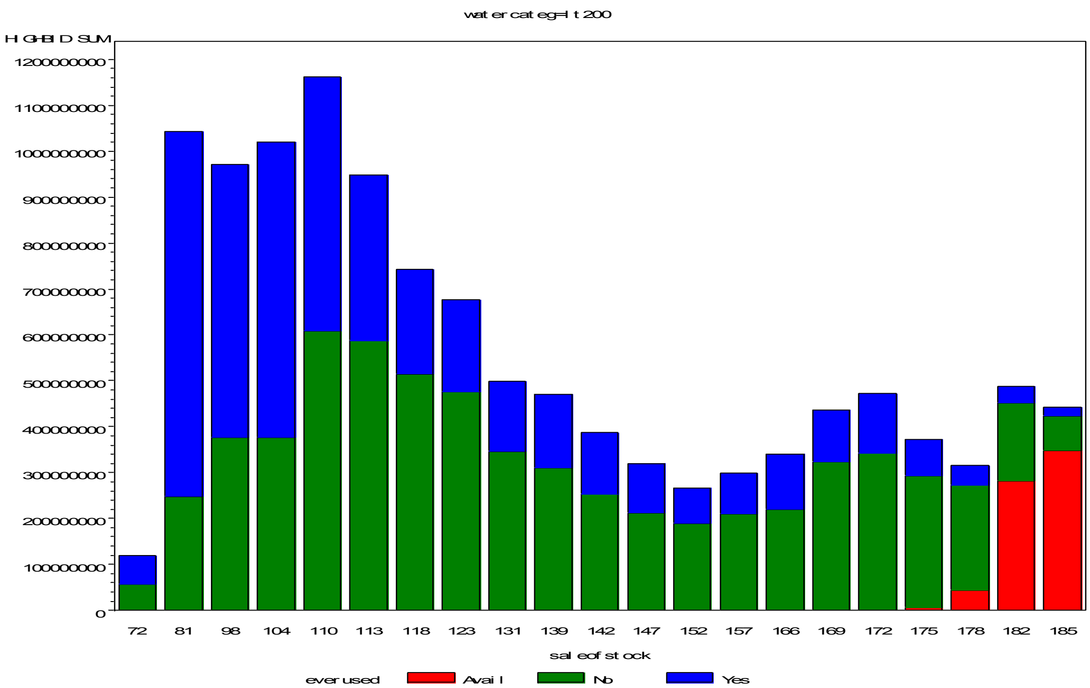

Figure 3.

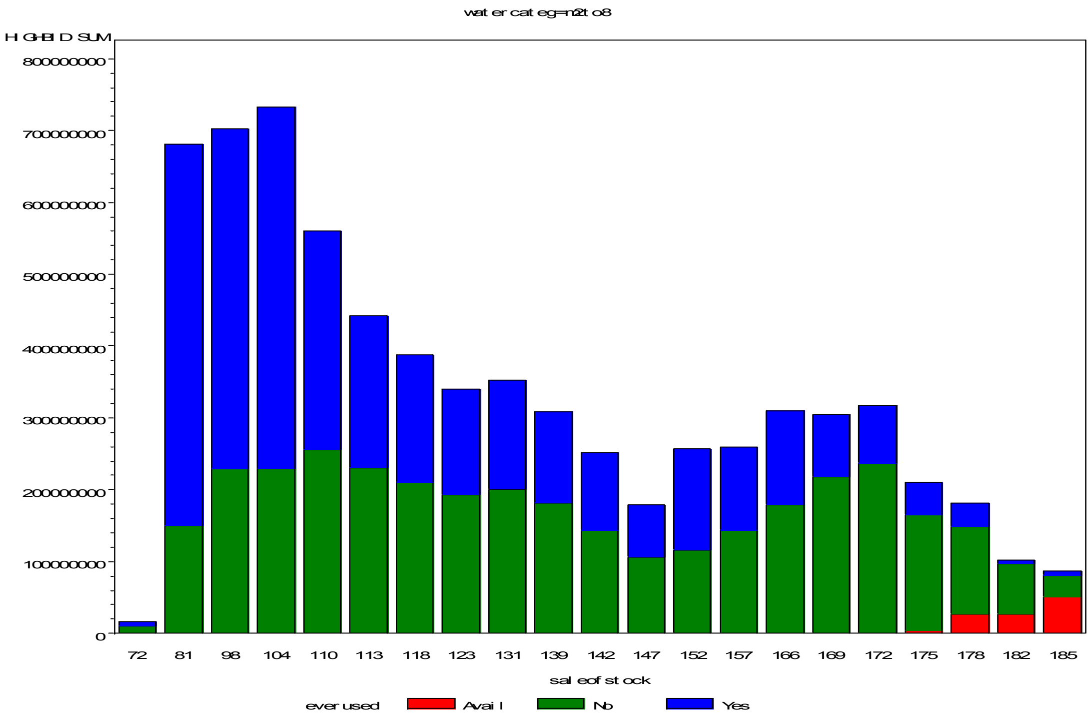
800000000
700000000
600000000
500000000
400000000
300000000
200000000
100000000
0
72
81
98
104
110
113
118
123
131
139
142
147
152
157
166
169
172
175
178
182
185
Yes

Figure 4.

