OPTIONS IN THE REAL WORLD: LESSONS LEARNED IN EVALUATING OIL AND GAS INVESTMENTS

JAMES E. SMITH and KEVIN F. MCCARDLE

Duke University, Durham, North Carolina

(Received August 1996; revisions received October 1996, April 1997, March 1998; accepted April 1998)

Many firms in the oil and gas business have long used decision analysis techniques to evaluate exploration and development opportunities and have looked at recent development in option pricing theory as potentially offering improvements over the decision analysis approach. Unfortunately, it is difficult to discern the benefits of the options approach from the literature on the topic: Most of the published examples greatly oversimplify the kinds of projects encountered in practice, and comparisons are typically made to traditional discounted cash flow analysis, which, unlike the option pricing and decision analytic approaches, does not explicitly consider the uncertainty in project cash flows. In this paper, we provide a tutorial introduction to option pricing methods, focusing on how they relate to and can be integrated with decision analysis methods, and describe some lessons learned in using these methods to evaluate some real oil and gas investments.

Uncertainty and complexity are business as usual in the upstream oil and gas business. Consequently, firms in this industry have long used quantitative tools for decision making. Many firms in this industry have made extensive use of decision analysis methods, and many have looked with interest at recent developments in option pricing theory. In this paper we provide a tutorial introduction to option pricing methods, describing their relationship with decision analysis techniques. How can the options approach be implemented in practice? What are the benefits of the options approach as compared to decision analysis methods? We illustrate the options approach by applying it on some real oil and gas investments.

This study was undertaken by the authors in conjunction with a major oil and gas company. Currently, almost every major capital expenditure at the firm is evaluated using decision analysis techniques. In a typical evaluation, the firm's analysts use sensitivity analysis to identify the key uncertainties, assess probabilities for these key uncertainties, and construct decision tree or simulation models of the project cash flows. They use these models to calculate expected net present values (NPVs) and distributions on NPVs, which then serve as important inputs into the decision-making process.

Though management is generally pleased with their decision analysis process, there were two concerns. First, there was concern that their analyses frequently do not capture some of the flexibilities associated with projects. Their decision models typically assume that management makes an initial investment decision, and then the project uncertainties are resolved and cash flows are determined. In reality, the firm makes a series of investment decisions as the uncertainties resolve over time. For example, when considering the development of a new oil field, if oil prices, production rates, or reserves exceed their expectations, or if production technology improves, the firm might be able to develop more aggressively or expand to nearby fields. Similarly, if prices, rates, or reserves are below expectations, the firm might be able to scale back planned investments and limit their downside exposure.

A second issue that has long concerned many at the firm is the way they discount cash flows. Many of their investments have time horizons as long as 30 or 40 years, and the NPVs for these investments are extremely sensitive to the discount rate used. Currently, the company calculates NPVs for these projects using a discount rate that reflects their cost of capital and desired rate of return. This discount rate is well above the rate for risk-free borrowing and lending (currently in the 6 to 7 percent range) and hence can be viewed as a "risk-adjusted" discount rate. There is concern, particularly among managers in the exploration and new ventures parts of the business, that the blanket use of such a risk-adjusted discount rate causes them to undervalue projects with long time horizons.

With these concerns in mind, many at this company have watched recent developments in option pricing theory with great interest (see, e.g., Dixit and Pindyck 1994, Trigeorgis 1996). In this approach, one views projects as analogous to put or call options on a stock and values them using techniques like those developed by Black and Scholes (1973) and Merton (1973) to value put and call options on stocks. These methods explicitly model and

Subject classifications: Decision analysis, applications: evaluating oil and gas investments. Finance: investments, real options.
Area of review: NUMBER GENERATORS.

0030-364X/99/4701-0001 $05.00

value the decision maker's ability to make decisions (e.g., "exercise the option") after some uncertainties are resolved and do not require the use of a risk-adjusted discount rate. Thus, these new techniques appear to have the potential to address both of management's concerns with their decision analysis process.

The analysts at this company were concerned, however, that the models described in the real options literature greatly oversimplify the problems they actually face.[1] For example, an undeveloped oil property is superficially analogous to a call option on a stock, but in reality there are many complications (uncertain production rates, development costs, construction lags, complex royalty and tax structures, the lack of a true underlying stock, etc.) that strain the analogy. Moreover, most of the articles describing the benefits of the options approach compare it to a traditional discounted cash flow approach based on point estimates of all cash flows. It was not clear what advantages the options approach would have compared to their decision analysis approach.

To better understand the potential of the options approach, the company formed an interdisciplinary team to see how option pricing methods compare with and could be integrated with their current decision analysis approach. This "Valuation Methods Improvement" (VMI) team consisted of six analysts from a variety of different operating companies within the firm, as well as the two authors. In addition, a Steering Committee, consisting of executives from corporate staff and several operating companies, was formed to oversee the VMI team. To facilitate comparisons between approaches, the option pricing methods were to be applied to projects for which the firm had done extensive decision analyses.

In this paper we describe some of the lessons learned by this VMI team with the goal of providing a tutorial on option pricing techniques and describing how they relate to, and can be integrated with, decision analytic methods. While our paper focuses primarily on the concerns and questions of a particular oil and gas company, we have heard similar concerns and questions from other oil and gas companies as well as from firms in a variety of other industries, especially electric utilities and pharmaceutical firms.

1. MODELING FLEXIBILITY

The first lesson we learned in this effort is that there are two distinct sets of issues associated with applying option methods. The first set of issues has to do with modeling project flexibilities: What options does management have now? What options will they have in the future? How should these options be modeled? The second set of issues concerns the valuation procedure used. We can contrast two different valuation approaches: the conventional risk-adjusted discount rate approach that this company and many others currently use, and the new valuation procedure underlying the option pricing approach. The two sets of issues really are distinct: One could do a great job modeling flexibilities and then value the risky cash flows using either the risk adjusted discount rate approach or the option valuation approach. Similarly, one could model no project flexibilities and use either valuation approach.

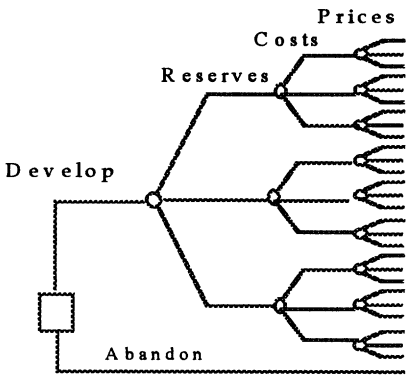

Figure 1. The original tree for Project X.

We begin by considering issues associated with modeling flexibility in this section and then consider the valuation issues in the next section. In this section, we will value cash flows using the conventional, risk-adjusted discount rate approach. In the next section, we consider the rationale of the risk-adjusted discount rate approach and compare and contrast this approach to the option valuation procedure.

1.1. Problem Structuring

In discussing issues associated with modeling flexibility, we will focus primarily on one of several projects considered by the VMI team. This project—which we will call Project X—is a large, undeveloped, offshore oil field. There had been a significant amount of exploratory drilling in this field and substantial reserves had been identified, though there was still substantial uncertainty about the extent of the field and the total reserves.

The original decision analysis study for Project X was based on the decision tree in Figure 1. The only decision considered in this analysis was whether to proceed with the project. Three uncertainties were modeled: reserves, prices, and costs. Each price scenario represents a sequence of oil prices, one for each year, going out for approximately 30 years. Similarly, the cost and reserve uncertainties represent a sequence of costs and production rates (and associated drilling expenditures) for each year, going out about 30 years. The distribution for reserves was calculated from a complex model that considered uncertainty about reserves in the as yet unexplored areas, the uncertainty in production rates, as well as many other factors. The values at the end of the tree represent NPVs of cash flows determined using an economic model that includes complex tax and royalty calculations. The results of this analysis showed a project with a positive expected NPV but a significant chance of having a negative NPV. In the end, the project was viewed as marginal because its

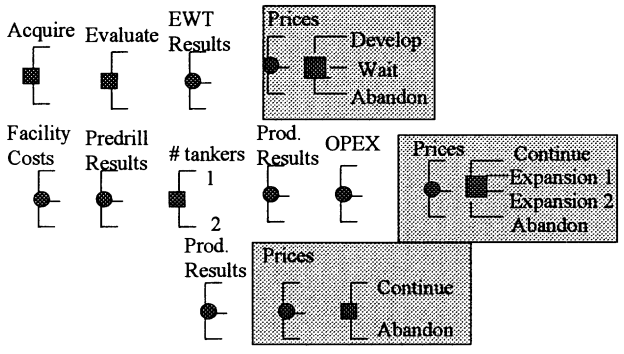

Figure 2. "Dream tree" for Project X.

expected NPV was small compared to the amount of capital required.

