

A case study in the design of an optimal production sharing rule for a petroleum exploration venture*

Philip Hampson

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

John Parsons

Baruch College, C.U.N.Y., New York, NY 10010, USA

Charles Blitzer

*Massachusetts Institute of Technology, Cambridge, MA 02139, USA
World Bank, Washington, DC 20433, USA*

Received January 1989, final version received March 1990

To improve on the design of the production-sharing rule in a contract for exploration and development negotiated between a state-owned oil resources authority and a U.S. oil company, we use the Grossman and Hart (1983) principal–agent model. In the original contract, the company was granted a share of production as an incentive to maximize the net return to the authority. The optimal sharing rule we develop increases the expected return to the authority by 6% by improving the company's incentives to choose an optimal exploration program.

1. Introduction

A key aspect of every partnership is the division of profits. Sometimes the parties agree to divide the profits on a simple 50–50 or 25–75 or similar basis. Other times the shares are contingent on the size of the realized

*This research has been supported in part by the Finance, Investment, and Contracts Program of the M.I.T. Center for Energy Policy Research and is based on Philip Hampson's master's thesis, 'The Design of Optimal Production Sharing Rules for a Petroleum Exploration Venture' (M.I.T., 1988). The research assistance of Raghuram Rajan and Sung-Hwan Shin is greatly appreciated. The final draft has benefited greatly from extensive comments and suggestions by Richard Ruback, the editor, and Kevin Murphy, the referee.

profits, so that one partner receives a greater percentage for a marginal project than for a successful one. The finance literature gives little guidance on the choice of the sharing rule that governs a particular partnership. Presumably management's stake should ensure that it makes the optimal investment and operating decisions. To say how large a stake is necessary for this task and to determine how it should vary with the partnership's profits, one must carefully model the effect of the profit share on management's decisions. Models of incentive problems have been applied in finance primarily to develop general insights about conflicts of interest between liability holders in the firm, rather than to improve the design and parameters of actual financial contracts. In this case study we show how an agency model can be used to shape and calibrate the production- or profit-sharing rule in an oil exploration partnership contract.

The contract we analyze was negotiated in mid-1986 when a state-owned oil resources authority hired a U.S. oil company to explore and develop one of its territories. Many of the important decisions involved in managing the program – how many exploratory wells to drill, which finds to develop – were to be made by the company, unmonitored by the host-country authority. To give the company an incentive to pursue an exploration and development program that would maximize the value of the host country's resources, the contract specified that the company would share in the oil production from the territory.

The agency literature emphasizes that the parameters of the sharing rule should be tailored to the project at hand, since the strength of the incentives the sharing rule creates depends on assumptions about exploration costs and the probability of finding various amounts of oil. The terms of this contract, however, were not based on information about the territory under agreement. Instead, many of the terms, including the parameters for the sharing rule, were copied directly from a contract negotiated in a neighboring country. The geology of the host country, however, was strikingly different from that of its neighbor. We were therefore suspicious that the terms of the contract gave the oil company the wrong incentives and that a better set of parameters for the sharing rule could be identified, based on the actual geology of the host country.

The analysis presented here confirmed these suspicions. According to our calculations, the optimal rule would give the company a much larger share of the small discoveries than the actual contract does. Under the optimal rule, the company's share of a discovery declines significantly as the size of the discoveries increases, rather than remaining relatively constant, as it does under the actual contract. Replacing the actual contract with the optimal sharing rule could increase the total expected return on the project by \$7.8 million, or 6% of the \$129 million net present value (NPV) the authority would receive under the original contract. The increased return results from

better incentives to choose an optimal exploration program and the company's financial interest in completing marginally profitable wells.

The paper proceeds as follows. In section 2 we briefly describe the negotiated contract sharing rule, and discuss how it can be redesigned to improve the incentives. One particular suggestion for an improved sharing rule is given and some alternative explanations for the shape of the original rule are discussed briefly. In section 3 we provide a numerical analysis of the incentives in the original contract and calibrate an alternative sharing rule using a principal-agent model. We compare the estimated incentive properties of the two sharing rules. In section 4 we discuss in detail some of the problems raised by the numerical analysis, and section 5 concludes.

2. An analysis of the contract

Under the contract the oil company obtains exploration and development rights in a 10,000-square-kilometer territory owned by the authority.¹ During the first three years of the contract the company agrees to drill at least one exploratory well and to produce a minimum amount of seismic data. The exploration period can be extended twice, to a total of $7\frac{1}{2}$ years, though the company relinquishes a portion of the territory during each extension. The contract is a hybrid profit-sharing and revenue-sharing rule. All exploration and development costs incurred by the company and the few expenses incurred by the authority are recorded in an account known as a cost recovery pool, and 40% of oil production is used to cover these costs. The remaining 60% of production is divided between the authority and the company according to a step function formula. The company receives 25% of the shared petroleum when total production lies between 0 and 25,000 barrels per day (bbl/day), 22% when total production lies between 25,000 and 50,000 bbl/day, and 20% when total production exceeds 50,000 bbl/day.

To make explicit the size of the company's share as a function of the project NPV, we developed a simulation model based on certain assumptions about the cost of exploration and development, the size of the find, and the net sale price of oil. We then calculated the cash flows to the company and to the authority under various scenarios for exploration programs and discovery sizes.

Table 1 lists different discovery sizes, the present values of the associated oil revenues, and the company's gross receipts, in absolute values and as a percentage of total revenues. Table 1 also reports the company's receipts net of its expenses in developing and operating the fields, again in absolute

¹The contract studied here is a private legal instrument, and as usual in such a case, it is impossible to reveal the identity of the two parties. We gained access to the terms of the contract as a byproduct of consulting work completely unrelated to the issues raised in this paper.

Table 1

An oil company's share of total oil revenues and of development NPV under the exploration and development contract with a state-owned oil resources authority analyzed in this paper. The company's share is expressed in absolute value and as a percent of the total revenues and development NPV.