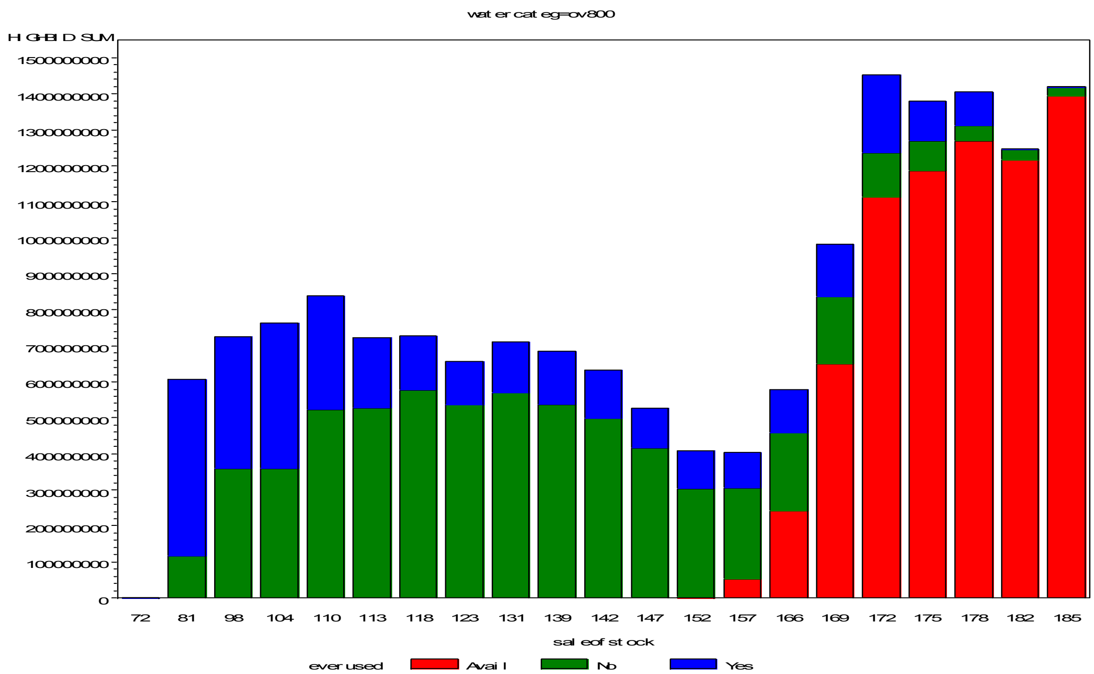

Figure 5.

3 Option Value of Future Information

Since most of the money spent to acquire leases is spent on leases never drilled, we needed to develop a theory of lease inventory behavior that incorporates an options approach. It is clear that firms operating in the Gulf of Mexico often buy the drilling rights to leases that they never drill. Any attempt to explain this behavior as rational for the firms involved will have to base the explanation upon the uncertainty faced by the firms.[2] A way to do this is to justify the purchase of a lease as an option to drill rather than as a purchase of a lease currently worth drilling. This makes particular sense in deepwater where the cost of drilling a lease is often much higher than the cost of the lease itself.

There are at least two potential uncertainties that might support an option-to-drill model of lease inventories. One of these is uncertainty about oil price. This approach is tempting in that it treats options to drill analogously to financial options to purchase a stock, something that has been well studied. If that analogy held, it would allow the application of the substantial amount of theoretical work that has been done on financial options. However, that approach has two serious problems.

---

[2] Another possibility, one we do not favor or pursue, is that while the firms do not behave rationally, the executives involved are responding rationally to incentives within the firms. The amount of money at stake makes it likely that at some level within successful firms controls will be applied to limit the irrationality of the firm's behavior.