After reviewing a number of different projects and decision analyses, Project X was selected as a candidate for our study because the VMI team and steering committee felt there were significant project flexibilities that were not captured in the original analysis. Because the project was marginal and its value was very sensitive to oil prices, it was viewed as being like a call option: Though the project was marginal at current prices, it could have considerable value if prices were to rise at some point in the future. There were also some expansion options in that one could use the platform and facilities constructed for Project X to develop (or "tie in") other nearby fields when production at the main field declined. Finally, there were abandonment options in that the field could be abandoned at any time if continued production appeared uneconomic. While the decision tree of Figure 1 contained fewer uncertainties than most of the models we reviewed, it was typical in its lack of delayed (or "downstream") decisions: Most of the models we reviewed had many uncertainties but used heuristic policies to determine when to expand production or shutdown the field rather than optimal policies depending on the then-prevailing prices, costs, etc. None of the models we reviewed explicitly modeled the decision concerning when to begin development.

The first step in our new analysis of Project X was to construct a new decision tree for the project that incorporates these previously unmodeled options (see Figure 2). The first row of this tree represents the predevelopment phase of the project. The firm's first decision is whether to acquire rights to develop the field. Next, they decide how and whether to test the field to learn more about production before drilling; for example, should they do extended well tests? They then observe the results of these tests. The boxes in the tree indicate repeated elements (or "do loops") in the tree. For example, if the firm decides not to begin development and not to abandon the field (by surrendering their development rights), they wait and observe oil prices and face the same decision in the next time period, say next year. If they choose to wait again in the next year, they observe prices again and face the same decision in the subsequent year. Once they choose to develop the field and construct the necessary facilities, they break out of this loop and move to the initial development phase and produce from the primary field. Here they observe production and costs as well as prices and enter into a new loop and repeatedly decide whether to continue production, abandon the field, or tie in the nearby fields. The final row of the figure represents the second development phase where they produce at these nearby fields. Depending on how prices and production rates evolve over time, they eventually abandon the field and salvage the offshore production facility.

1.2. Modeling Frameworks

In comparing Figures 1 and 2, we see that trees that take into account the flexibilities quickly become huge. In order to model these downstream decisions correctly, one must include not only the decisions in the tree, but also the information available at the time these decisions are made. While the tree of Figure 1 requires the evaluation of the economic model in a total of 27 ($= 3 \times 3 \times 3$) different scenarios, even if we allow only a few iterations in the "do loops," the tree of Figure 2 is much too large to be evaluated using off-the-shelf decision analysis software and today's personal computers. For these reasons, we referred to this tree (and others like it) as a *dream tree*—this is the tree we wished we could solve.

We considered three different approaches to evaluating this dream tree. The first approach we considered is to reduce the number of uncertainties and decisions modeled so the tree can be evaluated using commercially available decision analysis software. We did this in Project X by building a decision tree with five-year time increments and, reflecting our initial view of the project as a call option on oil prices, we focused our analysis on price uncertainties and development decisions. This simplified tree (shown in Figure 3) had approximately 52,500 endpoints and a complex spreadsheet-based economic model. It took about 90 minutes to run this model using commercially available decision analysis software (DPL™) on a 166-MHz Pentium-based personal computer.

An alternative approach to evaluating these flexible decision models is to use dynamic programming techniques. For example in Project X, to get a better understanding of the initial development decision, we constructed an infinite-horizon dynamic programming model corresponding to the first "do loop" of Figure 2. In this model, oil prices were uncertain and evolving over time according to a stationary Markov process (details in section 2.4 below). In this model, the costs of waiting and abandonment were directly specified. The possible values of the field at the time of development were calculated by repeatedly solving a tree similar to that of Figure 3 but assuming immediate development with varying initial price assumptions. This infinite-horizon dynamic program framework can easily handle small time steps (in this case we worked with time steps that were approximately two weeks in length), but in order to satisfy the additivity assumptions required by the

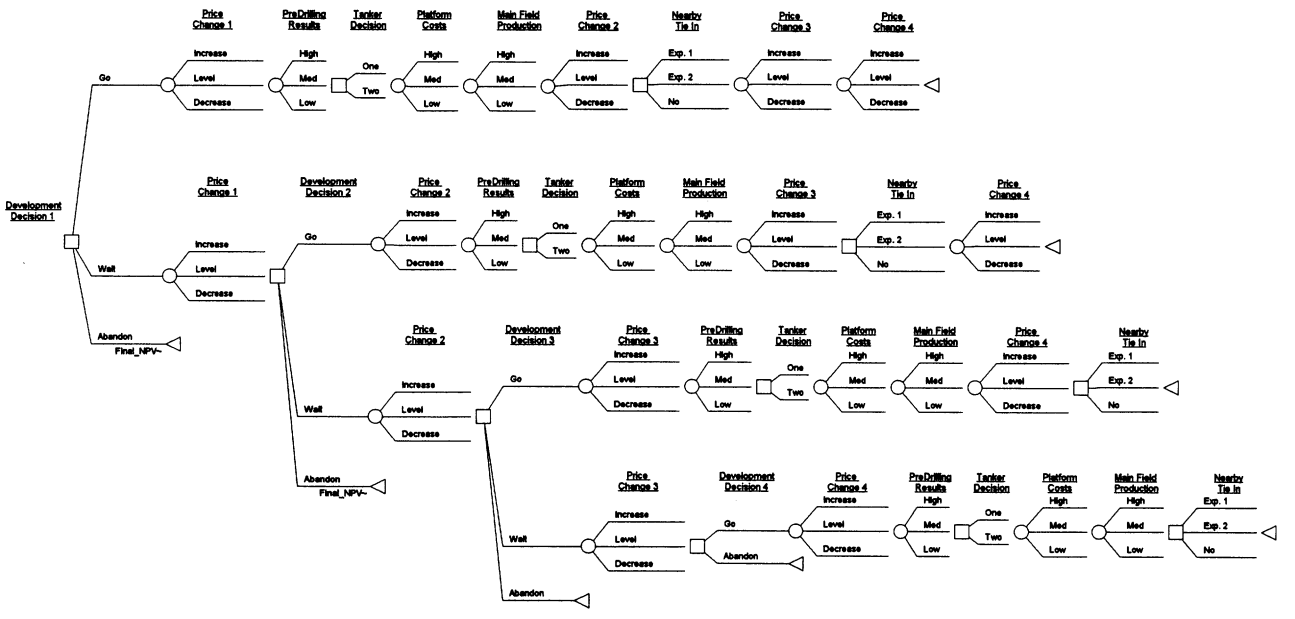

Figure 3. Simplified tree for Project X.

dynamic program, we had to make some simplifying approximations in the project's royalty and tax calculations.

The key to developing dynamic programming models of these kinds of problems is to identify a reasonably small set of state and decision variables that are sufficient to describe the value of the project over time. In the dynamic programming model of the initial development decision for Project X, we tracked only the evolution of prices over time. For an early phase exploration project, we developed a more complex dynamic programming model that considered prices, productivity, leasing costs, and drilling costs as uncertain and evolving over time and modeled drilling and leasing decisions on a well-by-well basis. The major barrier to implementing these dynamic programming models was the amount and level of computer programming required. While spreadsheets and off-the-shelf decision analysis software make it relatively easy to evaluate the decision tree models, we had to develop fairly sophisticated custom code to formulate and solve these dynamic programming models. For example, the dynamic programming model of the exploration problem was formulated as a large linear program and solved on a UNIX workstation using commercial LP software. The most time-consuming aspects of this effort were converting the model specification (especially the spreadsheet-based economic model) into the format required for the LP model and then converting the LP results to forms suitable for discussion with management.

Though we have focused on infinite-horizon dynamic programs, some problems would be more naturally formulated as finite-horizon dynamic program if the problem has a natural horizon (e.g., the options expire at some time) or if the transition probabilities and payoffs are nonstationary in that they explicitly depend on time. In these cases, we can formulate and solve the dynamic programs using lattice techniques, including the "binomial" or "trinomial trees" frequently used in the option-pricing literature. As with the infinite-horizon case, the key to managing these models is to define a relatively small set of state variables that evolve over time so that the lattice will not grow too large. Again, some fairly sophisticated programming is required to formulate and solve these problems, particularly when there are multiple state variables. In some cases, we may need to "paste together" finite- and infinite-horizon dynamic programs to model both expiring and nonexpiring options, perhaps using a finite-horizon model to determine payoffs for an infinite horizon model or vice versa.[2]

We also considered the possibility of using simulation techniques to solve these kinds of problems. While this approach is easy to implement using commercial software (such as @RISK or Crystal Ball®) and can easily handle many uncertainties, it is difficult to incorporate downstream decisions like those in Figures 2 or 3. The problem with the simulation approach is that it is difficult to determine the optimal policies for the downstream decisions: Though it is easy to calculate expected values and distributions of cash flows given a specific policy for all decisions, it is difficult to determine policies that maximize expected values given the information available at the time the decisions are made. While one could, in principle, use simulation to calculate expected values for all possible policies and then choose the optimal policy, practically, the number of possible policies grows much too quickly for this to be a viable approach.

