Modeling Dependence Among Geologic Risks in Sequential Exploration Decisions

J. Eric Bickel, SPE, Texas A&M University; **James E. Smith,** SPE, Duke University; and **Jennifer L. Meyer,** SPE, Strategic Decisions Group

Summary

Prospects in a common basin are likely to share geologic features. For example, if hydrocarbons are found at one location, they may be more likely to be found at other nearby locations. When making drilling decisions, we should be able to exploit this dependence and use drilling results from one location to make more informed decisions about other nearby prospects. Moreover, we should consider these informational synergies when evaluating multiprospect exploration opportunities. In this paper, we describe an approach for modeling the dependence among prospects and determining an optimal drilling strategy that takes this information into account. We demonstrate this approach using an example involving five prospects. This example demonstrates the value of modeling dependence and the value of learning about individual geologic risk factors (e.g., from doing a postmortem at a failed well) when choosing a drilling strategy.

Introduction

When considering a new prospect, it is important to consider its probability of success. In practice, this assessment is often decomposed into success probabilities for a number of underlying geologic factors. For example, one might consider the probabilities that the hydrocarbons were generated, whether the reservoir rocks have the appropriate porosity and permeability, and whether the identified structural trap has an appropriate seal [see, e.g., Magoon and Dow (1994)]. The overall probability of success is the product of these individual probabilities. Although these assessments may be difficult, for a single prospect, this risk analysis process is straightforward.

When considering multiple prospects in a common basin or multiple target zones in a single well, in addition to considering the probability of success for each prospect, we need to consider the dependence among prospects. For example, if hydrocarbons are found at one location, they may be much more likely to be found at another nearby location. Conversely, if hydrocarbons are not found at the first location, they may be less likely to be found at the other. When evaluating opportunities with multiple prospects, we should consider decision processes and workflows that exploit this dependence and use results from early wells to make more informed decisions about other locations. For example, if a postmortem analysis of core samples from a failed well reveals that there were no hydrocarbons present, then we may not want to continue drilling at nearby sites. On the other hand, if the postmortem analysis reveals that hydrocarbons were present, but the reservoir lacked a seal, then we may want to continue to explore other nearby sites. In this paper, we describe an approach for modeling dependence among prospects and developing a drilling strategy that exploits the information provided by early drilling results.

A Simple Two-Well Example. We can illustrate this problem by considering an example involving two wells. To keep things simple for now, we will assume that the wells simply succeed or fail and we do not obtain postmortem information in the case of a failure. We will assume that Well 1 has a 34.9% probability of being successful and an expected value (net of drilling costs, etc.) of $60 million US if the well is successful and an expected cost of $35 million if the well fails. The overall expected value of drilling this well is 0.349($60)+(1–0.349)(−$35) = −$1.86 million. Thus, this well would not be attractive in isolation. Well 2 has a 48.9% probability of success and an expected value of $15 million if successful, and an expected cost of $20 million if a failure. The overall expected value of Well 2 is −$2.88 million, so it also would not be attractive in isolation.*

Now consider the possibility of drilling Well 1, observing its results, and then deciding whether to drill Well 2; alternatively, we could reverse the order and drill Well 2 first, observe its results, and decide whether to drill Well 1. To evaluate these possibilities, we need to consider what the results from Well 1 tell us about the likelihood of success at Well 2: e.g., if Well 1 succeeds (or fails), what is the probability that Well 2 will succeed? To properly evaluate these possibilities, we need to consider the joint probabilities for the outcomes of both wells. A joint probability distribution for this example is shown in **Table 1.** The entries in the table represent the probabilities of a particular combination of outcomes for the wells. The probabilities shown at the right and bottom of the table are the "marginal" probabilities of success or failure for the individual wells; these are equal to the sum of the row and column entries in the table. These marginal probabilities for the individual wells must total one, and the joint probabilities inside the table must also total one.

There are a number of ways to specify this joint distribution. For example, one might assess the marginal probabilities of success for individual wells. As noted earlier, these are 0.349 and 0.489 for Wells 1 and 2, respectively. We can then complete the rest of the table by specifying one joint or one conditional probability. Here we will specify a conditional probability: suppose that if Well 1 is successful, the probability that Well 2 is also successful is 0.661. Note that our marginal probability that Well 2 is successful is 0.489: thus knowing that Well 1 succeeded leads to higher probability of success for Well 2. This conditional probability assessment implies that the probability that both wells are successful is $p(\text{Well 2 Succeeds}|\text{Well 1 Succeeds}) \times p(\text{Well 1 Succeeds}) = 0.661 \times 0.349 = 0.231$; this appears in the upper left corner of Table 1. The remaining joint probabilities for this two-well example can then be determined using the fact that the table entries must sum to the specified marginal probabilities.