The first problem is that, unlike financial stocks, which can be sold immediately after the exercise of an option, oil in the ground needs to be extracted over a period of decades. Thus, if a firm “exercises its option” by drilling the tract, it cannot sell most of the oil for many years. If oil prices were a true random walk, it could be argued that the current oil price is an unbiased estimate of future oil prices and thus reflects the future value of the oil in the ground. However, oil prices have been shown to be, most likely, a mean reverting process rather than a true random walk.[3] This implies that the current price of oil is not an unbiased estimate of future oil prices and, and in particular, that if the price of oil goes up, then it is more likely to go lower before the oil is extracted than it is to go up further. This interferes with the direct use of the stock option analogy.[4] In addition, there is no market for oil futures as far into the future as a reasonable extraction plan extends.

The second problem is even more serious. The principle reason that a lease not worth drilling now will become worth drilling in the future is not changes in the price of oil. Rather it is changes in the likelihood of finding oil if one drills. Drilling, especially in deep water, is expensive, and most tracts that are drilled do not end up producing oil. Thus, a relatively small change in the likelihood that exploratory drilling will lead to economic oil production can lead to major changes in the expected value of drilling a lease. Such changes in the likelihood of drilling leading to production can occur for two reasons. One such reason is the potential for improved oil-finding and oil-producing technology. The more important reason is information from the results of drilling other leases. If a lease is near another lease, especially a lease with similar geology, that is found to have oil, the chances that it too has oil too are increased substantially.[5]

4 Inventory Model

The lease inventory model described here is focused on this second reason for leases becoming worth drilling – the chance that new information from other drilling or improved diagnostic technology will increase the probability of finding economic quantities of oil sufficient to make the lease worth drilling. It is a model of single firm. In this model, the firm buys leases which, on a fully-costed basis, have value only because of the chance that in the future they may become worth drilling (we will call these leases grade B leases).[6]

---

[3] See Anderson 1985 and Schwartz, 1997.

[4] A related problem with this theory is that if it held, then all of the oil leases would become valuable at the same time. Since there is a limited number of drilling rigs and it takes time for new ones to be built, there would be a rig shortage. Not only would this further delay the recovery of the now more valuable oil, it would also result in a wealth transfer from the oil lease holders to the rig owners.

[5] Others have noted this previously in other contexts. See for example Peterson 1975, Gilbert 1979, and Smith and Thompson 2005.

[6] Since every unleased tract in the Gulf of Mexico is offered for lease each year, it is unusual for a firm to buy a lease that is worth drilling at once on a fully- costed basis. Occasionally, something will occur between sales that

In the model, the firm has a strictly increasing cost of drilling additional leases in a given year. This cost is intended to include both direct drilling costs and decreased information values from drilling less "interesting" leases. We have modeled this cost using a quadratic function of the number of leases drilled. The model assumes that each year a random number of grade B leases become grade A leases, i.e., leases which are fully worth drilling. The firms pay attention first to grade A leases. They drill as many of these as is profitable and "farm out" the rest.[7] If after dealing with the grade A leases, the firm finds it profitable on the margin to drill some grade B leases, it does so.

In our model, leases can leave the inventory of undrilled leases in three different ways. Leases can be drilled (either as grade A leases or as grade B leases); they can be farmed out; or they can reach the upper time limit on leases and be surrendered to the MMS undrilled. We assume that which leases become Grade A leases is random with respect to the duration of the lease. Thus, the Grade A leases that are drilled on average have the average age of all leases rather than the age of leases about to be surrendered. We assume that if a grade B lease is drilled, it is a lease that is in its last year of its lease.[8]

The model assumes that the firm has a policy to keep a particular number of undrilled leases in its inventory. (The firm will optimize that number.) Thus, each year it buys new leases equal to the number of leases that have left its inventory. It also pays the annual rent on the leases in its inventory.

The model first optimizes the firms' drilling and farm-out decisions. To do this, it first decides the maximum number of grade A leases that are worth drilling. If it pays to drill all of the A leases, it then considers for each number of grade A leases drilled, the optimum number, if any, of grade B leases to drill. If in a given year it happens that there are more grade A leases than it pays to drill, it farms out to other firms the extra grade A leases.

Finally, given this optimal drilling and farming-out policy, the firm optimizes its choice of inventory level. In this optimization, the firm is assumed to be a price taker.[9] Thus, the model is static and simple enough that all of the computations can be done in a spreadsheet or by a relatively simple computer program. The specification of this model follows.

Let the decision variables be:

---

makes an unleased tract worth drilling immediately. It will generally sell in the next sale at a high price and be drilled at once. Our model ignores these unusual cases.