	Discovery size q_i (bbl/day)									
	0	10	50	239	1,143	5,467	26,141	125,000	597,000	
Gross project oil revenues ^a (\$ millions)	0	0	1	6	27	128	610	2,918	13,954	
Company's gross receipts under the contract ^b (\$ millions)	0	0	0	2	10	50	239	1,063	4,957	
Company's receipts as a percent of project revenues	— ^c	32.6	32.6	32.7	37.9	39.2	39.2	36.4	35.5	
Total development NPV ^d (\$ millions)	0	0	0	4	20	98	467	2,233	10,676	
Company's development NPV ^e (\$ millions)	0	0 ^f	0	0	2	13	59	231	1,051	
Company's development NPV as a percent of total	—	—	—	—	10.4	13.0	12.7	10.3	9.8	

^aThe wellhead price for oil is \$11/bbl. Developed fields are operated for 20 years, from year 6 to year 25. Cash flows are discounted at 10%.
^bFigures shown are gross receipts earned by the company. Under the contract, however, the company's gross receipts depend in part on its documented expenses for exploration and development. All calculations are for a scenario in which ten exploratory wells have been drilled, yielding the discovery specified. Each exploratory well costs \$819,760, incurred in year 1. Expenses of \$5,443/bbl/day are incurred to develop finds of each given size, in each of years 4 and 5. Operating expenses are \$0.7452/bbl.

^cNo percentage is calculated when the denominator would be zero.

^dTotal project revenues from oil sales less development and operating expenses incurred by the company. Exploration expenses are considered sunk costs.

^eThe company's gross receipts under the contract less development and operating expenses incurred. This figure is not net of exploration expenses. Under the contract, however, recorded expenses for exploration determine in part the company's share of receipts from the project's oil sales.

^fFor discovery sizes between zero and 1,000 bbl/day the NPVs of development are negative and very small. Although the company receives marginal payments in these cases, they do not cover development expenses, and therefore the company will not develop these fields.




Fig. 1. The sharing rule implicit in the actual contract. The company's percentage share of the development NPV under the oil exploration and development contract with a state-owned oil resources authority as reported in table 1 is displayed here as a function of the size of a discovery.

values and as a percentage of the total development net present value – i.e., project NPV with exploration expenses excluded as sunk costs.

In absolute terms, the company's payoff constantly increases as the discovery size increases, whereas in percentage terms the company receives the largest share of intermediate-size finds. When ten exploratory wells are drilled, the company's share of revenues increases gradually from 32.6% to 39.2% and then declines to 35.5%. Similarly, its percentage share of the development NPV initially increases from 10% to 13% for discoveries between 1,143 and 5,467 bbl/day and then declines very slightly, to just below 10%. Fig. 1 graphs the company's percentage share of the development NPV as a function of discovery size.

At first glance it may seem that the increasing absolute payments to the company are a sensible way to induce an optimal exploration program. The contract sharing rule gives the company ever-larger bonuses – payments in excess of development expenses – for larger discoveries. Two factors, however, suggest to us that this design is suboptimal.

First, as is now well known, an optimal sharing rule does not always have increasing bonuses. Information on the geology of the territory is an important factor in the design of the optimal rule. Increasing bonuses make sense if the company's choice of exploration program will affect the relative probability of larger versus smaller finds. In the territory under contract, increasing the number of exploratory wells drilled increases the probability of making a discovery, but doesn't significantly affect the relative probability of making discoveries of various sizes. Since the size of the discovery is not affected by the company's choice of exploration program, there is no reason to make the company's bonus contingent on discovery size. A size-contingent bonus only loads the company with risk without simultaneously conferring any incentive benefits. Instead, the optimal sharing rule would give the company a constant bonus for each discovery, regardless of its size. This fixed bonus would yield a constantly declining share of the NPV from development. At small discovery sizes the company would receive virtually all of the project value, and at very large discovery sizes its share would be extremely small.²

A second consideration leading us to believe that the original contract is suboptimal is the completion problem. Once the exploration costs are sunk and a find has been made, the company must decide whether to develop the field for commercial production. If the decision were based on the benefits shared by all parties, even very small wells of 59 bbl/day would be worth developing, but since the contract never gives the company more than 40% of the total revenues, and since the company must pay all of the development and production expenses, there is a range of discoveries larger than 59 bbl/day that should be developed but that the company will leave undeveloped. In an optimal contract, the company would receive the full benefit of any small discoveries, to be sure it had an incentive to develop all marginally profitable wells.³

The shape of the sharing rule suggested by these two arguments is displayed in fig. 2. It contrasts sharply with the rule displayed in fig. 1.

Before proceeding to discuss in detail the parameters of the suggested optimal design, a few caveats are in order. It is possible that the original contract was written to address other agency problems than the one mentioned here – for example, decisions about the extraction rate at any time influence not only the time profile and cost of operation, but also the total amount of oil that can be taken from a given field. In that case, the actual contract may be optimal. Alternatively, our information about the effect of

²The relationship between the assumption about the effects of the agent's action on the relative probability of different outcomes and the special shape of the optimal sharing rule was pointed out in Holmström (1979) as a rationale for deductibles in insurance contracts.

³A comparable incentive problem was documented and analyzed by Wolfson (1985) for oil and gas partnerships in the United States.




Fig. 2. A suggested design for the optimal sharing rule. This design would give the company a fixed bonus payment for each find. For uneconomical finds the company would receive no payment. For marginally profitable finds the bonus would equal the full value of the discovery. For very large finds the bonus would be a very small percentage of the value of the discovery. The design is optimal for the cost and revenue assumptions reported in table 1, the discovery size probability matrix reported in table 2, and the company utility function assumptions reported in table 5.

additional exploratory wells on the relative probabilities of discoveries of different sizes may not be accurate, in which case the structure of the optimal sharing rule might be very different and possibly close in design to the actual contract.⁴ There are also cases in which the most efficient structure of incentives for an optimal exploration program conflicts directly with the structure necessary to ensure an optimal completion decision. Then the structure of the optimal sharing rule would be highly contingent on the relative significance of the two effects.