We can then use this probability information in a decision tree to consider the viability of different sequential drilling strategies; see **Fig. 1.** Here the initial decision, represented by the square node at the left, is whether to drill Well 1 or Well 2 first or to not drill either well. If we drill Well 1 or 2 first, we then observe whether it succeeds or fails. This uncertainty is represented by the next layer of (circular) chance nodes. The probabilities for the outcomes of the first well are simply the marginal probabilities in Table 1. These probabilities are shown above the branches corresponding to the outcomes; the values beneath the branches represent the payoffs received when that outcome occurs.

After observing the results of the first well, we then decide whether to drill the other well. To make this second drilling

This paper (SPE 102369) was accepted for presentation at the 2006 SPE Annual Technical Conference and Exhibition, San Antonio, Texas, 24–27 September, and revised for publication. Original manuscript received for review 27 June 2006. Revised manuscript received for review 24 September 2007. Paper peer approved 19 October 2007.

* These are two of the wells we will consider in the five-well example later in the paper. The precise probabilities assumed here are derived in the later example.

TABLE 1—JOINT PROBABILITIES FOR THE TWO-WELL EXAMPLE

		Well 1		
		Successful	Failed	
Well 2	Successful	0.231	0.258	0.489
	Failed	0.118	0.393	0.511
		0.349	0.651	

decision, we need to consider the probabilities for Well 2 (or Well 1) success given the results of the first well. These conditional probabilities may be calculated from Table 1. For example, the probability that Well 2 is successful, given that Well 1 is successful, is equal to the probability that both are successful (0.231) divided by the probability that Well 1 is successful (0.349) and is equal to 0.661. The other conditional probabilities may be calculated in the same way.

We can determine the optimal drilling strategy by working backward through the decision tree. For example, suppose we have drilled Well 1 first and it was successful. The expected value of drilling Well 2 is then $0.661(\$15)+0.339(-\$20)=\$3.14$. (These expected values are shown beneath the chance nodes in the decision nodes in the tree; the expected values beneath to the decision nodes represent the expected value of the most attractive alternative.) Thus, Well 2 is profitable given positive results at Well 1. On the other hand, if Well 1 fails, we revise the probability of success at Well 2 down to 0.397 and the expected value of drilling is −\$6.11, and we would decide to quit. Now that we know the action we would take and the corresponding "continuation value" for each possible Well 1 outcome, we can calculate the expected value of drilling Well 1 first as $0.349(\$60+\$3.14)+0.651(-\$35+\$0)=-\$0.76$. A similar calculation yields an expected value of \$1.91 for the strategy of drilling Well 2 first. Thus, the optimal strategy is to drill Well 2 first and then, if Well 2 is successful, to drill Well 1; if Well 2 fails, then we should quit. Though neither well is attractive in isolation, we see that we can exploit the dependence between the two wells to make the pair of wells attractive.

The General Problem. This simple example illustrates the nature of the problem we study, but it does not demonstrate its scale or complexity. In practice, explorationists will frequently consider plays with more than two prospects and, as indicated earlier, they will typically decompose the assessment of success probabilities into several underlying factors. For example, if we consider five prospects where each well may be either a success or a failure, the joint probability distribution corresponding to Table 1 will have five dimensions with $2^5=32$ possible outcomes whose probabilities must be specified. Many of these probabilities will be difficult to assess. For example, what is the chance that a well at location 5 would be productive, given that Wells 1 and 4 failed and 2 and 3 succeeded? If we have five wells and decompose the individual risk assessments into three underlying geologic factors each of which may succeed or fail, the full joint distribution must consider $(2\times2\times2)^5\approx33{,}000$ different possible outcomes.