1.3. Benefits of Modeling Flexibility

To capture flexibility, we must assess and solve more complex decision models. What are the benefits of modeling these project flexibilities? Decision theory suggests that incorporating flexibilities can only increase the values calculated for the project, because one could always choose

the base case alternative assumed in the nonflexible model. In practice, managers often took flexibilities into account informally and intuitively and incorporating flexibility would make a project more or less attractive depending on how the results of the analysis compared to these intuitive evaluations. For example, the "call option" feature of project X (the ability to wait for higher prices before developing the field) was initially viewed as potentially providing substantial value not captured in the tree of Figure 1; yet, as we will see shortly, our new analysis shows this option has no value. In general, these kinds of options are difficult to value intuitively, and one benefit of modeling flexibilities is that it improves the accuracy of these valuations and makes them more consistent across different managers.

A second and more important benefit is that in attempting to model project flexibilities we often identify new options and strategies. In constructing these models, we found it useful to ask questions like: "How could we use this information?", "What would we really do in this scenario?", and "What would I like to know before making this decision?" For example, in Project X, in thinking about how management might use the results from early drilling and well tests, one possibility that was identified was using this information to optimize the design of the production facility. While there was relatively little flexibility in the design of the production platform once construction had begun, there was some flexibility in the design of the system for transporting oil to market. If, for example, management were to learn from early drilling results that production rates would be less than expected, they could save substantial amounts by reducing the capacity of the transportation system (e.g., building one tanker instead of two). Here we are actually creating value for the project: While some of these options would be discovered in due course (in which case the benefit of identifying them now is through the improved measurement of the value of the project), some of them, like the flexibility in the transportation system, might be lost if management did not identify them up front and take steps to preserve these flexibilities.

A final benefit from modeling flexibilities is the set of optimal policies generated by the analysis. While the traditional analysis (e.g., that of Figure 1) generates an initial decision and value, the models that incorporate these downstream decisions generate an optimal policy that specifies, for example, when Project X should be developed and when production should be shifted to nearby fields. Such a policy might say, for example, "Begin Project X when prices reach $25 per barrel" or "If prices are below $15 per barrel and production at the main field is below 1,000 barrels a day, then move to the nearby field." These kinds of results provide management with "signposts" that suggest changes to (or at least a re-evaluation of) their operating procedures under certain conditions.

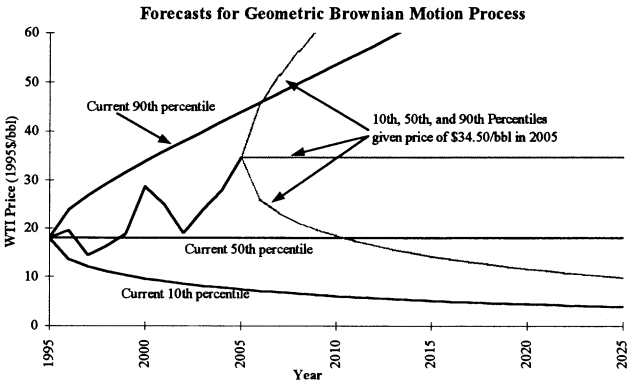

Figure 4. Probability forecasts for geometric Brownian motion process.

1.4. Stochastic Process Assumptions

To evaluate these flexible decision models, we had to ask and answer some new questions. While the decision tree of Figure 1 requires the specification of high, medium, and low scenarios for prices, production, and costs, the tree of Figure 2 requires a series of conditional probability distributions. Now, in addition to specifying a probability distribution for oil prices in the current year, in order to determine what the company should do if they wait for a year we need to specify distributions for prices the next year given prices from the first year. Similarly, we need to specify distributions for prices in year 3 conditioned on prices for years 1 and 2, and so on for subsequent years. Production rates and costs would be treated similarly. These conditional assessments were new questions for the company: While they construct corporate high, medium, and low price scenarios (of the kind required in Figure 1) for use throughout the corporation, they have rarely considered conditional price, production, and costs forecasts of the kind required by the tree of Figure 2. As we will see, the results of our analysis depend critically on the assumptions about these conditional distributions.

Most of the real options literature assumes the underlying uncertainty (in this case, oil prices) follows a random walk, specifically geometric Brownian motion (see Appendix 1 for a detailed description of this model). In this model, oil prices at any future time are lognormally distributed with the conditional distribution for later prices shifting by the amount of any (unexpected) change in prices in the early years. Figure 4 shows a representative price series and 10-, 50-, and 90-percent "confidence bands" for future oil prices generated by this process; the prices modeled here represent spot prices for West Texas Intermediate grade crude oil for delivery in Cushing, Oklahoma. The parameters for the process are based on historical estimates using annual data from 1900–1994.[3] The first set of bands fanning out from the current (1995) price of $18.00 per barrel show probability forecasts conditioned on today's prices. For example, in the year 2000, there is a 90-percent chance of oil prices being less than $33.79 per barrel; a 50-percent

chance of being less than $18.00 per barrel, and a 10-percent chance of being less than $9.59 per barrel; these confidence bands diverge rapidly and become extremely wide for distant years. The jagged path emanating from the current price represents one randomly generated set of future oil prices and the second set of confidence bands represent probability forecasts conditioned on prices being $34.50 in 2005. Here we see that a price increase in the early years lifts all three bands and increases the rate of divergence of the outer confidence bands.

Though this price model is the most frequently used model in the real options literature, the assumptions underlying this process were not consistent with the beliefs of the managers in the firm undertaking the study. They argued that, when prices are high compared to some long-run average (or equilibrium price level), new production capacity comes on line, that older production expected to come off line stays on line, and prices tend to be driven back down toward this long-run average. Conversely, if prices are lower than this long-run average, less new production comes on line, older properties are shut down earlier, and prices tend to be driven back up. Thus oil prices should be "mean reverting" in that prices tend to revert to some long-run average. Both historical spot and futures prices tend to support this view. Looking at the historical data from 1900–1994, prices have averaged about $18.00 per barrel in 1995 dollars; and deviations from this average, either above or below, have been followed by reversions back to this level. Similarly, in the futures markets, when spot prices are high compared to historical levels the 6- and 12-month futures prices tend to be lower than spot prices, and when prices are low the 6- and 12-month futures prices tend to be higher than spot prices.

To capture the phenomenon of mean reversion, we used a mean-reverting stochastic process for oil prices where future prices are expected to drift back to a specified long-run (perhaps inflation or growth-adjusted) average price. The particular form we used assumes that the logarithm of oil prices follows an "Ornstein-Uhlenbeck" process and was chosen both for its analytic tractability and its ability to fit historical and futures price data (see Appendix 1 for details). We again use annual data from 1900–1994 to estimate the model parameters. The confidence bands for this process

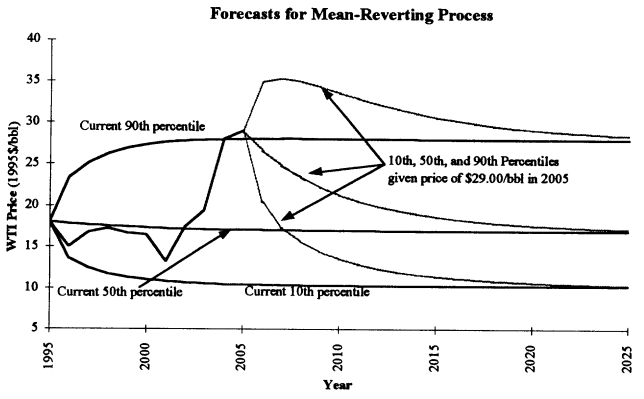

Figure 5. Probability forecasts for mean-reverting price process.

are shown in Figure 5. Like Figure 4, the first set of confidence bands are based on today's prices and the set of bands starting in 2005 are conditioned on a price of $29 in 2005. Comparing these bands to those of Figure 4, we see that there is less fanning out in the confidence bands, and when prices are above the long-run average, the bands tend to point back down to this average. In the limit, the two sets of bands for the mean-reverting process converge and prices in the distant future are independent of current prices. The rate of mean reversion in the model can be interpreted as being like a "half-life": Using historical estimates, we find that deviations from the long-run average are expected to decay by half their magnitude in about four years.