[7] While farming out tracts does occur, in our model farming out also serves as a highly simplified way of representing two other tactics that a firm may have available to it to cope with the opportunity presented by having unusually many tracts achieve Grade A status in a single year. The firm can also hire additional drilling capacity on a temporary basis and it can hold a grade A tract for a year and drill it the following year. We do not bother to model explicitly "farming in" by the firm or the farming out of its drilling capacity in years in which relatively few tracts become grade A tracts.

[8] If many grade B leases were drilled, this would be a biased approximation, but that is not generally the case in the examples we have worked out.

[9] However, in doing analysis on royalty relief policies, we examine parametrically, the effect of collective changes in inventory levels on the price level of tracts.

X = the inventory level,
X_t = the inventory level of "cohort" age t,
a = the number of high value (A) leases to drill,
f = the number of high value leases to "farm out," and
b = the number of low value (B) leases to drill.

The parameters are:

V_a = the value, pre-drilling, of a high value lease,
V_f = the value of "farming out" a high value lease,
V_b = the value, pre drilling, of a low value lease,
p = the annual probability that a low value lease in inventory becomes a high value lease,
y = the number of leases that become high value leases in a given year,
$F(y)$ = the probability distribution of y,
L = the lease term for leases (10 years for deepwater leases),
r = the annual rent on a lease,
C = the cost of purchasing a low value lease for the inventory, and
α and β are parameters of the quadratic drilling cost function.

The inventory variables obey the following relationships:

$$\Sigma_{t=0 \text{ to } L-1} X_t = X,$$
$$X_t = X_{t-1}(1-p) = X_0(1-p)^t, \text{ and}$$
$$X_0 = X\, p/[1-(1-p)^L].$$

With this notation, the value of a policy in a given year is given by

$$V(X,a,f,b) = (a-f)V_a + V_f f + bV_b - rX - \alpha(a+b-f) - \beta(a+b-f)^2 - C[a+b+(1-p)X_{L-1}]. \quad (2)$$

In this expression, the first term is the value of A leases drilled, the second term is the value of A leases farmed out, the third term is the value of B leases drilled, the fourth term is the annual rent for the leases in the inventory, the fifth and sixth terms give the cost of drilling, and the final term gives the cost of replacing the leases that were drilled, farmed out, or reached their lease term.

In the static model, the decision variable X is an overall one-time policy decision. However, the decision variables, a, f, and b are chosen by the company in each year with knowledge of that year's number of grade A leases. The decision variables, a and f, must be chosen so as to sum to this number, while b is constrained to be between zero and X_{L-1} minus it. In other words, the company must first decide how many of the A leases to drill and how many to farm out. Then, given the number of grade A leases to be drilled, it must decide how many of the remaining grade B leases to drill. In principle, these decisions should be made simultaneously, but as a practical matter, they can be made sequentially. First, the company decides if it pays to drill all of its grade A leases or to farm some of them out. Once it has made that decision, it can then decide if the marginal drilling costs are low enough to make it worthwhile to drill any grade B leases and, if so, how many. (This will always get the same answer as the simultaneous optimization of the decisions as long as it never pays to drill a grade B lease while farming out a grade A lease, i.e., whenever $V_a > V_b + V_f$.)

A simple marginal condition determines how many of available grade A leases to drill. If x grade A leases have been designated for drilling, then the $x+1^{st}$ grade A lease should be drilled if $\alpha + \beta(2x+1) + V_f < V_a$. Otherwise, it should be farmed out. Similarly, once it has been determined that a grade A leases and x Grade B leases will be drilled, the $x+1^{st}$ grade B lease should be drilled if $\alpha + \beta(2x+ 2a+1) < V_b$ and not otherwise.

If we use this optimization of $V(X, a, f, b)$ to fix, a, f, and b each year for the number, y, of leases that become grade A leases that particular year and call the result of that optimization $V^*(X,y)$, then the expected annual value of following inventory policy X is given by

$$E_X(V) = \Sigma_y f(y) V^*(X,y). \tag{3}$$

If leases becoming grade A leases is an independent process, $f(y)$ would be the binomial distribution on $0, 1, \ldots, X$ with mean pX. However, the independence assumption is not a necessary one, and any distribution on that range with mean pX could be used. Given the ability to calculate $E_X(V)$, it is straightforward to calculate it for a wide range of values of X to locate numerically the optimal value of X.