If we did somehow manage to specify a joint probability distribution over all of the possible outcomes, we then need to build a decision tree to determine the optimal drilling sequence. The structure of the tree is straightforward—we decide which well to drill first, if any; we then observe the results for that well and decide which well (if any) to drill next and so on, for all possible well outcomes and possible sequences of wells—but there are many possible scenarios to consider. For example with five wells, if we only learn whether a well succeeded or failed, the decision tree would include a total of 9,496 scenarios. If we consider a more detailed model that considers the success or failure of three underlying geologic factors, then each well would have 2^3 different possible outcomes (all possible combinations of success or failure on the three individual factors) and, with five wells, the decision tree would contain approximately 5,000,000 scenarios.

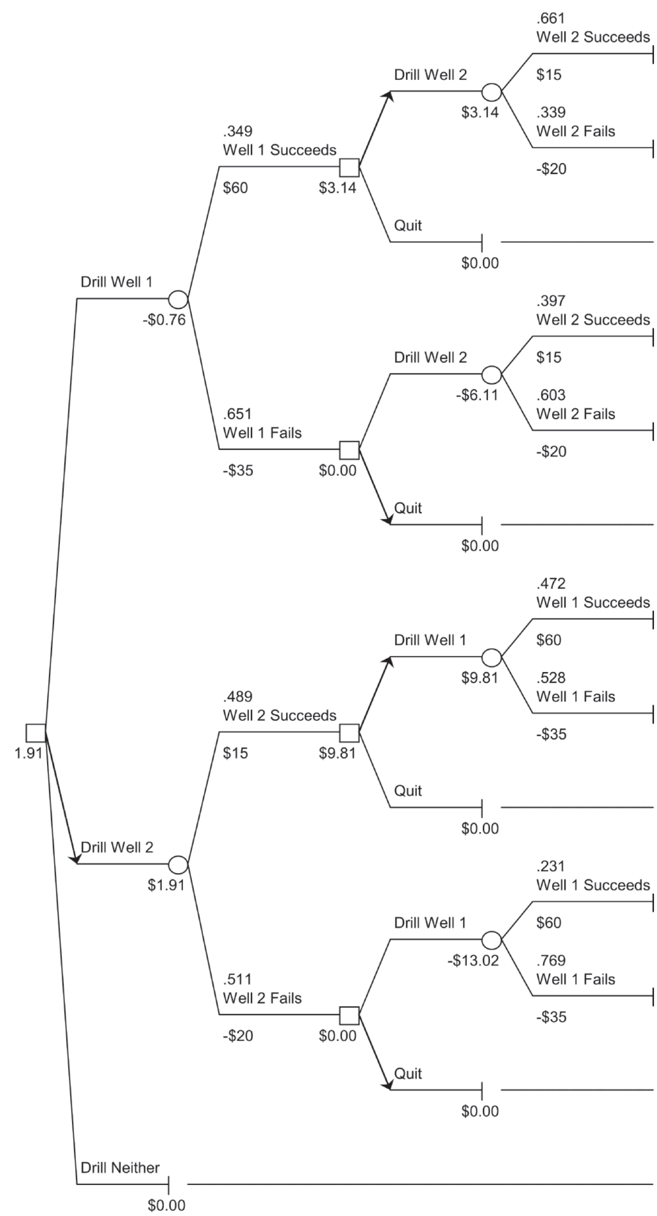

Fig. 1—Decision tree for the two-well example.

Our goal in this paper is to develop a practical approach for modeling dependence among prospects and determining an optimal drilling strategy that exploits the information provided by early drilling results; our interest in this problem stems from a consulting project the first and third authors did for a large oil and gas company. To accomplish this goal we must (1) simplify the assessment of the required joint probabilities, while still capturing important dependencies, and (2) develop a decision model that can efficiently solve for the optimal exploration strategy in situations involving a realistic number of wells. Specifically, we will assume that the explorationists can provide marginal probabilities of success for each well on each factor and pairwise assessments like those in the two-well example, but cannot provide more complex assessments involving the outcomes of three or more wells. We then use techniques from the field of information theory to estimate a complete joint probability distribution that is consistent with these assessments. Finally, we develop a "dynamic programming" model that carries out the same calculations as in the decision tree of Fig. 1, but takes advantage of the fact that different paths through the tree will lead to exactly the same state of information.