1.5. Results

These different price processes lead to markedly different valuations and strategies for Project X; the results are summarized in Table I. In the nonmean-reverting case (i.e., assuming geometric Brownian motion), using the tree of Figure 3, a 10-percent discount rate, and a current oil price of $18.00 per barrel, we find that the expected NPV of the field is $1,623MM. The distribution of NPVs shows the field has tremendous upside potential (there is a 10-percent chance of exceeding $4,870MM in NPV), reflecting the possibility of sustaining high future prices implied by this nonmean-reverting model of oil prices. The optimal strategy from this tree is to develop the field now. To get a

Table I
Results for Project X

	Brownian Motion Price Model		Mean-Reverting Price Model	
	with Flexibility	without Flexibility	with Flexibility	without Flexibility
Expected Value ($mm)	1,633	770	740	421
Optimal Development Threshold ($/bbl)	17	—	7	—
Percentiles from Distribution of NPVs ($mm)				
10th Percentile	−750	−665	−150	−340
50th Percentile	865	480	725	405
90th Percentile	4,870	2,550	1,610	1,165

better sense of the exact optimal development threshold, we used the dynamic programming model of the initial development decision (described in Section 1.2) with time steps that were approximately 2 weeks in length, as compared to the 5-year increment assumed in the tree of Figure 3. According to this model, the optimal development strategy is to develop the field when oil prices exceed $17.00 per barrel (in 1995 dollars). Thus, with current prices at $18.00 per barrel, they should immediately develop the field, but if prices were below this $17 threshold, they should wait for higher prices. Thus, with prices at their current level, the option to wait and develop the field has no value. At lower prices this option has some value but, until prices drop below about $14 per barrel, the additional value given by waiting is slight.

To illustrate the value of other options associated with the project, Table I also shows results for the case where we remove all of the downstream decisions from the tree of Figure 3 and assume the project is developed immediately, two tankers are used, and the nearby fields are not developed; these assumptions mirror the assumptions made in the original analysis of Figure 1. The result in this case is a much lower expected value ($770MM instead of $1,623MM) and a distribution with a similar downside but much less upside. The difference in upsides and expected values reflects the omission of the option to develop the nearby fields. The two investments have similar downsides as the flexible decision model will choose not to develop the nearby fields in most of those cases where development is economically unattractive.

If we assume that oil prices follow the mean-reverting process rather than the Brownian motion process, using the tree of Figure 3 and maintaining all other assumptions, we find a much lower expected value of $740MM (versus $1,623MM) and much less uncertainty in value. This reduction in risks and expected values reflects the decreased probability of sustaining either high or low oil prices. In this case, the values are not so sensitive to current prices and, as long as prices are above about $7 per barrel, they should develop the field. This insensitivity to current prices is a result of long development lags in the project: It is about eight years from the initial "go" decision until oil is first produced, and if prices are high when development begins, one would expect prices to revert before production begins. In this case, the optimal decision was always to go with two tankers, and in the tie-in decision the optimal decision was always to develop the neighboring field. This is again a reflection of mean reversion: Just as in the initial development decision, the critical question is whether the two tankers and nearby fields are economic at long-run prices. The difference in expected values in the flexible and nonflexible cases shows that the ability to develop the nearby fields is worth approximately $320MM in expected value.

The lessons here are twofold: (1) mean reversion greatly decreases the value of waiting to develop, particularly when facing long lead times; and (2) because it implies narrower long-run confidence bands with a decreased probability of sustaining high or low future prices, mean reversion implies there is significantly less risk and value associated with long-term oil projects than implied by the non-mean-reverting model. When contemplating the development of long-term projects, the mean-reverting model suggests that the critical question is whether the project is profitable at long-run average prices; current prices are not particularly relevant.

One should not draw the conclusion that mean-reversion eliminates the value of all flexibilities. For example, suppose management had the ability to quickly adjust production rates, perhaps by rapidly drilling and completing more wells or temporarily curtailing production. These kinds of flexibilities may have substantial value even in a mean-reverting price environment. If, for example, prices drop below some threshold, it could be optimal to temporarily shut-in production at certain wells or, if above some other threshold, it could be optimal to drill additional wells. The real lesson is that we need to think carefully about the conditional distributions (or stochastic processes) in the model and focus our analysis on options that can take advantage of the learning that takes place over time.

2. VALUATION METHODOLOGY

Insofar as modeling flexibility is concerned, the option pricing and decision analysis approaches are identical—the issues discussed in the previous section pertain equally to both approaches. Where the two methods differ is in how they value risky cash flows. In the decision analysis approach, firms typically attempt to incorporate any risk premiums required by shareholders by adjusting the discount rate used in calculating NPVs. In the option pricing approach, one uses futures and options prices to estimate risk-adjusted probabilities and discounts at the risk-free rate. In this section, we briefly describe each approach, then consider some practical issues in estimating these risk-adjusted probabilities and compare results given by the two approaches for some real projects.

2.1. The Risk-Adjusted Discount Rate Approach

Like many others, the company undertaking this study uses decision analytic techniques in an effort to determine the value-maximizing strategy for managing a given project; by "value," they mean market value or shareholder value. Traditionally, they have attempted to align their preferences for cash flows over time with shareholder's expectations of returns by using a discount rate equal to the firm's "weighted average cost of capital." This weighted average cost of capital is determined by considering the firm's financial structure, its marginal tax rate, its borrowing rate, and the expected rate of return on the firm's stock as estimated using, for example, the capital asset pricing model (CAPM). The justification of this approach relies primarily on the model used to determine the expected

return on the equity. If we use the CAPM to determine the expected return on equity, the project is valued as if it were a stock that satisfies the assumptions of the CAPM. While these assumptions might be appropriate for valuing companies as a whole, some of these assumptions—especially the assumption that project returns are normally distributed and jointly normal with the market as a whole—seem inappropriate when applied to flexible projects that might possess highly asymmetric distributions of returns.

To illustrate the determination of a weighted average cost of capital (or WACC), we calculate it for a hypothetical firm similar to the one we worked with; the values are chosen to be representative of a major oil and gas company. If we consider a firm with only debt and equity,[4] the WACC would be given by:

$$\begin{aligned}\text{WACC} &= (D/V)(1 - T_c)r_d + (S/V)r_e \\ &= (7/27)(1 - 34\%)8.5\% + (20/27)11.5\% \\ &= 9.97\%,\end{aligned}$$

where D (= \$7 billion) is the market value of the firm's interest-bearing debt, S (= \$20 billion) is the market value of the equity, $V = D + S$ (= \$27 billion) is the market value of the firm, T_c (= 34%) is the corporate tax rate, and r_d (= 8.5%) is the pre-tax yield on the firm's debt, and r_e (= 11.5%) is the firm's expected return on equity as given by the CAPM. This expected return on equity is given by the CAPM as:

$$\begin{aligned}r_e &= r_f + (r_m - r_f)\beta = 7\% + (13\% - 7\%)(0.80) \\ &= 11.50\%,\end{aligned}$$

where r_f (= 75) is the risk-free rate (given as the yield on long-term government bonds), r_m (= 13%) is the expected return on the market portfolio (say, the S&P 500), and β (= 0.80) is the "beta" of the firm, a measure of the correlation between the return on the firm's stock and the return on the market portfolio. The weighted average cost of capital in this example is 9.97 percent—we will round off to 10 percent—which is then applied to after-tax, then-current cash flows. The firm we worked with uses its WACC for all projects and for all scenarios considered. For a given project, the NPVs in each scenario are weighted by their respective probabilities—reflecting the firm's beliefs about future oil prices, production rates, costs, etc.—to determine the expected NPV for the project. Though the firm looks at other financial measures associated with a project (e.g., internal rates of return and various productivity indices), this expected NPV is taken to represent the value of the project.

While this cost-of-capital-based discounting rule may, in some sense, be right "on average" for the company, it can lead to trouble when applied to projects that are significantly different from the firm as a whole. If you are going to use risk-adjusted discount rates, you should use different discount rates for different projects, evaluating each on the basis of their own cost of capital. To do this, you need to somehow estimate the correlation between the project

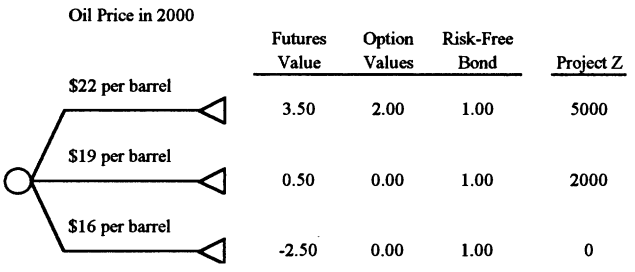

Figure 6. A simple numerical example.

returns and the market as a whole, either by identifying betas for firms that are "similar in risk" to the project or by making a difficult, subjective estimate of the beta (see Brealey and Myers 1991, p. 181–183, also p. 919). Given a flexible project, you might need to go one step further and use different discount rates for different time periods and different scenarios as the risks of a project may change over time, depending on how uncertainties unfold and management reacts. For example in Project X, the risks associated with the later cash flows are very different in the case where they choose to expand development as compared to the cases where they abandon the project after the main field declines. While, in principle, one could use time- and state-varying discount rates to value flexible projects, it becomes very difficult to determine the appropriate discount rates to be used in this framework.