We initially set up the lease inventory model for a single steady state number of leases, and computed the associated value of $E_X(V)$. For example, if we assume a constant inventory of 10 deepwater leases, we can determine the expected annual value given a set of input parameters. We set the cost of acquiring a lease at $500,000, which closely resembles the observed median high bid in deepwater lease sales prior to the DWRRA. We also computed the present value of the expected economic benefit of different field sizes in the deepwater region of the Gulf of Mexico.[10] A critical assumption underlying the lease inventory model was to ensure a notable, but reasonable difference between V_a ($30,000,000) and V_b ($20,000,000), as well as keeping the value of farmouts, V_f ($1,000,000) low enough to avoid companies simply farming out all leases. We included drill cost parameters of $10,000,000 for alpha, and $1,500,000 for beta, based on our best estimate of single well drilling costs. Finally we included the typical annual rent of $43,200, a lease term of 10 years, and a 5 percent binomial probability of a low-value lease becoming a high-value lease. Using these input parameters, we can determine the optimal number of grade A leases to drill, and the expected annual value $E_X(V)$.[11]

---

[10] The field-level net present value calculations are based on earlier research by IIC, Inc., reported in the MMS study, MMS 2004-018.

[11] It is important to note that although the optimal number of grade A leases to drill is 7, there is an extremely low-probability of having 7 grade A leases at low inventory levels. Where X=10, the probability of having 6 grade A leases is 8.04×10^{-6} percent!

The use of the model is not restricted to determining $E_X(V)$ for one inventory level. Rather, the model can be used to determine the optimal inventory by maximizing $E_X(V)$ over different values of X. In determining the optimal inventory level for the aforementioned input parameters, the model iterates over a series of X values ranging from 10 to 1,000 and computes the expected annual value associated with each level. As shown in Figure 6, the highest expected annual value ($E_X(V)$ equals $47.9mm) is associated with an inventory level of 143.

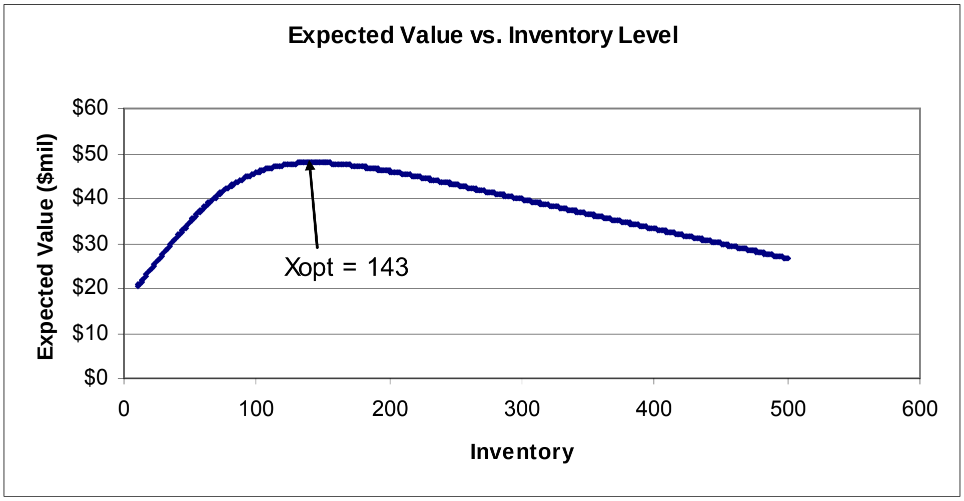

Figure 6.

As policy initiatives are implemented, there are associated changes in the input parameters, and thus, the optimal inventory level. For example, if a royalty relief program is implemented such that the values of high, low and farmout leases are increased by 20 percent, the model will generate a higher optimal inventory level, shown in Table 2, and higher expected values for all inventory levels, as depicted in Figure 7.

Table 2.

Variable	Initial	Royalty Relief
C	$500,000	$500,000
Va	$30,000,000	$36,000,000
Vb	$20,000,000	$24,000,000

Vf	$1,000,000	$1,200,000
α	$10,000,000	$10,000,000
β	$1,500,000	$1,500,000
r	$43,200	$43,200
p	5.00%	5.00%
L	10	10
Xopt	143	198
$E_{Xopt}(V)$	$47,912,401	$88,695,002