These caveats, however, should not be construed as devaluing a careful modeling of the principal-agent problem. On the contrary, the sensitivity of the sharing rule design underscores the value of using an explicit model. Modeling exposes the assumptions made about the environment and clarifies the reasons for a given sharing rule. If the actual contract is optimal, the

⁴For example, if increased exploration increased the relative probability of larger discoveries, the optimal sharing rule could give the company a sharply increasing share of the development NPV – see Grossman and Hart (1983, propositions 7–9).

analysis highlights the need to revise either the assumptions or the model itself. Without the application of some rigorous principal–agent model it would be virtually impossible to identify precisely the important incentive effects implicit in the probability model used or to make a compelling argument for or against any particular sharing rule. These points also highlight the general conclusion that when agency problems are central, the financial contract between two parties must be carefully tailored to the cost and payoff structure of the project. In our opinion, the subject of agency theory in the field of corporate finance is precisely the exploration of the relation between the characteristics of projects and the structure of the financial contracts.

Finally, formalizing the intuition by using a detailed principal–agent model offers important advantages. For example, even if the system of increasing bonuses in the original contract makes sense as a device for giving the company strong incentives to make larger finds, the size of the bonus needs to be determined. If, for example, the company's share increased continually to 20% for the largest discoveries instead of falling to 10%, as it does under the actual contract, the company might have an even greater incentive to pursue an exploration program that raised the probability of large discoveries. It is our objective to derive the precise values for the contract sharing rule that will induce the company to pursue an optimal exploration program and thereby maximize the value of the project to the state authority.

3. A principal–agent model

To analyze the incentives embedded in the sharing rule and to more carefully derive an optimal incentive contract, we apply the Grossman–Hart (G-H) (1983) principal–agent model to our problem. The model requires both a specification of how the agent's choice of exploration program affects the probability of a discovery of any given size, and a specification of the principal's and the agent's preferences.

To formalize the probability relation between the exploration program and discovery size, we use a method that is consistent with approaches used in the exploration geology literature [e.g., Adelman et al. (1983)]. We specify twenty possible exploration programs, a_j , $j = 1, \dots, 20$, representing decisions to drill one through twenty wells, respectively. The probability of making a discovery of size q given an exploration program a_j is the product of (i) the probability of finding any oil, $G(j)$, and (ii) the probability, given a find, that the discovery is of size q , $H(q)$. A binomial distribution is used to relate the number of exploratory wells drilled to the probability of finding oil: $G(j) = 1 - (1 - \delta)^j$, where δ is called the wildcat probability. The probability that the find is a discovery of size q , $H(q)$, is modeled as a lognormal density with parameters μ and σ and is assumed independent of the number of wells

drilled. We approximate the lognormal density over discovery size with an eight-point discrete probability distribution. Each point is exactly one standard deviation apart from the next on the log scale. This process yields a 20×9 matrix that we refer to as the exploration–discovery matrix, in which the first element in each row denotes the probability that no discovery is made conditional on action a_j and the j th element of the matrix, $p_i(a_j)$, denotes the probability of a discovery of size q_i conditional on action a_j . Associated with each discovery size, q_i , we calculate a development NPV, $\pi(q_i)$, representing the value of developing the find for commercial production, exclusive of the sunk exploration costs.⁵ Table 2 shows part of an exploration–discovery matrix when the wildcat probability, δ , is 0.4, the median discovery size, e^μ , is 2,500 bbl/day, and the standard deviation of the log of discovery sizes, σ , is 20%. In fig. 3 three rows from this matrix are graphed, showing the probability distributions across discovery sizes as the number of wells drilled increases.

The exploration–discovery matrix contains all the information necessary to calculate the efficient or first-best choice of an exploration program and the expected NPV for the project. The firm simply chooses the exploration program that yields the greatest expected development NPV less exploration costs. Table 3 presents the results for a number of alternative parameters describing the probability distributions. For example, if the wildcat probability, δ , is 0.4 and the median discovery size, e^μ , is 2,500 bbl/day, then the first-best number of wells is nine; the cost of this exploration program is \$7.38 million; the expected development NPV is \$164.36 million; and the ex ante expected NPV of the exploration program is \$156.98 million. The data indicate that the optimal number of exploratory wells increases as the median discovery size increases. On the other hand, as the wildcat parameter, δ , increases, the optimal number of exploratory wells drilled falls, since the probability increases that the total benefits of exploration will be captured with the first well. The marginal value of each subsequent well falls and the expense could be applied to a different territory with greater marginal return. A rough rule of thumb seems to be that doubling the wildcat probability parameter causes optimal drilling effort to fall by a factor of one-half. These first-best choices for an exploration program are unattainable in the presence of the agency costs discussed here, and our task is to find the least-cost incentive contract with which we can approach these first-best solutions.

To complete our specification of the G-H model we need to identify a utility function for the principal, the state authority, and one for the agent, the company. In the G-H model the principal's utility function is defined over the space of possible profit levels – in our case the development NPV. The principal is typically modeled as risk-neutral, although it is possible to

⁵We set $\pi(q_i) = 0$ whenever development would be a negative NPV decision.

Table 2

Extract from a discovery size probability matrix. Given a number of exploratory wells drilled, the probability of not finding oil is calculated along with the probability of making finds of various sizes. Also calculated is the cost of each exploration program, the expected value of developing the discoveries that might be made, and the expected value net of the exploration cost.

Number of wildcat wells drilled	Discovery size q_i (bbl/day)										Expected project NPV (\$ millions)	Expected development NPV (\$ millions)	Exploration costs (\$ millions)	Expected project NPV (\$ millions)
	Development NPV π_i (\$ millions)													
	0	10	50	239	1,143	5,467	26,141	125,000	597,720	10,676				
1	0.6000 ^a	0.0005 ^b	0.0086	0.0544	0.1365	0.1365	0.0544	0.0086	0.0005	0.0005	0.8	066.4	065.6	
4	0.1296	0.0011	0.0186	0.1184	0.2971	0.2971	0.1184	0.0186	0.0011	0.0011	3.3	144.5	141.2	
8	0.0168	0.0013	0.0210	0.1337	0.3356	0.3356	0.1337	0.0210	0.0013	0.0013	6.6	163.2	156.6	
12	0.0022	0.0013	0.0214	0.1357	0.3406	0.3406	0.1357	0.0214	0.0013	0.0013	9.8	165.7	155.9	

^aThe probability of making no discoveries is equal to $1 - (1 - \delta)^j$, where δ is the wildcat probability and j is the number of exploratory wells drilled. For this table we set the wildcat probability equal to 40%.