2.2. The Option Valuation Approach

Rather than risk-adjusting discount rates to capture risk premiums, the option pricing (or "contingent claims") approach uses information from securities markets to value market risks more precisely. We illustrate this valuation procedure by considering the simplified example illustrated in Figure 6. In this example, we focus on oil price risks in the year 2000 and assume that there are three possible spot prices for oil in that year: \$16, \$19, or \$22 dollars per barrel. The approach assumes that investors can buy or sell securities in any desired quantity (including fractional and negative amounts) at market prices with no transactions costs. There are three securities in this example, defined and priced as follows:

Futures Contract: The futures contract obligates the buyer to buy (and the seller to sell) a barrel of oil in the year 2000 for a fixed price of, say, \$18.50. Thus, if the oil price is \$22 per barrel in the year 2000, this contract will then be worth \$22.00 − \$18.50 = \$3.50, as the contract would deliver a barrel of oil for \$18.50, which could then be sold on the spot market for \$22.00. Similarly, if the oil price is \$19 or \$16 per barrel in 2000, the futures contract is worth \$0.50 or −\$2.50, respectively.

Call Option: The call option gives the investor the right, but not the obligation, to buy a barrel of oil in the year 2000 at a price of \$20 per barrel, and it may be purchased for a current price of \$0.40 per contract. If the oil price

turns out to be \$22, the holder of the option would exercise the option, by buying oil at \$20 per barrel and selling it on the spot market at \$22 per barrel, for a profit of \$2.00 per barrel. If prices are \$19 or \$16 dollars per barrel, the holder of the option would decline to exercise the option and let it expire worthless.

Risk-Free Bond: The risk-free bond may be purchased for \$0.7629 today and returns a certain \$1.00 in the year 2000. (This is a "zero coupon" bond that pays no interest in intermediate years.) The current year is assumed to be 1996, so the price for the bond corresponds to a risk-free interest rate (r_f) of 7 percent per year (viz., $1/(1 + r_f)^4 = 0.7629$).

The basic idea of the option valuation procedure is to use prices for traded securities to determine the market value of related cash flows. We can, for example, determine the value of a portfolio that pays \$1.00 if the price of oil is \$16.00 in the year 2000, and \$0.00 in the other states. To determine this portfolio, we let w_1, w_2, and w_3 denote the shares of the futures contract, call option, and risk-free bond, respectively, and solve the following set of linear equations specifying portfolio values in each state:

$$\$16 \text{ state:} \quad w_1(-\$2.50) + w_2(\$0.00) + w_3(\$1.00) = \$1.00;$$

$$\$19 \text{ state:} \quad w_1(\$0.50) + w_2(\$0.00) + w_3(\$1.00) = \$0.00;$$

$$\$22 \text{ state:} \quad w_1(\$3.50) + w_2(\$2.00) + w_3(\$1.00) = \$0.00.$$

The solution is a portfolio consisting of $w_1 = -1/3$ futures contracts, $w_2 = 1/2$ share of the call option, and $w_3 = 1/6$ share of the risk-free bond. The value of this portfolio can be interpreted as a "state price" representing the present value of \$1 paid in the year 2000 if and only if the oil price is then \$16 per barrel. The state price for the \$16 price scenario is $-1/3(\$0.00) + 1/2(\$0.40) + 1/6(\$0.7629) = \$0.3271$.[5] We can similarly calculate state prices of \$0.2357 and \$0.2000 for the \$19 and \$22 price scenarios, respectively.

Using these state prices, we can determine the market value for any project whose payoffs depend only on the price of oil in the year 2000. For computational purposes, it is convenient to renormalize these state prices by dividing by $1/(1 + r_f)^t$ where r_f is the risk-free rate and t is the time when the claims are paid. Since the present value of the risk-free bond must equal its future value discounted at the risk-free rate, these normalized state prices must sum to one and we can interpret these normalized state prices as risk-adjusted probabilities. In the example, we divide the state prices by $1/(1 + .07)^4 = 0.7629$ (the current price of the risk-free bond) and obtain risk-adjusted probabilities of 0.4288, 0.3090, and 0.2622 for the \$16, \$19, and \$22 price states, respectively. In this interpretation, the value of any project or any security is given by calculating its expected future value using these risk adjusted probabilities and discounting at the risk-free rate.

To illustrate this approach, consider a hypothetical project, call it Project Z, that can, if the firm chooses, produce 1,000 barrels of oil in the year 2000 at a cost of \$17 per barrel. To keep the example as simple as possible, we assume that this is a one-shot deal: if the firm chooses not to produce in 2000, the project generates no cash flows at any other time. In the \$16 price state, assuming the firm chooses not to produce, project Z would be worth \$0. In the \$19 price state, Project Z would be worth \$2,000 (= 1000 (\$19 − \$17)); in the \$22 state, it would be worth \$5,000. The value of project Z is then equal to:

$$\frac{1}{(1 + .07)^4}(0.4288(\$0) + 0.3090(\$2{,}000) + 0.2622(\$5{,}000)) = \$1{,}471. \tag{1}$$

These risk-adjusted probabilities can be viewed as providing a shortcut method for computing the market value of a portfolio that exactly matches the project payoffs in all price states. In this example, the project is exactly replicated by a portfolio of 666.67 futures contracts, 500 call options, and 1,666.7 risk-free bonds; the current market value of this portfolio is exactly the value given by using the state prices in Equation (1).

Note that the firm's probabilities and risk preferences are not used anywhere in the options approach. In this framework, it is the market's beliefs and preferences that are important and these are reflected in the risk-adjusted probabilities. Also note that you need not adjust the probabilities or discount rate depending on the features of the project being valued. While in the risk-adjusted discount rate approach, you should use different discount rates for different projects or even different states of the world, here you use the same risk-adjusted probabilities to value a futures contract, a call option, or a real project with embedded options. This is a key advantage of the option valuation approach over the risk-adjusted discount rate approach when valuing complex, flexible projects.

In order for this option valuation approach to be practical for real projects, we need to be able to value projects that cannot be replicated by portfolios of existing securities. While there are well-developed financial markets for managing oil and gas price risks, there are no securities for hedging project-specific risks like the production at Project X. The classic option valuation theory assumes that markets are complete in that all project risks can be perfectly hedged by trading securities, perhaps dynamically over time. With incomplete markets, we can extend the option pricing approach to distinguish between "market risks" that can be hedged by trading securities (e.g., oil price risks) and "private risks" that cannot be hedged by trading existing securities (e.g., production risks, cost risks, etc.). In this integrated approach, we use option valuation techniques to value market risks and traditional decision analytic techniques to value private risks.

This integrated approach works as follows. Assuming the firm is risk-neutral (as was the case with the firm in this study), we use the firm's probabilities to determine the expected value of the project conditioned on the occurrence of a particular market state. The value of the project is then given by using the market-based, risk-adjusted probabilities to calculate the expected value of these market-state-contingent values, and discounting at the risk-free rate. For example, let us reconsider Project Z and assume that production is equally likely to be either 500 or 1,500 barrels of oil (independent of the price of oil) rather than being 1,000 barrels for sure. The market-state-contingent expected values are then $0, $2,000 and $5,000 in the $16, $19, and $22 price states (as in the case where production was assumed to 1000 barrels for sure), and the overall value is given using the risk-adjusted probabilities as in Equation (2). Any dependence between the market and private risks is captured by conditioning the probabilities and expected values for the private risks on the outcome of the market risks.

This option valuation procedure and its extension can be applied recursively in a multiperiod setting. To do this, construct a decision tree or dynamic program that uses risk-adjusted probabilities for the market risks and ordinary probabilities for the private risks, taking care to model how the private risks depend on the market risks. Again assuming the firm is risk-neutral, roll back the tree or solve the dynamic program by calculating expected values in the usual way using these mixed probabilities, making decisions to maximize these expected values, and discounting at the risk-free rate. The values generated by this procedure can be interpreted as present certainty equivalent values: taking into account all project decisions and risks, as well as all trading opportunities related to the project, the value generated by this procedure is the amount such that the firm is just indifferent between undertaking the project and receiving this amount as a lump sum, with certainty, today. Thus, the firm would want to invest in Project Z if and only if it costs less than $1,520 in present dollar terms.[6]

2.3. Estimating "Risk-Adjusted" Probability Distributions

To apply the option valuation technique with real projects, we need to determine the appropriate risk-adjusted probabilities and, more generally, a risk-adjusted stochastic process describing the evolution of these probabilities over time. To do this, we will assume a particular functional form for the risk-adjusted stochastic process and estimate the parameters for that process from the available futures and options prices. While the standard Black-Scholes option pricing model assumes the risk-adjusted stochastic process follows geometric Brownian motion process (reflecting the assumption that the true price process has that form), we will use the mean-reverting price model described in the previous section and estimate its parameters to match the futures and options prices listed in the *Wall Street Journal*. On August 15, 1995 (the date the analysis was done), the *Wall Street Journal* listed prices for 24 futures contracts, one for each month from September 1995 to March 1997, plus five contracts ranging out as far as December of 1999. There were prices for 31 different option contracts, with strike prices ranging from $16.00 to $18.50 and expiration dates ranging from October to December of 1995. Here we are constructing "implied" estimates of the parameters of the risk-adjusted stochastic process, backing them out from current prices for futures and options. Alternatively, one could use historical futures and options prices to estimate the parameters for this risk-adjusted process.