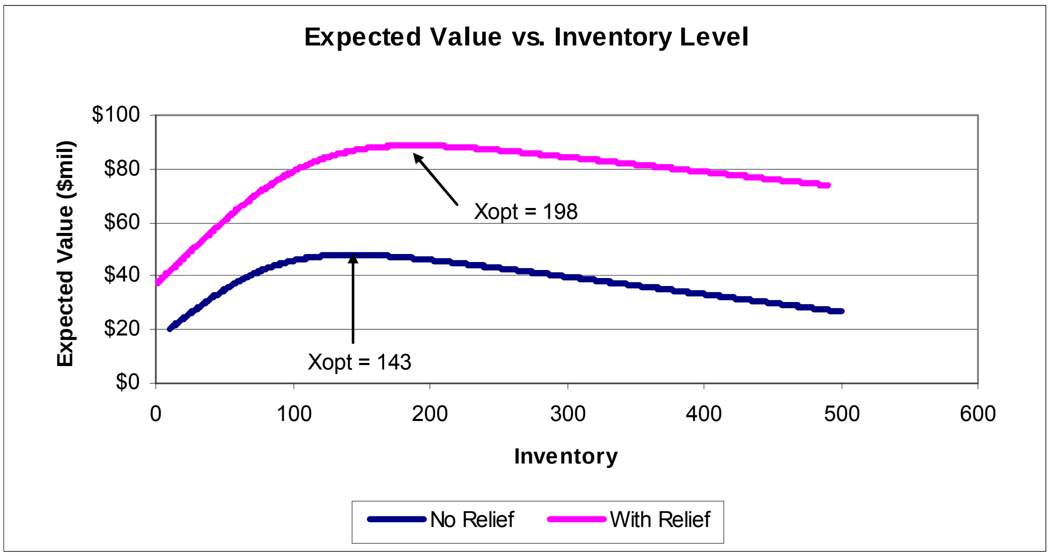

Figure 7.

Modeling royalty relief by increasing the value of the high, low and farmout leases leads to an optimal inventory that is over 33 percent larger than the original base case. In order to achieve this new optimal inventory, companies or individuals should purchase up to the new inventory level in subsequent lease sales. However, we also must consider how the other input parameters might be affected by implementation of a royalty relief program.

The results in Table 2 indicate that royalty relief leads to an increase in the value of the average lower value lease. This is determined by examining the lease cost plus the expected value at the optimal X divided by the optimal X. We observe an increase of 13 percent when including royalty relief and it is important to include our knowledge that bidders will discount the increase in value due to royalty relief consistent with the optimal bid fraction when determining lease cost. Analysis of the optimal bid fractions using the widely used methodology of Capen et al. (1971) indicates the lease cost should rise by less than a third of the observed increase in average lease value. In other words, given the low probability under the present leasing system of encountering a competitive bid, it is in the bidder's best interest to offer less than one third of the value of the lease, and we apply this observation to the increase in value of the lease due to royalty relief. Table 3 illustrates the observed differences when we increase the lease cost commensurate with the bidding model results. Interestingly, the optimal inventory does not change due to the $32,000 increase in lease cost, although the optimal economics decreases by $456,132.

Table 3

Variable	Initial	Royalty Relief	Royalty Relief (Increased Tract Cost)
C	$500,000	$500,000	$532,000
Va	$30,000,000	$36,000,000	$36,000,000
Vb	$20,000,000	$24,000,000	$24,000,000
Vf	$1,000,000	$1,200,000	$1,200,000
α	$10,000,000	$10,000,000	$10,000,000
β	$1,500,000	$1,500,000	$1,500,000
r	$43,200	$43,200	$43,200
p	5.00%	5.00%	5.00%
L	10	10	10
Xopt	143	198	198
$E_{Xopt}(V)$	$47,912,401	$88,695,002	$88,238,870

Analysis of the average high bid on deepwater leases shows that during the DWRRA time period the average bid increased 40 percent from previous years.[12] While some of this increase may possibly be due to other changes such as improved technology or a perception that the proportion of good leases was slightly higher, using this real-world lease cost difference still leads to a higher optimal inventory level than the initial case, but slightly lower than the attempts at modeling royalty relief. We performed a sensitivity analysis to observe the impact of a higher lease cost, the result of a bidder assuming greater relief than there actually is. The results are presented in Table 4 and Figure 8. It is important to remember that different bidders often have different opinions of lease value, which can lead to large variances in bidding behavior.

[12] The average high-bid for leases in water depths greater than 800 meters was $648,166 prior to the DWRRA, compared with $911,437 in the first three years after implementation.

Table 4.

Variable	Initial	Royalty Relief (Increased Tract Cost)	Royalty Relief (Higher Tract Cost)
C	\$500,000	\$532,000	\$700,000
Va	\$30,000,000	\$36,000,000	\$36,000,000
Vb	\$20,000,000	\$24,000,000	\$24,000,000
Vf	\$1,000,000	\$1,200,000	\$1,200,000
α	\$10,000,000	\$10,000,000	\$10,000,000
β	\$1,500,000	\$1,500,000	\$1,500,000
r	\$43,200	\$43,200	\$43,200
p	5.00%	5.00%	5.00%
L	10	10	10
Xopt	143	198	184
$E_{Xopt}(V)$	\$47,912,401	\$88,695,002	\$83,599,907