^bFor every exploration program the probability of making a discovery of any given size is the product of two numbers: (i) the probability of making any discovery given that number of exploratory wells, δ^j , and (ii) the probability of making a discovery of size q_i given that some discovery is made, $H(q_i)$. The distribution over discovery sizes, H , is independent of the number of exploratory wells drilled. For this table we specify H as an approximation of the lognormal distribution with a median discovery size $e^{\mu} = 2,500$ bbl/day and a standard deviation $\sigma = 20\%$.

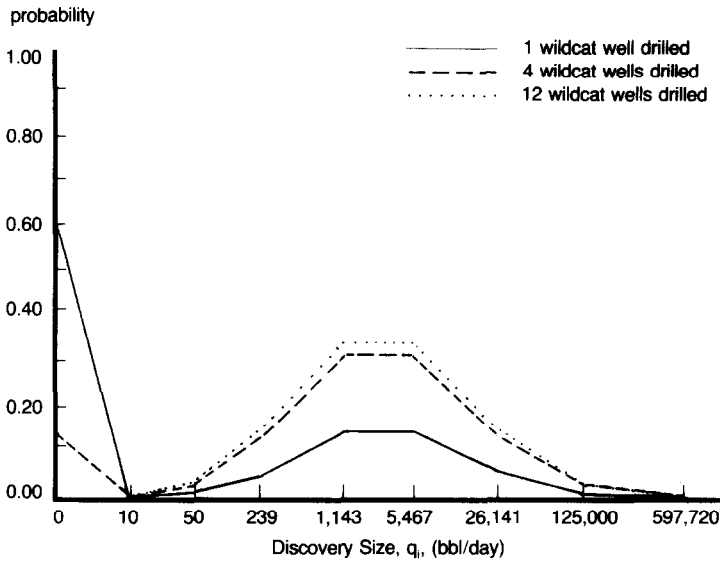


Fig. 3. Discovery size probability as a function of exploration effort. Increasing the number of wells drilled lowers the probability of not making any discoveries and raises the probabilities of making discoveries of every size. The relative probability of discoveries of any given size remains constant, independent of the number of wells drilled. The graph shows a part of the matrix reported in table 2.

incorporate risk aversion. Under the assumption of risk neutrality, the objective function for our problem is expected profit maximization. The agent is described by a utility function defined over exploration programs, a_j , and over the agent's realized compensation, I_i , $U(a_j, I_i) = V(I_i - D(a_j))$, where $D(a_j)$ is the agent's expense incurred in pursuing the exploration program a_j . The agent's compensation is a vector $I = I_1, \dots, I_9$, whose elements are the payments made to the agent in the event of each possible realized outcome, π_1, \dots, π_9 . The values for I_i assume that the authority reimburses the company for all development expenses, but not for exploration expenses. Therefore the payments I_i should be interpreted as the company's share of the development NPV.

The company's utility function must show aversion to idiosyncratic risk, $V'' < 0$. This requirement raises two problems: (i) the motivation for this unconventional assumption and (ii) how parameters for the utility function are to be set as well as how sensitive the results are to the choice of parameters. To retain the focus of this section on the mechanics of the model, we postpone until section 4 a discussion of these two issues. For now we simply note that the results reported here are based on the utility function $U(a_j, I_i) \equiv \ln(30 + I_i - D(a_j))$. We have repeated the calculations for the case

Table 3

First-best results. The optimal number of exploratory wells to be drilled, the cost of this exploration program, the expected value of developing the discoveries that might be made, and the expected value net of the exploration costs are calculated for a range of wildcat probability values and median discovery sizes.

Median discovery size e^μ (bbl/day)	First-best exploration effort (number of wells)	Exploration cost (\$ millions)	Expected discovery NPV (\$ millions)	Expected project NPV (\$ millions)
Wildcat probability, $\delta = 0.02$				
500	8	6.56	16.96	10.40
1,000	12	9.84	46.32	36.48
2,500	17	13.94	162.29	148.35
5,000	21	17.21	418.09	400.88
10,000	26	21.31	1,089.41	1,068.10
25,000	32	26.23	3,948.29	3,922.05
Wildcat probability, $\delta = 0.04$				
500	5	4.10	18.80	14.70
1,000	7	5.74	48.35	42.61
2,500	9	7.38	164.36	156.98
5,000	11	9.02	420.45	411.43
10,000	13	10.66	1,091.30	1,080.64
25,000	16	13.12	3,950.31	3,937.19
Wildcat probability, $\delta = 0.06$				
500	3	2.46	19.08	16.62
1,000	4	3.28	47.38	44.10
2,500	6	4.92	165.35	160.43
5,000	7	5.74	421.29	415.55
10,000	8	6.56	1,091.98	1,085.42
25,000	9	7.38	3,950.38	3,943.00

of exponential utility and the qualitative shape of the derived sharing rule does not change, although the precise values calculated vary slightly.

With the actual contract sharing rule, $I^\Delta = I_1^\Delta, \dots, I_9^\Delta$, the company chooses the exploration program that maximizes its utility:

$$a^\Delta \text{ maximizes } \sum_{j=1}^9 [p_j(a_j)U(a_j, I_j^\Delta)].$$

The exploration program the contract induces the company to choose and the expected returns to the authority are displayed in table 4. For the scenario in which the median discovery size, e^μ , is 2,500 bbl/day and the wildcat probability, δ , is 0.4, the contract sharing rule achieves an expected profit of \$129.05 million for the authority, or roughly 82% of the first-best expected level. Although some of the lost profit may be the result of the poor

Table 4

Expected results of the contract described in table 1. For a range of wildcat probability values and median discovery sizes and given the terms of the actual contract, we calculate the number of exploratory wells drilled that will maximize the return to the company. Also calculated is the cost of these exploration programs, the expected payment made to the company under the contract, and the net value remaining for the authority after it makes the payment.