The remainder of the paper is organized as follows. After reviewing related literature in the remainder of this section, in the next section, we introduce an example involving five wells and

TABLE 2—WELL DATA FOR THE FIVE-WELL EXAMPLE							
	Probability of Success By Factor				Expected Values		
Well	Charge	Rock	Seal	Overall	Given Success	Given Failure	Intrinsic Value
1	0.73	0.81	0.59	0.349	60	–35	–1.86
2	0.77	0.87	0.73	0.489	15	–20	–2.88
3	0.73	0.83	0.87	0.527	30	–35	–0.74
4	0.65	0.81	0.63	0.332	40	–20	–0.10
5	0.55	0.57	0.57	0.179	80	–20	–2.13

three geologic factors that we will use to demonstrate our proposed approach; although this example is disguised, it captures the essential features of the actual application that motivated this research. In the third section, we describe our technique for constructing a joint probability distribution. In the fourth section, we describe our dynamic programming model for determining an optimal drilling strategy. The final section discusses some possible extensions of this basic model and offers some concluding remarks. An Excel spreadsheet that implements the procedures described in this paper is available free of charge from the authors.*

Literature Review. The problem of modeling dependence in multiprospect drilling programs has long been of interest in the petroleum engineering literature. For example, Newendorp (1975) argued that wells in a common basin are typically dependent and discussed some of the challenges in modeling this dependence. Megill (1984) described how to calculate the probability of geologic success in a multireservoir prospect when the reservoirs are either independent or perfectly correlated. Lerche (1992) used Bayes' theorem to update the probability of an underlying geologic success factor (reservoir fracturing) at one prospect within a basin, given the drilling results at two other prospects.

More recently, Murtha (1996) discussed modeling dependence among features of a prospect with multiple layers, suggesting that one assess the full joint distribution. Stabell (2000), Delfiner (2003), and Keefer (2004) discussed the complexity of this general approach and described frameworks that distinguish between risks that are shared (i.e., perfectly correlated across prospects or targets) and those that are assumed to be independent. Wang et al. (2000) proposed a simple model that assumes all wells are "exchangeable," meaning the wells all have identical probabilities of success and the conditional probabilities for later wells depend on how many wells have succeeded or failed, but not on which wells succeeded or failed. Although these shared risk or exchangeable models may be appropriate in certain settings, these models are fairly restrictive and inappropriate for modeling sequential drilling decisions in general.

Murtha and Petersen (2001) proposed assessing pairwise correlations and use of the "black box" correlation procedures that come with commercial Monte Carlo simulation packages to generate correlated samples. However, as these authors note, in the binary variable setting that we consider, the correlation procedure [Iman and Conover (1982)] that is typically used in these Monte Carlo simulation packages does not actually generate samples that match the specified correlation coefficients. Moreover, the Monte Carlo framework is not well suited for determining optimal drilling strategies. Monte Carlo analyses of multiprospect exploration opportunities typically consider some fixed sequence or strategy. For example, Moore and Mudford (1999) assumed that drilling stops after three failures; Kokolis et. al. (1999) require the user to specify "if-then-else" logic to simulate the decisions that would be made during exploration.

The approach we describe here is based on Bickel and Smith (2006). The information-theoretic approach we use to construct the joint probability distribution can be traced to the seminal paper by Jaynes (1968) and has been used in the decision analysis literature by Smith (1993) and Abbas (2006), among others. Genrich and Sommer (1989) used information-theoretic techniques to study the reduction of uncertainty in reservoir properties within the context of a waterflooding example. Information-theoretic approaches have been shown to lead to reasonable distributions in a variety of contexts and possesses many desirable theoretical properties [see, e.g., Jaynes (1982) or Cover and Thomas (1991)]. Dynamic programming is a standard modeling technique in operations research dating back to Bellman (1957) and used extensively in "real options" analysis and in other contexts. Here, we extend Bickel and Smith (2006) to consider the possibility of learning about the underlying geologic factors; Bickel and Smith consider the success or failure of a well without considering underlying geologic factors. Our discussion here is self-contained though the interested reader will be referred to Bickel and Smith for some more technical discussions and extensions, as well as a study of the accuracy of the approach.