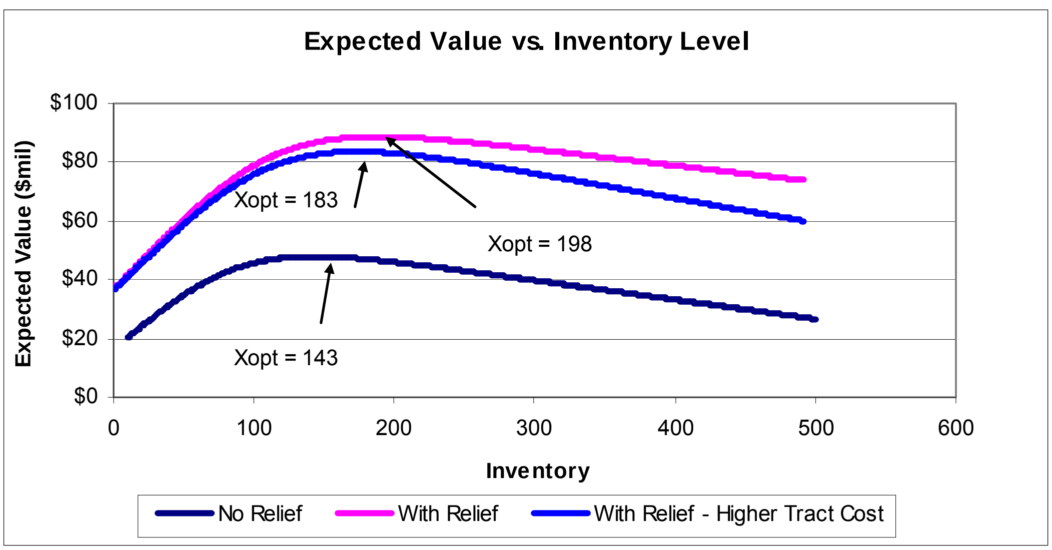

Figure 8

We expand our analysis of royalty relief one step further, analyzing lease inventory behavior from a drilling perspective. It is important to address the idea that given a finite resource base, there may be a lower probability that a low-value lease becomes a high-value lease as more leases are held in inventory. Reducing the probability from 5 to 2 percent has a profound impact on both the optimal inventory level and the expected annual value, as depicted in Figure 9. The optimal inventory rises substantially, from 198 to 355. The intuition behind this result is that it takes more leases to keep the drilling capacity employed drilling largely grade A leases. Even with the extra inventory, there are 29 percent fewer grade A leases to drill.

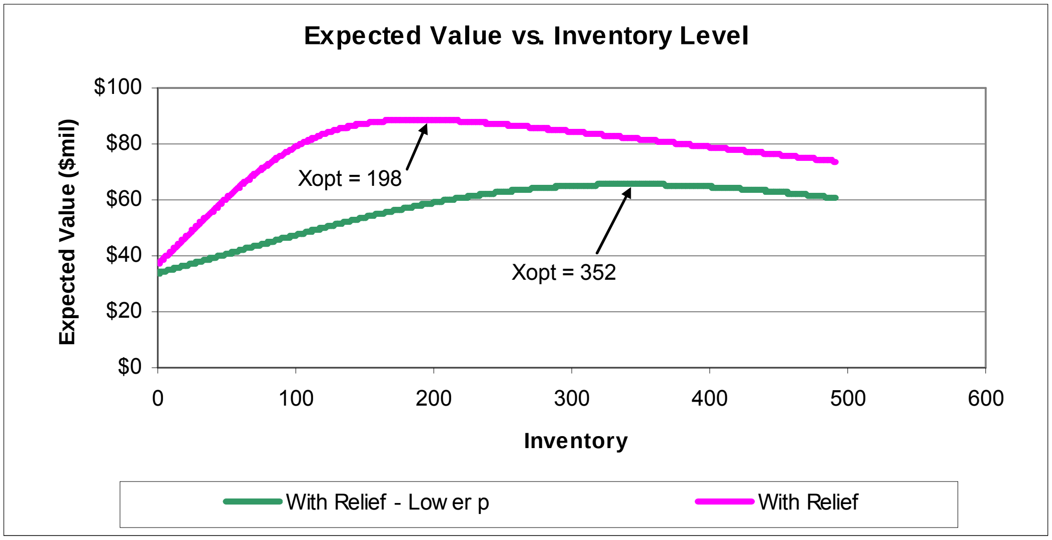

Figure 9.