Median discovery size e^μ (bbl/day)	Exploration effort ^a (number of wells)	Exploration cost (\$ millions)	Expected payment ^b to the company (\$ millions)	Expected return to the authority (\$ millions)
Wildcat probability, $\delta = 0.02$				
500	0	0.00	0.00	0.00
1,000	0	0.00	0.00	0.00
2,500	6	4.92	12.55	110.00
5,000	9	7.38	40.31	325.04
10,000	12	9.84	106.28	911.33
25,000	16	13.12	383.54	3,456.65
Wildcat probability, $\delta = 0.04$				
500	0	0.00	0.00	0.00
1,000	2	1.64	2.93	28.90
2,500	4	3.28	15.46	129.05
5,000	6	4.92	42.55	359.75
10,000	7	5.74	109.58	952.53
25,000	8	6.56	385.93	3,499.12
Wildcat probability, $\delta = 0.06$				
500	0	0.00	0.00	0.00
1,000	2	1.64	3.84	37.94
2,500	3	2.46	16.33	139.07
5,000	4	3.28	43.02	368.16
10,000	5	4.10	110.86	970.65
25,000	6	4.92	390.31	3,544.92

^aThe number of exploratory wells that maximizes the company's utility given the terms of the actual contract: $U(I, a) = \log(30 + I - D(a))$, where I is the payment received by the company and $D(a)$ is the cost of the chosen exploration program.

^bThe expected payment is calculated as the company's share of the project revenues less its share of the development expenses. The company's expenses for the exploration program are not included and ex ante the company must expect the payment to cover them.

contract design, some of it is an unavoidable cost of giving the company the right incentives. The only way to determine the unavoidable portion is to use our model to set parameters for the optimal sharing rule, and it is to this problem that we now turn.

The optimal or second-best sharing rule, $I = I_1^{SB}, \dots, I_9^{SB}$, is derived using a pair of mathematical programming problems. With the first program, we choose, for each number of exploratory wells drilled, (i) the least-cost incentive contract, $I^*(a_j)$, for which (ii) it is in the company's interest to accept the contract as against the best available alternative, and (iii) it is in

the company's interest to choose the specified number of exploratory wells as against any other number of wells. That is,

$$\forall a_j, \quad I^*(a_j) \equiv [I_1^*(a_j), \dots, I_9^*(a_j)] \text{ minimizes } \sum_{I_1, \dots, I_9} \sum_{i=1}^9 [p_i(a_j)I_i],$$

such that

$$\sum_{i=1}^9 [p_i(a_j)U(a_j, I_i)] \geq U_{\min},$$

and such that

$$\forall a_k \neq a_j, \quad \sum_{i=1}^9 [p_i(a_j)U(a_j, I_i)] \geq \sum_{i=1}^9 [p_i(a_k)U(a_k, I_i)].$$

For most of the results presented in the paper we have added the constraint that the payment to the company be both nonnegative and bounded above by the total development NPV, $0 \leq I_i \leq \pi_i$. With the second program, we find the number of exploratory wells that gives the authority the greatest development NPV net of the payments it must make to the company. That is,

$$a_{\text{SB}} \text{ maximizes } \sum_{a_j} \sum_{i=1}^9 [p_i(a_j)[\pi_i - I_i^*(a_j)]].$$

In table 5 the optimal sharing rule is displayed for the case in which the median discovery size is 2,500 bbl/day and the wildcat probability is 0.4 – the

Table 5

The optimal sharing rule. The payment – contingent on discovery size – to be made to the company under the calculated optimal contract is displayed in absolute value and as a percentage of the development NPV for the discovery.

	Discovery size q_i (bbl/day)								
	0	10	50	239	1,143	5,467	26,141	125,000	597,720
	Development NPV π_i (\$ millions)								
	0	0	0	4	20	98	467	2,233	10,676
Payment to the company ^a									
in \$ millions, I_i	0	0	0	4	20	30	30	30	30
as a % share, I_i/π_i	0	0	0	100	100	30.6	6.4	1.3	0.3

^aPayments are derived under the assumption that the wildcat probability of making a discovery is 40%, that the median discovery size is 2,500 bbl/day, and that the standard deviation on the log of discovery size is 20%. Also assumed is the company utility function $U(I, a) = \log(30 + I - D(a))$, where I is the payment received by the company and $D(a)$ is the cost of the chosen exploration program. The payment is net of all development expenses incurred, but not net of the expenses incurred for exploration. Ex ante the company must expect the payment to cover these expenses.

same sharing rule was graphed in fig. 2. A careful look at the payments made under this rule will help us understand the incentive problem for this project. Of course when no discovery is made, the company receives no payment. The company also receives a zero payment for discoveries of 10 or 50 bbl/day, i.e., for discoveries that should not be developed. For discovery sizes 239 and 1,143 bbl/day, the company receives a payment equal to the full development NPV: our ceiling constraint on the payment is binding for these two events. For discovery sizes ranging from 5,467 to 597,720 bbl/day the company receives payments that are approximately invariant at \$30 million. The company's share of the development NPV falls from 100% for discoveries of 239 and 1,143 bbl/day to 30.6% for a discovery of 5,467 bbl/day and to 0.3% for the largest discovery.

There is a clear rationale for this shape of the optimal sharing rule. The efficient agency contract is to pay the company a fixed bonus whenever a commercial discovery is made. Our sharing rule has this characteristic, except that at low discovery sizes the constraint that the incentive payment cannot exceed the total NPV is binding. The optimality of a fixed bonus is a consequence of two earlier assumptions. First, the authority is risk-neutral while the company is risk-averse. Therefore it is optimal for the authority to bear all of the risk except when making the company's income contingent on the development NPV provides an incentive for more exploration. Second, the probability matrix relating exploration levels to outcomes has been specified so that the probable size of any discovery is independent of the exploration effort. Therefore, the absolute size of the discovery gives the authority no information about the company's choice of an exploration program, so there are no incentive benefits to making the company's income contingent on discovery size. Instead, it is optimal to give the company a fixed bonus – approximately \$29.97 million – whenever a commercial discovery of any size is made.⁶