A Five-Well Example

Suppose we have obtained a 3D seismic survey and are considering drilling wells at five locations that are known to have suitable traps. The probability of success at each location is decomposed into the assessment of three independent geologic factors: hydrocarbon charge (C), reservoir rock (R), and seal (S). We will assume that these three geologic factors are independent, so that the probability of overall geologic success at Well i is given by the product of probabilities for these individual factors. **Table 2** shows our assumed probabilities for each location. The expected values of the well given geologic success or failure are also shown there; these expected values are in millions of dollars and represent the net present value (NPV) of a successful or failed well. The intrinsic values shown in Table 2 are the unconditional expected values for each prospect: For example, for Well 1, the probability of geologic success is $0.73 \times 0.81 \times 0.59 = 0.349$ and the intrinsic value is $0.349 \times (\$60) + (1\text{-}0.349) \times (-\$35) = -\$1.86$ million. This intrinsic value represents the value of the well if it were considered in isolation of the other wells. In this example, these intrinsic values are all negative—none of these prospects would be attractive by itself. (Note that the assumptions for Wells 1 and 2 here match the assumptions in the simple two-well example considered in the introduction.)

To model the dependence among these five prospects, we must specify a joint probability distribution over the $2^5 = 32$ possible outcomes for each geologic factor. Though it would likely be difficult to assess all of these probabilities, it is not too difficult to assess pairwise conditional probabilities for each factor. For example, one might assess the probability of finding a charge at prospect i given a charge at prospect j for each pair of prospects. There are a total of 10 such pairs of assessments for each factor. **Table 3** shows a complete set of such assessments for our example. Alternatively, one could assess pairwise correlation coefficients and then calculate the required conditional or joint probabilities. The correlation coefficients corresponding to the condi-

* The spreadsheet is available at both the first and second authors' university websites.

TABLE 3—CONDITIONAL PROBABILITIES FOR THE FIVE-WELL EXAMPLE

Charge						
	Conditional Assessments $p(C_j=1 \mid C_i=1)$					Marginal
i / j	1	2	3	4	5	$p(C_i)$
1		0.80	0.77	0.70	0.60	0.73
2			0.78	0.72	0.61	0.77
3				0.68	0.61	0.73
4					0.65	0.65
5						0.55
Rock						
	Conditional Assessments $p(R_j=1 \mid R_i=1)$					Marginal
i / j	1	2	3	4	5	$p(R_i)$
1		0.95	0.90	0.92	0.69	0.81
2			0.90	0.91	0.65	0.87
3				0.90	0.65	0.83
4					0.70	0.81
5						0.57
Seal						
	Conditional Assessments $p(S_j=1 \mid S_i=1)$					Marginal
i / j	1	2	3	4	5	$p(S_i)$
1		0.87	0.94	0.80	0.80	0.59
2			0.95	0.80	0.76	0.73
3				0.72	0.65	0.87
4					0.85	0.63
5						0.57

TABLE 4—PAIRWISE CORRELATIONS FOR THE FIVE-WELL EXAMPLE

Charge					
	Pairwise Correlation Coefficient				
i / j	1	2	3	4	5
1		0.12	0.15	0.17	0.17
2			0.21	0.27	0.22
3				0.10	0.20
4					0.27
5					
Rock					
	Pairwise Correlation Coefficient				
i / j	1	2	3	4	5
1		0.49	0.38	0.58	0.50
2			0.48	0.66	0.42
3				0.51	0.36
4					0.54
5					
Seal					
	Pairwise Correlation Coefficient				
i / j	1	2	3	4	5
1		0.38	0.25	0.42	0.56
2			0.39	0.58	0.63
3				0.48	0.42
4					0.74
5					

tional probabilities of Table 3 are shown in **Table 4.*** Generally, given n prospects, $n(n-1)/2$ pairwise assessments of correlations or conditional probabilities will be required. (We will discuss the possibility of omitting some of these assessments later in the paper.)

The pairwise conditional probabilities of Table 3 or correlation coefficients of Table 4, together with the marginal probabilities given in Table 2, are sufficient to specify the joint probability distribution for the possible outcomes for any pair of prospects, but are not sufficient to specify the full joint distribution for all combinations of outcomes for all five locations. In the next section, we describe an approach for constructing a complete joint probability distribution based on these marginal and pairwise probability assessments. We then describe how to use this joint probability distribution to determine an optimal drilling strategy.