As demonstrated by the above figures, the lease inventory model supports what was actually happening in the deepwater environment after the implementation of the DWRRA. As illustrated in Figure 2, the number of deepwater leases held in inventory increased dramatically during the initial years of the DWRRA before reaching a plateau in 1999. This is consistent with the lease inventory model, where a change in input parameters to model royalty relief leads to a higher optimal inventory, providing companies with a new inventory objective. Although for certain input parameters, the lease inventory model might predict an increase in inventory of only 25 percent, we must remember that different companies will employ different inventory strategies. Many bidders might not have the initial capital required to purchase a significant number of deepwater leases in any one lease sale. Alternatively, larger and more financially secure companies may benefit from economies of scale associated with drilling costs, thus reducing the drilling cost parameters. Finally, the ability to maintain an inventory of neighboring leases might provide valuable information, which might lead to a different assumption about low-value leases becoming high-value leases.

In the case of a small company unable to achieve the optimal inventory level, the lease inventory model provides information on the optimal number of low-value leases to drill given an inventory level. If we assume that a smaller company faces higher drilling costs, it is not surprising that the optimal number of low value leases to drill decreases. Given a constant inventory of 20 leases, if we have zero high-value leases in a given year, then the optimal number of low-value leases to drill is three. However, if we change the drilling parameter beta, increasing it from $1.5 million to $2.5 million, the corresponding optimal number of low-value leases to drill decreases to two. If we implement royalty relief, even with higher drilling costs, the optimal number of low-value leases to drill rises back to three, assuming zero high-value leases that year. At this stage, there is insufficient data to assess how actual drilling behavior compares to the drilling behavior illustrated by the inventory model. As leases sold in the DWRRA program age, drilling data will become more robust, and will allow examination of whether these companies are indeed drilling leases consistent with the theoretical conclusions of the lease inventory model.

5 Discussion

This model provides the first logically coherent explanation for the holding of large lease inventories that never get drilled by profit-maximizing firms. Of course, the numerical results achieved through the constant inventory level model are clearly dependent on the input assumptions for lease value, lease cost and drilling parameters, as well as the probability of low-value leases becoming high-value leases. Determining precisely what those inputs should be under different circumstances will depend upon the specifics of the situation. However, although numerical results may differ between companies or individuals, the general patterns observed in the lease inventory model are consistent with the results that we observe in the Gulf of Mexico during the years immediately after the implementation of the DWRRA. It is important to remember that the lease inventory model is a single period, constant inventory model, and does not necessarily capture the dynamic nature of inventory changes in multiple years following the start of DWRRA.[13]

Companies buy drilling rights to tracts that are never exercised. Our analysis strongly suggests that the number of oil leases sold in a given year is importantly determined by the stock of leases, including undrilled leases that the companies have at a given point in time. Optimal lease inventory management is thus an important factor in understanding why firms buy these leases at a given lease sale. We present a simple, coherent, option-value-based model in which firms hold leases on the chance that during the term of the lease the leases will become valuable enough to be worth drilling. In our model, it is generally rational to buy more leases than are likely to be drilled, because new information about the quality of particular leases is possible in the future. The number of leases to add to inventory, to hold in inventory, to drill, to farm out, and to allow expire are optimized according to this model. The model can be used to show how government polices that affect that prospective profits of extracting oil from leases affect optimal inventory management. This is a key to understanding why the demand for leases increased sharply in the first few years following a large reduction in the production royalties for deepwater leases provided by the Deepwater Royalty Relief Act of 1995.

[13] Future modifications to the inventory model may capture the dynamic situation that occurs over a series of years in response to different policy and other changes.

6 References

Anderson, R. (1985). “Some Determinants of the Volatility of Futures Prices,” *Journal of Futures Markets* **5**(3): 331-348.

Capen, E., R. Clapp, and W. Campbell (1971). “Competitive Bidding in High Risk Situations,” *J. Petroleum Technology* **23**, 641-653.

Gilbert, Richard J. (1979). “Search Strategies and Private Incentives for Resource Exploration,” *Advances in the Economics of Energy and Resources*, vol 2.

Innovation and Information Consultants, Inc. (2004). “Effects of Royalty Incentives for Gulf of Mexico Oil and Gas Leases,” final report to the Minerals Management Service on contract number 71722, November.

Peterson, Fredrick M. (1975). “Two Externalities in Petroleum Exploration,” in *Studies in Energy Tax Policy*, G. Brannon, Ed., Ballinger, Cambridge, MA.

Schwartz, E. (1997). “Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging,” *Journal of Finance* **52** (3): 922-973.

Smith, James L. and Rex Thompson (2005). Diversification and the Value of Exploration Portfolios,” working paper, Dept. of Finance, Southern Methodist University.