In table 6 we display the optimal sharing rules for alternative parameters defining the exploration–discovery matrix. The first two columns list these parameters. The number of exploratory wells to be drilled in the second-best exploration program is shown in the third column. The remaining nine columns list the payments, net of development costs, to be made to the company for each of the eight possible discovery sizes as well as the event of

⁶In addition to the argument of Holmström (1979) referred to earlier, the proof used for proposition 9 of Grossman and Hart (1983) helps explain the reason for this special shape of the sharing rule. In the proof they describe certain situations in which the expression $[1/V'(I_{i+1}) - 1/V'(I_i)]$ must be proportional to the expression $(\pi_{i+1}(a_i)/\pi_{i+1}(a_{SB})) - (\pi_i(a_j)/\pi_i(a_{SB}))$. Although all of the conditions of the proposition do not obtain in our case, over one range of outcomes – $i = 2, \dots, 8$ – the equation used in the proof is also determinant for our problem. Since over the range of $i = 2, \dots, 8$ this second expression is constant, so too must the first expression be constant, and this can be true only if $I_i = I_{i+1}$ for $i = 2, \dots, 8$, which is the case for our sharing rule.

Table 6

Optimal sharing rules with different parameters. The payment – contingent on discovery size – to be made to the company under the calculated optimal contract is displayed in absolute value for a range of wildcat probability values and median discovery sizes. Also displayed is the number of exploratory wells the company would find it optimal to drill under the contract, a_{SB} .

Median discovery size e^μ (bbl/day)	Wildcat probability δ (%)	Chosen number of exploratory wells a_{SB}	Index of the discovery size i								
			1	2	3	4	5	6	7	8	9
500	0.2	4	0	0	0	0.44	4.80	13.56	13.54	13.55	15.17
500	0.4	3	0	0	0	0.44	4.80	7.87	7.86	7.85	8.28
500	0.6	2	0	0	0	0.44	3.93	3.93	3.93	3.95	3.80
1,000	0.2	6	0	0	0	1.57	8.95	18.57	18.55	18.54	20.73
1,000	0.4	4	0	0	0	1.57	8.95	11.40	11.40	11.42	10.88
1,000	0.6	2	0	0	0	1.57	3.73	3.73	3.73	3.75	3.60
2,500	0.2	9	0	0	0	4.21	20.42	29.03	29.08	29.13	33.62
2,500	0.4	6	0	0	0	4.21	20.42	29.99	29.98	29.71	29.56
2,500	0.6	3	0	0	0	4.21	8.65	8.65	8.66	8.60	9.65
5,000	0.2	12	0	0	0.28	6.94	38.11	51.43	51.36	51.67	48.03
5,000	0.4	6	0	0	0.28	6.94	25.02	25.02	25.00	24.77	24.87
5,000	0.6	5	0	0	0.28	6.94	20.16	20.15	20.17	20.27	18.90
10,000	0.2	14	0	0	0.96	11.27	65.93	65.92	66.09	67.03	51.09
10,000	0.4	8	0	0	0.96	11.27	61.87	61.99	62.02	62.70	173.52 ^a
10,000	0.6	4	0	0	0.96	11.27	19.33	19.33	19.33	19.29	19.30
25,000	0.2	17	0	0	2.32	21.40	111.52	116.78	122.24	104.44	5465.30
25,000	0.4	8	0	0	2.32	21.40	57.31	57.33	57.52	57.33	167.36
25,000	0.6	5	0	0	2.32	21.40	41.87	41.87	41.85	41.94	41.65

^aValues for the largest payment to the company are sometimes extreme – for example, when $e^\mu = 10,000$ and $\delta = 0.4$ or when $e^\mu = 25,000$ and $\delta = 0.2$. This occurs because (i) the probability of this event is very small, and (ii) rounding is done in the process of constructing the discrete probability matrix. The construction of the optimal sharing rule plays heavily on the small changes in relative probabilities across events. Since the probability of these events is so small, however, the large numbers never significantly change the total expected payment to the company.

Table 7

Second-best results. For a range of wildcat probability values and median discovery sizes and given the terms of the calculated optimal contract, we calculate the number of exploratory wells drilled that will maximize the return to the company. Also calculated is the cost of these exploration programs, the expected payment made to the company under the contract, and the net value remaining for the authority after it makes the payment.

Median discovery size e^μ (bbl/day)	Second-best exploration effort (number of wells)	Exploration cost (\$ millions)	Expected payment to the company (\$ millions)	Expected return to the authority (\$ millions)
Wildcat probability, $\delta = 0.02$				
500	4	3.28	5.01	7.03
1,000	6	4.92	9.26	27.44
2,500	9	7.38	19.12	124.69
5,000	12	9.84	36.93	356.03
10,000	14	11.48	54.52	989.81
25,000	17	13.94	98.23	3,757.80
Wildcat probability, $\delta = 0.04$				
500	3	2.46	4.42	11.57
1,000	4	3.28	7.81	35.49
2,500	6	4.92	21.48	136.81
5,000	6	4.92	20.96	381.34
10,000	8	6.56	52.83	1,021.60
25,000	8	6.56	57.32	3,835.80
Wildcat probability, $\delta = 0.06$				
500	2	1.64	2.83	14.29
1,000	2	1.64	2.82	38.96
2,500	3	2.46	7.35	148.09
5,000	4	3.28	17.45	393.83
10,000	4	3.28	17.40	1,047.70
25,000	5	4.10	37.75	3,874.90

no discovery. Since the absolute sizes of the eight discoveries are different for each set of probability parameters, the columns are indexed only by the ordinal size of the discovery.

Table 7 summarizes the project results under the second-best sharing rule for the same set of alternative parameters. For the scenario in which the median discovery size, e^μ , is 2,500 bbl/day and the wildcat probability, δ , is 0.4, the optimal sharing rule yields an expected profit of \$136.81 million for the authority, or roughly 87% of the first-best expected level. The 13% difference is the cost the principal must pay to give the agent the appropriate incentives. Incentive costs – measured as a fraction of the first-best profit – decrease as median discovery size and the wildcat probability increase. The incentive cost is 32% of first-best expected profits for a median discovery size of 500 bbl/day with a wildcat probability of 0.2 and decreases to 4% of first-best profits for a median discovery size of 25,000 bbl/day.