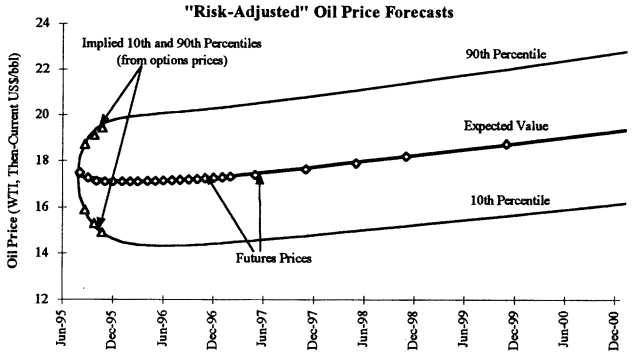

Figure 7. Risk-adjusted oil price forecasts.

In the option valuation approach, the value of each security should be equal to its expected future value, where expectations are calculated using these risk-adjusted probabilities and discounting is done at the risk-free rate. Accordingly, we selected our parameters for the mean-reverting price model to minimize the squared errors in futures and options prices, where the errors are the differences between the discounted expected values calculated by the model and the prices listed in the *Wall Street Journal*. The results are summarized in Figure 7 and the details are described in Appendix 2. In this approach, the futures prices should be equal to the expected (risk-adjusted) oil price. In Figure 7, we see that the expected values of the mean-reverting process (shown with the bold line) provide a very good fit to the futures prices; the model correctly mimics the initial decline in prices for near-month futures, followed by an increase in the longer term futures prices. The option prices provide information about the uncertainty in these risk-adjusted price forecasts. To place the option prices back on the same scale as the futures prices, we have used the listed options prices to estimate confidence bands (10th and 90th percentiles) for the risk-adjusted distribution for oil prices in the month of expiration, using the current price for options expiring in that month. Comparing these implied confidence bands to those from the mean-reverting model (or comparing the direct estimates of put and call prices), we see that the estimated put and call prices generated by the mean-reverting model are very close to their true prices.

The parameter estimates and price forecasts for this risk-adjusted stochastic process are quite different from the unadjusted forecasts based on historical, annual price data from 1900–1994. Compared to the unadjusted historical estimates (see Figure 5), the risk-adjusted estimates have lower expected values, have narrower confidence bands, and revert much faster. Because the risk-adjusted forecasts reflect both market opinions and risk premiums, it is difficult to discern the reasons for these differences. It could be that the market does not view the historical data as a good predictor of the future (this could explain the differences in reversion rates and confidence bands) or it could be a reflection of the market risk premiums for oil price exposure. The difference in means is most likely a combination of both these factors. In reviewing a number of different oil price forecasts provided to the firm by government sources and private consultants, we found some forecasts above and some below our historical expected values, but no forecasts below the futures prices. This is evidence that there is some risk premium embedded in the risk-adjusted price forecasts and that the difference in means does not simply reflect a change in beliefs.

One major problem in using the futures and options markets to generate the risk-adjusted oil price forecasts is that the maturities of the exchange-traded futures and options contracts are much shorter than the time horizons of the projects we are interested in evaluating. While the projects may last 30 or 40 years, the futures contracts go out less than 5 years and the options contracts go out only 4 months.[7] Thus, we need to somehow extrapolate from these shorter term risk-adjusted forecasts. In performing this extrapolation, it is important to remember that we are not attempting to forecast what oil prices will be after the year 2000. Instead, we are asking what an oil futures or option contract maturing in say, 2010, would trade for today: it is not the firm's projections of future oil prices that matters, so much as the current market assessment. Here, we extrapolate using our mean-reverting price model, estimating its parameters with the near-term market data and assuming these estimates hold going forward.

2.4. Project Valuation

Before we consider results for some actual projects, to demonstrate the effects of the different valuation methodologies let us first consider the value of a hypothetical project that produces a single barrel of oil in a specified year. To isolate the effect of the valuation methodology, we will assume that there are no costs associated with this production, no uncertainty about the amount produced, no royalties or taxes, and no basis risks: The project produces one barrel of West Texas Intermediate grade oil in Cushing, Oklahoma. The results of this comparison are summarized in Figure 8. The bold line indicates values generated using the option valuation approach: Here we calculate expected net present values using the risk-adjusted probabilities and discounting at a risk-free rate of 7 percent, reflecting the yield on long-term Treasury securities in

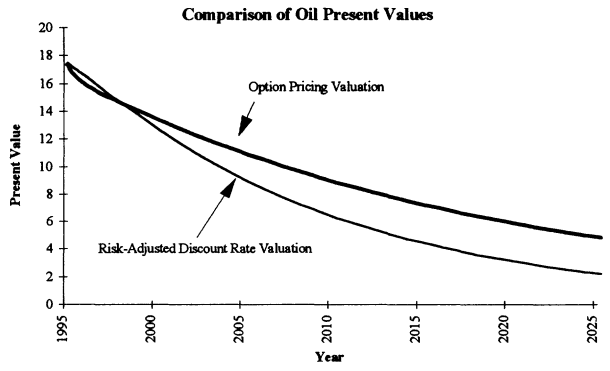

Figure 8. Present value of a barrel of oil produced in different years.

August of 1995. The lighter line indicates values calculated using the risk-adjusted discount rate approach with historically based probabilities and discounting at the firm's weighted average cost of capital of 10 percent. In both cases, we have used our mean-reverting model for oil prices.

In Figure 8, we see some support for management's hypothesis that the blanket use of a 10-percent discount rate biases evaluations against the long-term projects: While the risk-adjusted discount rate approach slightly overestimates the values of oil produced in the near future, it severely underestimates the value of distant production. For example, the options approach shows the present value of a barrel of oil produced in the year 2025 to be \$4.78 and the risk-adjusted discount rate approach shows a value of \$2.19. Thus we see that the market-required risk premiums do not grow as fast as those implied by compounding the risk-adjusted discount rate. While these specific numbers reflect the particular price forecast used with the risk-adjusted discount rate, the effect is fairly robust: even if we double the price forecasts, the risk-adjusted discount rate approach would give a present value of \$4.38 for a barrel of oil delivered in 2025, which is still less than the value given by the option valuation approach.

We applied these valuation techniques to two real projects, Project X (which we discussed in the previous section) and another project—a large undeveloped field in a remote region—which we will call Project Y. As before, both evaluations use the mean-reverting price process in decision tree models; we used the model of Figure 3 for Project X and model of similar complexity for Project Y. The only difference between the risk-adjusted discount rate and option valuations is in the parameters of the oil price process and the discount rate. The assumptions in both cases are exactly as in Figure 8.

For project X, we find values of \$740MM and \$1,265MM for the risk-adjusted discount rate approach and option valuation approach, respectively. This is what one might expect given the results of Figure 8. Here, most of the capital expenditures occur in the first eight years, followed by a long stream of oil production extending out

to approximately 2030. The options procedure values this future production more highly than the risk-adjusted discount rate procedure and leads to a higher project value. Though the optimal development policy is the same in the two methods (they should begin development immediately, always use two tankers, and expand to the nearby fields), the difference in values suggests very different behavior in an acquisition or divestiture setting.

The change in values for Project Y are more extreme: The risk-adjusted discount rate approach leads to value of $102MM, while the options analysis leads to a value of $721MM. Project Y, like Project X, is a long-term project with a significant delay before production begins. The difference in values here reflects a substantial change in optimal policies for the project. This project requires a big investment in infrastructure (pipelines, processing facilities, etc.) that would be shared with other producers in the area and could generate tariff revenue after Project Y's main field begins to decline. The model for the project includes the ability to adjust the size of the equity position in this infrastructure (depending on prices, costs, and early well tests) and the change in the valuation procedure leads to a substantial change in optimal investment in this infrastructure. In the risk-adjusted discount rate approach, this infrastructure investment was seen as a necessary evil required to bring the reserves to market in a timely manner. In the options approach, the tariff revenue itself looks much more attractive, and the optimal policy calls for taking a larger equity position in the infrastructure. Thus, in this case, we see that the valuation procedure leads to substantial changes in policies as well as changes in values.

While in both examples the options approach leads to higher values, in other cases the options approach leads to lower values. While the options approach values future oil production more highly (as indicated in Figure 8), it also discounts the future costs less. While in Projects X and Y, we saw a substantial increase in the overall present value of the project, in other cases, the reduced discounting of costs may dominate the effect due to the increased valuation of production. For example, in the case of a mature oil field that we examined the production costs were quite high (secondary and tertiary recovery methods were being used) and the reduced discounting of these costs lead to a lower overall project value.