Constructing a Joint Probability Distribution

We will construct a joint probability distribution by making the well results as close as possible to independent while respecting the given marginal and conditional assessments. By choosing a joint distribution to minimize dependence in this way, we are being conservative about what we assume about how much information each well provides about the other prospects. We will measure how close we are to independence by considering the relative entropy or Kullback-Leibler (KL) distance between the constructed joint distribution and the joint distribution one would obtain if the wells were assumed to be independent. More specifically, our goal is to construct a joint distribution according to the following criteria:

- Minimize the KL distance between our distribution and the independent joint distribution subject to the constraints that the joint probabilities:
  - Sum to one.
  - Match the specified marginal probabilities for each well.
  - Match the specified pairwise probabilities.

Because we have assumed independence among the geologic factors, we will construct three separate joint distributions (one for each factor) using this information-theoretic approach. The full joint distribution is then given by the product of these three factor-specific joint distributions.

Method. To describe our approach more formally, we need to introduce some notation; we will adopt a generic notation that is applicable to each geologic factor. We let w_i be a binary variable such that $w_i = 1$ if Well i succeeds on this factor (e.g., the geologic feature is present) and $w_i = 0$ otherwise. We let $\boldsymbol{w} = (w_1, \ldots, w_n)$ be a vector of n binary random variables. For instance, in our five-well example $\boldsymbol{w} = (1,0,0,1,1)$ would denote a scenario where Wells 1, 4, and 5 succeeded on this factor and Wells 2 and 3 failed. Our goal in this section will be to construct a joint probability distribution $\pi(\boldsymbol{w})$ over the set of all possible values of $\boldsymbol{w}$.*

We will assume that we are given the n marginal probabilities $p_i \equiv p(w_i = 1)$ and the $n(n-1)/2$ pairwise joint probabilities $p_{ij} \equiv p(w_i = 1, w_j = 1)$. These pairwise joint probabilities can be calculated from the marginal and conditional probabilities shown in Tables 2 and 3. For example, the probability of the charge being present at both locations 1 and 2 is given by

* The correlation between binary random variables A and B is

$$\rho_{AB} = \frac{p(A)(p(B|A) - p(B))}{\sqrt{p(A)(1 - p(A))p(B)(1 - p(B))}}$$

where p(A) is the probability of success at A and p(B| A) is the probability of success at B given success at A.

* To streamline our notation, we will not explicitly list the ranges for $\boldsymbol{w}$ in our summations below: the vector $\boldsymbol{w}$ will range over the 2^n possible combinations of outcomes of the n propsects. Similarly, the marginal probabilities p_i and corresponding Lagrange multipliers (introduced later) will range from i=1, . . . , n, i.e., over the prospects. Similarly, the pairwise joint probabilities p_{ij} and corresponding Lagrange multipliers will range over the $n(n-1)/2$ unique pairs of prospects, which may be indexed as i=1, . . . , $n-1$ and j=i+1, . . . , n.

$p(C_2|C_1)\times p(C_1)=0.80\times 0.73=0.584$. We will assume these assessments are consistent in that $0<p_i<1$ and $0<p_{ij}<p_i$, but make no other specific assumptions about the specified probabilities.

The KL distance between π and the independent joint probability distribution $\pi_0(\boldsymbol{w})$ is formally defined as:

$$KL(\pi,\pi_0)\equiv\sum_{\boldsymbol{w}}\pi(\boldsymbol{w})\ln\left(\frac{\pi(\boldsymbol{w})}{\pi_0(\boldsymbol{w})}\right), \quad (1)$$

where:

$$\pi_0(\boldsymbol{w})\equiv\prod_i (p_i)^{w_i}(1-p_i)^{1-w_i}. \quad (2)$$

If the binary events are independent under π (i.e., $\pi=\pi_0$), we will have $KL(\pi,\pi_0)=0$; otherwise $KL(\pi,\pi_0)>0$. [For more on KL and information theory, see, e.g., Cover and Thomas (1991)].

In the two-well example using the probabilities of Table 1, Eq. 1 becomes:

$$\begin{aligned}KL(\pi,\pi_0)&=0.231\ln\left(\frac{0.231}{0.489\times 0.349}\right)+0.258\ln\left(\frac{0.258}{0.489\times 0.651}\right)\\&\quad+0.118\ln\left(\frac{0.118}{0.511\times 0.349}\right)+0.393\ln\left(\frac{0.393}{0.511\times 0.651}\right)\\&=0.032.\end{aligned}$$

Here the $\boldsymbol{w}$s range over the cells inside Table 1 and the $\pi(\boldsymbol{w})$ are the table entries. The independent joint distribution $\pi_0(\boldsymbol{w})$ (appearing in the denominator above) is given by multiplying the row and column marginal probabilities for each cell.