Table 8
A comparison of the results of the optimal and actual contracts.^a

	Number of wells drilled	Cost of the exploration program chosen (\$ millions)	Expected payment to the company (\$ millions)	Expected return to the authority (\$ millions)
First-best	9	7.38	7.38	156.98
Actual contract	4	3.28	15.46	129.05
Optimal contract/ second-best	6	4.92	21.48	136.81

^aFigures are taken from tables 4, 5, and 7 for the case in which the wildcat probability is 40% and the median discovery size is 2,500 bbl/day.

Measured in absolute value, however, the incentive cost increases from \$3.37 million to \$164.25 million. As we change the wildcat probability from 0.2 to 0.6, keeping the median discovery size constant at 500 bbl/day, the incentive cost falls from 32% of first-best expected profits to 15% – from \$3.37 million to \$2.33 million.

It is now possible to compare the results of the optimal contract design with the results of the actual contract sharing rule – see tables 4 and 7. One can see that the ex ante expected profits to the authority are significantly lower under the actual contract. For example, for the scenario in which the median discovery size, e^{μ} , is 500 bbl/day, regardless of the wildcat probability the project is worthless under the actual contract, whereas under the optimal sharing rule the authority realizes between \$7 and \$14 million profit. For the scenario in which the median discovery size is 2,500 bbl/day and the wildcat probability is 0.4, the actual contract gives the authority \$7.76 million less in expected profit than the optimal sharing rule, a loss of 5.7% of the total project value. For larger values of the median discovery size and the wildcat parameters, the dollar loss with the actual contract increases significantly – up to \$336 million. A summary comparison of the results for the first-best rule, the actual contract, and the optimal contract is given in table 8.

The actual contract is less favorable to the authority than the optimal sharing rule primarily because incentives in the actual contract are expensive. In the actual contract the size of the payment to the company is highly correlated with the size of the discovery. But given that a find is made, the absolute size of the discovery is an exogenous random variable unrelated to the company's choice of an exploration program, so the randomness in the payment does not increase the company's incentive to expand its exploration program. Moreover, the average payment to the company must be increased to compensate for the risk the company bears under the actual contract. In

the optimal sharing rule the company receives nothing when no discovery is made, and a fixed bonus when a discovery of any size is made. Since a discovery conveys some information about the extent of the exploration program, the bonus increases the company's incentive to expand the exploration program, but since the bonus doesn't vary with the size of the discovery, the risk borne by the company is minimal.

4. Discussion

The explicit use of a utility function for the agent/company in this paper raises an interesting problem in the application of some agency models to corporate finance. It is often assumed that managers of a firm choose projects as if the firm were neutral toward idiosyncratic risk. This assumption corresponds correctly with the decision rule that results if shareholders have diversified their investments and capital markets are perfect. If the management of the corporation, or the agent, is neutral toward the project's idiosyncratic risk, then, as Harris and Raviv (1979) show, the efficient incentive contract takes on a trivial form. The principal, in this case the authority, should simply sell the rights to exploration and development to the agent, the company, for a fixed fee and allow the company to bear all of the risk. The actual contract in the case under study does not fit this description. Thousands of comparable contracts are negotiated with similarly complex designs that do not correspond to the predictions of models based on the agent's risk neutrality. We believe that in many cases these modeling assumptions and the incentive contract they predict do not correspond very closely to the real context.

Our assumption that the company is to some degree averse to the idiosyncratic risks of the project can be justified in the framework of an appropriate equilibrium model of a capital market with imperfect information. The assumption of risk neutrality is a strong one that makes sense only in a model of relatively perfect information and frictionless capital markets. If, however, the management of each firm possesses private information about the projects in which it is investing and about its own management of these projects, we believe the consequent imperfections in the financial markets lead the firm rationally to behave as if it were averse to the idiosyncratic risk associated with each particular project. Holmström (1982) makes a similar suggestion. The environment we have in mind could even incorporate individual investors who are risk-neutral and who hold diversified portfolios. Rather than focus on modeling the environment that might generate this risk aversion, we choose in this paper to take it for granted and focus instead on the interesting problem of how to design and calibrate an optimal sharing rule between the authority and the company in this environment.

Of course, having decided to recognize the risk aversion that probably characterizes most corporate activities, we then face the problem of specifying the utility function. This problem has not been seriously explored in the finance literature, so there is dangerously little basis on which to specify the model. The results presented here assume that the agent's utility function is $\log(30 + I - D(a))$: we began with this specification because it describes a project manager with a limited exploration budget of \$30 million, a budget more than adequate to cover any of the possible first-best exploration programs, and in some cases much more. Fortunately in our case the key result of the analysis is unaffected by the parameters of the utility function: although the exact values of the payments made to the agent are obviously subject to the choice of utility function, the shape of the optimal sharing rule does not change. The actual contract cannot be made optimal by altering the parameters of the agent's utility function. It is possible to prove that under all parameter choices the optimal sharing rule involves a zero payment for small finds and a fixed bonus for large finds, with the size of the bonus unaffected by the size of the find. The motivation for this shape was provided in section 2 above. Of course, in the application of this model to other cases, the parameters of the agent's utility function may be significant in determining the optimal shape of the sharing rule. This issue therefore requires attention in the future.

We have modeled the problem as if the principal were risk-neutral. This assumption is made in the theoretical literature primarily because risk aversion is not necessary for interesting results, and the problem and its solution are much clearer if the principal is risk-neutral. The Grossman–Hart model can easily be adapted for a risk-averse principal. Many arguments can be made for modeling the principal in any given case as risk-averse or risk-neutral, but as Holmström (1982) points out, the principal's aversion to the idiosyncratic risk of the project does not immediately follow from the capital-market imperfections that make the agent risk-averse.