3. CONCLUSIONS

Our first major conclusion from this study is that the option pricing and decision analysis approaches are equally capable of modeling flexibility. In both approaches, the evaluation models correspond to constructing a decision tree or dynamic programming model that describes the sequence of decisions to be made and the resolution of uncertainties over time. Despite the ubiquity of options in business and everyday life, in practice we find that embedded options are often overlooked in the formulation and evaluation of decision problems, even when uncertainties are explicitly modeled (see Howard 1996 for similar observations). One possible reason for this is the difficulty of evaluating decision problems that include many downstream decisions. To properly evaluate these downstream decisions, you must model not only the downstream decisions, but also the information available at the time these decisions are made. While decision analysts have developed techniques for assessing probabilities for simple random variables, with flexible decision models, we need to consider some complex conditional probability or stochastic process assessments. In our example of project X, we examined the pivotal role of the price processes; in other cases, the key questions concern learning about production or costs over time.

Our second main conclusion from our experience to date is that the option pricing and decision analysis techniques should be viewed as complementary modeling approaches that can be nicely integrated. While we are not comfortable with applying off-the-shelf option pricing models to real projects (for example, using the Black-Scholes formula to value an undeveloped oil field), the options approach provides a simple technique for incorporating market information into project values that can be easily incorporated into decision-analytic or dynamic programming models. To implement this approach, you risk-adjust the probabilities for market risks and use a risk-free discount rate and solve the models using standard dynamic programming techniques. Viewing the option pricing methods as a refinement on decision analytic and dynamic programming methods, we can draw on the experience and expertise of the operations research and management science communities when applying these options techniques to real projects.

While we have focused on oil and gas applications, most of the methodology and lessons have broader applications. One could easily imagine similar models and issues for producers of other commodities (e.g., gold, copper, coal, etc.) or consumers of these commodities (for example, an electric utility that consumes oil, gas, and/or coal). Similar issues arise in manufacturing: products may be manufactured and distributed in a variety of different countries. These firms may be able to use securities markets to help value exchange rate risks while relying on their own judgment to value private risks, like the demand for their product. In these applications, and others like them, it seems that there is much to be gained by integrating the option pricing and decision analytic methods of evaluation.

APPENDIX 1: OIL PRICE MODELS

Geometric Brownian Motion. The geometric Brownian motion model of oil prices is described by two parameters, μ_p and σ_p, representing the expected rate of change and volatility of the process. Given the current price $p(0)$, future oil prices $p(t)$ follow a stochastic process described by the following stochastic differential equation:

$$d(\ln(p(t))) = \mu_p \, dt + \sigma_p \, dz_p(t),$$

where dz_p represents increments of a standard Brownian motion process. Given $p(0)$, this process implies that $\ln(p(t))$ is normally distributed with mean $\ln(p(0)) + \mu_p t$ and variance $\sigma_p^2 t$. As a log-normal random variable, $p(t)$ has mean $p(0)\exp(\mu_p t + \sigma_p^2 t/2)$ and variance $p(0)^2\exp(2\mu_p t + \sigma_p^2 t)(\exp(\sigma_p^2 t) - 1)$. The conditional distributions for $p(t)$ given $p(\tau)$ are similar, with $p(0)$ replaced by $p(\tau)$ and t replaced by $t - \tau$. Using annual oil prices from 1900–1994 adjusted back to 1995 dollars, we find an estimated mean growth rate (μ_p) equal to 0 percent per year and a volatility (σ_p) equal to 22 percent per year. We discretized this process for use in the decision tree and dynamic programming models, using three-point approximations for the conditional distributions with probabilities and values selected to match the mean and variance of the exact conditional distributions for the time step of the model. This process describes movements in real prices; these were inflated to nominal prices in the model by assuming a constant inflation rate of 3 percent.

Mean-Reverting Model. Our mean-reverting model of oil prices assumes that the logarithm of oil prices, $\pi(t) = \ln(p(t))$, follows an Ornstein-Uhlenbeck process, i.e., future oil prices are described by the stochastic differential equation:

$$d\pi(t) = \kappa(\bar{\pi} - \pi(t))\, dt + \sigma_\pi\, dz_\pi(t), \tag{A1}$$

where $\bar{\pi}$ denotes the long-run mean to which log-prices revert, κ describes the strength of mean reversion, σ_π describes the volatility of the process, and $dz_\pi(t)$ represents increments of a standard Brownian motion process. This implies that, given $\pi(0)$, $\pi(t)$ is normally distributed with mean $\bar{\pi} + (\pi(0) - \bar{\pi})e^{-\kappa t}$ with variance $\sigma_\pi^2(1 - e^{-2\kappa t})/2\kappa$. In this form, we see that deviations from the long-run mean are expected to decay following an exponential decline. Given a current log-price $\pi(0)$ away from the long-run mean $\bar{\pi}$, in $-\ln(0.5)/\kappa$ years we would expect $\pi(t)$ to have reverted half way back to the long-run mean; thus we can interpret $-\ln(0.5)/\kappa$ years as representing the "half-life" of the mean-reverting process. Notice that in the limit as $t \to \infty$, the distribution for $\pi(t)$ is independent of the initial price and has mean $\bar{\pi}$ and variance $\sigma_\pi^2/2\kappa$.

Using annual oil prices from 1900–1994 adjusted to 1995 dollars, we estimate a long-run mean ($\bar{\pi}$) of 2.83 (corresponding to a long-run median oil price of $e^{2.83}$ = \$16.89 in 1995 dollars), a mean-reversion coefficient (κ) of 16.9 percent per year (corresponding to a half-life of 4.11 years), and a volatility (σ_π) of 22.8 percent per year. These parameter estimates imply that the long-run distribution of oil prices is lognormal with a mean of \$18.25 (in 1995 dollars) and a standard deviation of \$7.47. These distributions were discretized and inflated for use in the decision tree and dynamic programming models in the same way as in the geometric Brownian motion case.

APPENDIX 2: ESTIMATING RISK-ADJUSTED PROBABILITIES FROM FUTURES AND OPTIONS

In this appendix, we describe the details of how we estimated parameters for our mean-reverting oil price process based on listed futures and options prices. The key to this estimation process is developing formulas for determining futures and options prices using this price process. We then choose parameters of the process to fit the listed prices. There were a total of five parameters estimated for the process including the three parameters described above ($\bar{\pi}$, κ, and σ_π). In addition, we estimated the current spot price $p(0)$ and an implied inflation/growth rate (α). In this framework, we assume that $\pi(t) = \ln(p(t)) - \alpha t$ follows the mean-reverting process described by (A1). The spot price at time t is then given by $p(t) = e^{\pi(t)+\alpha t}$.

Valuing Futures Contracts. Using the risk-neutral valuation framework, the current futures price is equal to the expected spot price at the time the contract expires, where expectations are calculated using the risk-adjusted probability distribution (see Duffie 1992, p. 122). Assuming the risk-adjusted oil price process has this mean-reverting form, the futures price for a contract with t years before expiration is

$$\begin{aligned} \mathrm{E}[p(t)] &= \exp(\mathrm{E}[\pi(t)] + \mathrm{Var}[\pi(t)]) \\ &= \exp(\bar{\pi} + (\pi(0) - \bar{\pi})e^{-\kappa t} + \alpha t \\ &\quad + 0.5\sigma_\pi^2(1 - e^{-2\kappa t})/2\kappa). \end{aligned} \tag{A2}$$

The equation follows from the fact that $p(t)$ is a log-normal random variable with $\ln(p(t))$ having the mean and variance given by the assumed stochastic process for $\pi(t)$.

Valuing Futures Options Contracts. A "European" call option gives its owner the right to buy a futures contract at some specified "strike price" (K) on a specified exercise date (t). On the exercise date, the value of the call option is equal to $(F(t) - K)^+ = \max(F(t) - K, 0)$ where $F(t)$ denotes the price of the futures contract at the time the option expires. In the risk-neutral valuation framework, the current price of the call option is given as the expected value of the futures contract at the time the option expires, discounted to present values at the risk-free rate: $e^{-rt}\mathrm{E}[(F(t) - K)^+]$. In general the option contracts expire before the underlying futures contracts, and for short-term options it is important to recognize that these options represent options on futures contracts rather than options on the spot price of oil.

The key to valuing these options on futures contracts is to note that, if the spot price $p(t)$ follows our mean-reverting process, the value of the futures contract $F(t)$ follows the same form of process with transformed parameters. If we let T denote the time the futures contract expires and let $\Delta t = T - t$ denote the difference between the futures and option expirations, then from (A2) we see that at the time the option expires, given the then-prevailing (log) spot price $\pi(t)$, the futures price is then given by

$$\ln(F(t)) = \pi(t)e^{-\kappa\Delta t} + \bar{\pi}(1 - e^{-\kappa\Delta t}) + \alpha\Delta t + 0.5\sigma_\pi^2(1 - e^{-2\kappa\Delta t})/2\kappa.$$

Since the only uncertainty in this formula is $\pi(t)$, $\phi(t) \equiv \ln(F(t))$ is a linear transformation of a normally distributed random variable ($\pi(t)$), with mean and variance:

$$\mu_\phi(t) \equiv \mathrm{E}[\pi(t)]e^{-\kappa\Delta t} + \bar{\pi}(1 - e^{-\kappa\Delta t}) + \alpha\Delta t + 0.5\sigma_\pi^2(1 - e^{-2\kappa\Delta t})/2\kappa,$$

$$\sigma_\phi^2(t) \equiv \mathrm{Var}[\pi(t)]e^{-2\kappa\Delta t}.$$

$F(t)$ then is log-normally distributed with mean $\exp(\mu_\phi(t) + 0.5\sigma_\phi^2(t))$.