Our goal then is to choose a joint distribution π to minimize the KL distance $KL(\pi,\pi_0)$ subject to the constraint of matching the specified marginal and pairwise joint probabilities. To formalize this optimization problem, we define the following probability constraint functions:

$$\Omega_0(\boldsymbol{w})=1$$
$$\Omega_i(\boldsymbol{w})=w_i \qquad \text{for all } i$$
$$\Omega_{ij}(\boldsymbol{w})=w_iw_j \qquad \text{for all } i, j.$$

We let $E_\pi[f(\boldsymbol{w})]$ denote the expected value of a function $f(\boldsymbol{w})$ when $\boldsymbol{w}$ has distribution π. Using this notation, the constraint that probabilities must sum to one, for example, is written as $E_\pi[\Omega_0(\boldsymbol{w})]=1$ and we can formulate our optimization problem as

$$\min_\pi\sum_{\boldsymbol{w}}\pi(\boldsymbol{w})\ln\left(\frac{\pi(\boldsymbol{w})}{\pi_0(\boldsymbol{w})}\right), \quad (3)$$

subject to: $E_\pi[\Omega_0(\boldsymbol{w})]=1$

$$E_\pi[\Omega_i(\boldsymbol{w})]=p_i \qquad \text{for all } i$$
$$E_\pi[\Omega_{ij}(\boldsymbol{w})]=p_{ij} \qquad \text{for all } i,j.$$

These constraints ensure that the probabilities sum to one, and match the specified marginal and pairwise probabilities. For instance, with the probabilities in the two-well example, the second constraint given previously is:

$$\pi(1,1)\times 1+\pi(1,0)\times 1+\pi(0,1)\times 0+\pi(0,0)\times 0=p_1=0.349,$$

$$\pi(1,1)\times 1+\pi(1,0)\times 1+\pi(0,1)\times 1+\pi(0,0)\times 0=p_2=0.489.$$

In Table 1, this corresponds to requiring the sum of the probabilities in the "successful" row and column to match the marginal probabilities shown on the right and bottom, respectively.

Bickel and Smith (2006) show that this optimization problem leads to a optimal joint probability distribution of the form:

$$\pi^*(\boldsymbol{w},\boldsymbol{\lambda})=\pi_0(\boldsymbol{w})\exp\left(-1+\lambda_0+\sum_i\lambda_i\Omega_i(\boldsymbol{w})+\sum_{i,j}\lambda_{ij}\Omega_{ij}(\boldsymbol{w})\right), \quad (4)$$

where λ_0, λ_i, and λ_{ij} denote the "Lagrange multipliers" associated with the constraints in Eq. 3; we let $\boldsymbol{\lambda}$ denote the vector of these Lagrange multipliers. These Lagrange multipliers may be found by solving a simpler optimization problem:

$$\max_\lambda\left(-\sum_{\boldsymbol{w}}\pi^*(\boldsymbol{w},\boldsymbol{\lambda})+\lambda_0+\sum_i\lambda_ip_i+\sum_{i,j}\lambda_{ij}p_{ij}\right). \quad (5)$$

This optimization problem has $1+n+n(n-1)/2$ decision variables (the Lagrange multipliers), has no constraints, is concave in $\boldsymbol{\lambda}$, and is straightforward to solve numerically. In our example with $n=5$ wells, this optimization problem can be solved in one or two seconds using Solver® in Excel.

The distribution $\pi^*(\boldsymbol{w},\boldsymbol{\lambda})$ in Eq. 4 has a nice intuitive structure: the exponential term in Eq. 4 multiplies, or "boosts," the independent probabilities ($\pi_0(\boldsymbol{w})$) based on how the other wells turn out. Such multiplicative "boosts" to the independent probabilities are commonly proposed as a simple way to adjust probabilities to take into success or failure at nearby wells; Murtha (1996) and Delfiner (2003) discuss some of the problems with simply using a constant multiplier (e.g., 2 or 1.25) to take into account nearby successes. Here, our process of calculating the optimal Lagrange multipliers ensures that the resulting distribution is consistent and avoids these difficulties.