Instead of debating whether the principal should be modeled as averse to the idiosyncratic risk, we merely note that a change in this assumption has important consequences for the design of the optimal sharing rule and the comparison of the actual contract with it. Table 9 shows the optimal sharing rule when the authority shows the same risk aversion as the company. The payments to the company increase very moderately with discovery size, so the company shares a portion of this exogenous risk. The company's share of the development NPV declines significantly as the discovery size increases, moving from 100% to 0.2%. The structure of the sharing rule is therefore very similar to the one displayed in table 5, the case of the risk-neutral principal. The authority's risk aversion leads to a sharing rule in which the authority also receives a portion of the value of small discoveries: the company's share of a discovery of 1,143 bbl/day is only 38.2% instead of

Table 9

The optimal sharing rule with a risk-averse principal. The payment – contingent on discovery size – to be made to the company under the calculated optimal contract under the assumption that the authority is also risk-averse is displayed in absolute value and as a percentage of the NPV of developing the field. The contract simultaneously provides incentives to the company and optimally shares the risk between the company and the authority.

	Discovery size q_i (bbl/day)								
	0	10	50	239	1,143	5,467	26,141	125,000	597,720
	Development NPV π_i (\$ millions)								
	0	0	0	4	20	98	467	2,233	10,676
Payment to the company ^a									
in \$ millions, I_i	0	0	0	4	8	23	44	37	27
as a % share, I_i/π_i	0	0	0	100	38.2	23.2	9.4	1.7	0.2

^aPayments are derived under the assumption that the wildcat probability of making a discovery is 40%, that the median discovery size is 2,500 bbl/day, and that the standard deviation on the log of discovery size is 20%. Also assumed is the company utility function $U(I, a) = \log(30 + I - D(a))$, where I is the payment received by the company and $D(a)$ is the cost of the chosen exploration program. The payment is net of all development expenses incurred, but not net of the expenses incurred for exploration. Ex ante the company must expect the payment to cover these expenses. The authority is assumed to be risk-averse with a utility function $\log(30 + \pi - I)$.

100% and its share of a discovery of 5,467 bbl/day is 23.2% instead of 30.6%. Of course, the company must receive a slightly increased share in other events to compensate for these reductions. Although the exact values for the payment to the company do change, the general structure of the optimal sharing rule remains as we have described it. In particular, under all cases (i) it is important that the company receive a large fraction of the development NPV when the discovery is small, and (ii) although the absolute dollar value of the payment to the company may be increasing, it is always a sharply decreasing share of the development NPV. This sharing rule remains in sharp contrast to the one implicit in the actual contract, and the expected returns to the authority remain far greater.

Another important assumption implicit in our use of the Grossman–Hart model is that the company makes an essentially one-shot decision. In one respect this assumption is reasonable, since most of the uncertainty in an oil exploration venture is resolved in a relatively short period. At the end of that time a decision is made whether to develop the territory for commercial production. After this point the operating wells can in some cases be properly viewed as cash cows. In some important respects, though, the exploration decision is not a one-shot problem. For example, one could view the exploration program as involving a Bayesian updating process in which the results from the first wells drilled provide information about whether

additional wells should be drilled or whether the program should be abandoned. The number of wells drilled is then the result of a dynamic stochastic program and not a simple *ex ante* choice variable, as in our model. Simplifying the decision problem as we have done obviously sacrifices some accuracy and some interesting components of the problem, but we think the value of our results justifies this compromise. An interesting research problem would be to analyze the company's choice of exploration program as a dynamic agency problem and to consider the actual contract as an attempt to provide the optimal incentives for this type of problem.

As a final caveat, we treat the oil price as if it were known *ex ante* and constant over the life of the contract. To allow for an uncertain future oil price, we must modify several aspects of our analysis. First, we must be more careful about the exact terms in which the contract is written. The actual contract specifies shares in quantity of oil that each party will receive: hence, when the oil price changes, the absolute value of the payments to the company will change even though its share of the oil produced does not. On the other hand, we have solved for the optimal sharing rule in terms of the absolute value of the dollar payments that should be made to the company in the event of different discovery sizes. When the oil price changes, these payments should remain fixed; they will represent changing fractions of the development NPV. Given that oil prices are variable, in which terms should the optimal contract be written? The company's choice of an exploration program affects the probability of a discovery, and not the present value of that discovery. Therefore the company should receive an incentive payment whenever it makes a discovery, and the payment should not depend on the realized value of the discovery and thereby on the price of oil. This conclusion is modified slightly when the principal is risk-averse, since the risk implicit in the variable price then needs to be shared between the two parties – although not for incentive purposes. Solving for the optimal sharing rule in the face of an uncertain oil price gives results virtually identical to those presented for a fixed oil price. There may be other reasons to write the contract in terms of shares of oil, but from the standpoint of incentives for exploration in a model such as ours this is not the preferred form.

5. Conclusion

We analyze a classic capital budgeting case in the form of an oil exploration project, and we value in the classic manner the contract written between the owner of the territory and the manager of the exploration and development program. We then analyze the same contract using a principal–agent model, quantifying the incentives given to the manager of the exploration and development program, and valuing the project again in light of these incentives. We solve for an optimal incentive contract and contrast it

with the actual one. Our results demonstrate that it is possible to improve significantly on the design of the original contract, using an existing principal-agent model. In the course of this demonstration we highlight several difficulties in the practical application of the model.

References

- Adelman, Morris A., 1986, The competitive floor to world oil prices, *Energy Journal* 7, 9–35.
- Adelman, Morris A., John C. Houghton, Gordon Kaufman, and Martin Zimmerman, 1983, *Energy resources in an uncertain future* (Ballinger, Cambridge, MA).
- Grossman, Sanford J. and Oliver D. Hart, 1983, An analysis of the principal-agent problem, *Econometrica* 51, 7–45.
- Harris, Milton and Artur Raviv, 1979, Optimal incentive contracts with imperfect information, *Journal of Economic Theory* 20, 231–259.
- Holmström, Bengt, 1979, Moral hazard and observability, *Bell Journal of Economics* 10, 74–91.
- Holmström, Bengt, 1982, Moral hazard in teams, *Bell Journal of Economics* 13, 324–340.
- Wolfson, Mark A., 1985, Empirical evidence of incentive problems and their mitigation in oil and gas tax shelter programs, in: J.W. Pratt and R.J. Zeckhauser, eds., *Principals and agents: The structure of business* (Harvard Business School Press, Boston, MA) 101–125.