The value of the call option is given by $e^{-rt}\mathrm{E}[(F(t) - K)^+]$. To calculate these values we note that $\mathrm{E}[(F(t) - K)^+]$ can be evaluated analytically, resulting in a call option value given as

$$e^{-rt}\{\mathrm{E}[F(t)](1 - \mathrm{N}(\ln(K), \mu_\phi(t) + \sigma_\phi^2(t), \sigma_\phi^2(t))) - K(1 - \mathrm{N}(\ln(K), \mu_\phi(t), \sigma_\phi^2(t)))\},$$

where $\mathrm{N}(d, \mu, \sigma^2)$ indicates the tail probability for d (i.e., $\mathrm{P}(X < d)$) for a normally distributed random variable (X) with mean μ and variance σ^2. (The standard Black-Scholes formula can be written in this same form with the appropriate substitution for the means and variances.) Following a similar analysis, we find an analogous formula for determining the value of the put options:

$$e^{-rt}\{-\mathrm{E}[F(t)](1 - \mathrm{N}(\ln(K), \mu_\phi(t) + \sigma_\phi^2(t), \sigma_\phi^2(t))) + K(1 - \mathrm{N}(\ln(K), \mu_\phi(t), \sigma_\phi^2(t)))\}.$$

Throughout this discussion, we have assumed that the option contracts are European in nature allowing exercise only at expiration, when, in fact, the exchange traded options on oil futures contracts are American options that allow the holder to exercise at any time. The valuation of American contracts is more difficult than their European counterparts and typically requires the construction of a binomial tree or lattice approximation of the stochastic process (as in Cox et al. 1979). For very short-term options, like the exchange traded options on oil futures contracts, the European options usually provide a good approximation of the American values (Barone-Adesi and Whaley 1987). While the approximation may not be appropriate for longer term options, the long-term options that we have seen traded over the counter are, in fact, European option contracts. In valuing the put and call options, we used a risk-free rate of 5.60 percent, corresponding to the then-prevailing interest rate for three-month Treasury securities.

Results. Given these formulas for valuing futures and option contracts, we then chose parameter values to minimize squared errors in observed prices. Specifically, we chose parameters to minimize the sum of squared errors in prices, where the errors are defined as the difference between the "settle" price listed in the *Wall Street Journal* and the values given by our model, and each of the 55 securities was weighted equally in calculating the total squared error. The parameter estimates given by this process were as follows: $\bar{\pi} = 2.798$, $\kappa = 2.194$ per year (corresponding to a mean-reversion half-life of 3.79 months), $\sigma_\pi = 26.4$ percent per year, $\alpha = 2.90$ percent, and $p(0) =$ \$17.36. These parameters imply the long-run median oil price is \$16.41 ($= e^{\bar{\pi}}$) in 1995 dollars, where future dollars are discounted at $\alpha = 2.90$ percent.

With these parameters, the model provides a very good fit to the actual security prices. The largest error is 14 cents and occurs in the near-month futures contract: the model gives a price of \$17.32 and the actual price is \$17.48. After that the next largest error is 8 cents, with a mean absolute error of 3.3 cents on the 24 futures contracts (including the 16 cent error on the near term contract) and a mean absolute error of 2.1 cents on the 31 option contracts.

ACKNOWLEDGMENTS

The authors are grateful for the support of the National Science Foundation under grant no. SBR-9511364 to Duke University and gratefully acknowledge the contributions of their colleagues on the VMI working team and steering committee to the work described in this paper. They also thank an anonymous referee and associate editor for helpful comments.

ENDNOTES

1. Examples of real options applications in the oil and gas area (or natural resource extraction more generally) include Brennan and Schwartz (1985), Paddock et al. (1988), Lehman (1989), Trigeorgis (1990), Kemna (1993), and Smith and McCardle (1996).

2. Bertsekas (1996) discusses finite- as well as infinite-horizon dynamic programming in general (including how to formulate them as linear programs). Luenberger (1997) provides a nice discussion of lattice methods for the cases of one- and two-state variables. The stochastic differential equations often presented in the real options literature can be viewed as a limiting case of discrete-time dynamic programs with infinitesimal time steps (see Dixit and Pindyck 1994).

3. These oil price forecasts and all others presented and used in this paper are based on publicly available price data. These forecasts are intended to provide a basis for discussion and do not represent the company's or anyone else's recommended or actual price forecast.

4. For more on the weighted average cost of capital and the capital asset pricing model, see, for example, Brealey and Myers (1991) p. 465–470 or Copeland et al. (1990), p. 171–205. These books talk about how to incorporate more complex debt instruments as well as issues associated with estimating these parameters.

5. Though the futures contract commits the buyer and seller to potentially make payments in the future, the contract requires no initial payment and, hence, the current price of the contract is \$0.

6. The traditional option valuation theory (sometimes called "contingent claims valuation," "risk-neutral valuation," or "valuation by arbitrage") is discussed at an introductory level in Dixit and Pindyck (1994), and at a more advanced level by Harrison and Kreps (1979) and Duffie (1992). Note what we call "risk-adjusted probabilities" are often referred to as "risk-neutral probabilities" or "equivalent martingale measures." The extension to handle private risks is developed in Smith and Nau (1995) and Smith (1996) and discussed in Luenberger (1997). These references also discuss how to handle the case where the firm is risk-averse rather than risk-neutral.

7. Longer maturity futures and options are currently traded over-the-counter, though prices for such contracts are not readily available to those not active in those markets.

REFERENCES

Barone-Adesi, G., R. E. Whaley. 1987. Efficient analytic approximation of American option values. *J. Finance* **42** 301–320.

Bertsekas, D. P. 1995. *Dynamic Programming and Optimal Control: Volumes I and II*. Athena Scientific, Belmont, MA.

Black, F., M. Scholes. 1973. The pricing of options and corporate liabilities. *J. Political Economy* **81** 637–659.

Brealey, R., S. Myers. 1991. *Principles of Corporate Finance*. Fourth Edition. McGraw-Hill, New York.

Brennan, M. J., E. S. Schwartz. 1985. Evaluating natural resource investments. *J. Business* **58** 135–157; see also M. Brennan and E. Schwartz 1985. A new approach to evaluating natural resource investments. *Midland J. Applied Corporate Finance* **3** 37–47.

Copeland, T., T. Koller, J. Murrin. 1990. *Valuation: Measuring and Managing the Value of Companies*. John Wiley and Sons, New York.

Cox, J., S. Ross, M. Rubinstein. 1979. Option pricing: a simplified approach. *J. Financial Econom.* **7** 229–263.

Dixit, A. K., R. S. Pindyck. 1994. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.

Duffie, D. 1992. *Dynamic Asset Pricing Theory*. Princeton University Press, Princeton, NJ.

Harrison, J. M., D. M. Kreps. 1979. Martingales and arbitrage in multiperiod securities markets. *J. Econom. Theory* **20** 381–408.

Howard, R. A. 1996. Options. *Wise Choices*. R. J. Zeckhauser, R. L. Keeney, J. K. Sebenius (eds.). Harvard Business School Press, Boston, MA.

Kemna, A. G. Z. 1993. Case studies on real options. *Financial Management* **22** 259–270.

Lehman, J. 1989. Valuing oilfield investments using option pricing theory. *Proc. Soc. Petroleum Engineers* **18923** 125–136.

Luenberger, D. G. 1997. *Investment Science*. Oxford University Press, New York.

Merton, R. C. 1973. The theory of rational option pricing. *Bell J. Econom. Management Sci.* **4** 141–183.

Paddock, J. L., D. R. Siegel, J. L. Smith. 1988. Option valuation of claims on real assets: the case of offshore petroleum leases. *Quarterly J. Economics* **103** 479–508; see also D. R. Siegel, J. L. Smith, J. L. Paddock. 1987. Valuing Offshore Oil with Option Pricing Models. *Midland Corporate Finance J.* **5** 22–30.

Smith, J. E. 1996. Fisher separation and project valuation in partially complete markets. Working Paper. Fuqua School of Business, Duke University, Durham, NC.

——, R. F. Nau. 1995. Valuing risky projects: option pricing theory and decision analysis. *Management Sci.* **41** 795–816.

——, K. F. McCardle. 1996. Valuing oil properties: integrating option pricing and decision analysis approaches, *Oper. Res.* **46** 198–217.

Trigeorgis, L. 1990. A real options application in natural resource investments. *Advances in Futures and Options Res.* **4** 153–164.

——. 1996. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. MIT Press, Cambridge, MA.