Letting $\boldsymbol{w}_{-i}$ denote the vector of outcomes for all wells other than Well i, from Eq. 4 we find that the conditional log-odds for success at well i can be written as

$$\ln\left(\frac{p(w_i=1|\boldsymbol{w}_{-i})}{1-p(w_i=1|\boldsymbol{w}_{-i})}\right)=\ln\left(\frac{p_i}{1-p_i}\right)+\lambda_i+\sum_{j\neq i}\lambda_{ij}w_j. \quad (6)$$

Here, $\ln(p_i/(1-p_i))$ is the prior log-odds for Well i, i.e., the logarithm of the odds of success when we do not know the outcome of any other well and $\ln(p(w_i=1|\boldsymbol{w}_{-i})/(1-p(w_i=1|\boldsymbol{w}_{-i})))$ are the log-odds for well i conditional on the outcome of all the other wells. Eq. 6 shows that learning the outcome of other wells has a linear effect on these log-odds. The marginal Lagrange multipliers λ_i (which are typically negative) describe the adjustment to the prior log-odds in the scenario that all other wells fail. The joint Lagrange multipliers λ_{ij} describe the increase in the log-odds of success for Well i because of success at Well j. These joint Lagrange multipliers λ_{ij} should be positive if success at one well increases the probability of success at the others.

To summarize, the workflow associated with our method for constructing a joint probability distribution is as follows. For each factor:

(1) Specify marginal probabilities p_i describing the probability of succeeding on this factor at Well i. In the example, these are shown in Table 2 and repeated in the right margin in Table 3.

(2) Specify pairwise joint probabilities p_{ij} describing the probability of success on this factor at both prospects i and j. These probabilities can be calculated using the marginal probabilities p_i with either the conditional probabilities (e.g., those shown in Table 3) or the pairwise correlation coefficients (shown in Table 4).

(3) Solve the optimization problem (Eq. 5) to determine the optimal Lagrange multipliers. For problems of reasonable size (up to, say, 15 wells), this can be done using standard spreadsheet-based optimization packages such as Excel's Solver.

(4) Given these Lagrange multipliers, we have fully specified the joint probability distribution $\pi^*(\boldsymbol{w},\boldsymbol{\lambda})$ of Eq. 4 and can calculate the probability of any combination of outcomes for all prospects $\boldsymbol{w}$.

We illustrate the results of this procedure in our five-well example and then describe how to use these joint probability distributions to determine an optimal drilling sequence. These calculations are available in the authors' previously mentioned spreadsheet.

Example Results. Table 5 shows the Lagrange multipliers associated with our five-well example. Notice that the pairwise Lagrange multipliers are all positive. Following Eq. 6, this implies that the outcomes of all the individual wells are positively related, meaning success (failure) on a particular geologic factor at one well increases the probability of success (failure) at the other locations.

TABLE 5—LAGRANGE MULTIPLIERS FOR THE FIVE-WELL EXAMPLE

Charge

	Lagrange Multiplier λ_{ij}					Marginal
i / j	1	2	3	4	5	λ_i
1		0.20	0.53	0.57	0.48	–1.11
2			0.80	1.05	0.66	–1.60
3				0.01	0.69	–1.31
4					0.95	–1.68
5						–1.98
					λ_o	3.32

Rock

	Lagrange Multiplier λ_{ij}					Marginal
i / j	1	2	3	4	5	λ_i
1		0.80	0.44	1.39	2.60	–2.70
2			1.22	2.49	1.76	–3.12
3				1.34	0.85	–2.52
4					3.61	–4.74
5						–7.92
					λ_o	6.17

Seal

	Lagrange Multiplier λ_{ij}					Marginal
i / j	1	2	3	4	5	λ_i
1		0.23	0.09	0.05	2.36	–1.51
2			0.62	1.07	2.97	–2.13
3				3.22	1.50	–1.58
4					3.15	–5.16
5						–7.14
					λ_o	4.42

Table 6 shows an example of a conditional probability calculation for Well 5 in the scenario in which Wells 1 and 2 succeeded on rock (R) and seal (S) but failed on charge (C); Well 3 succeeded on C an R but failed on S, and Well 4 succeeded on all three factors. Using Eq. 6 and the Lagrange multipliers from Table 5, we find the posterior odds and probabilities shown on the right side of Table 6. For example, for Charge shown in the first row of the table, the posterior log odds that Well 5 succeeds on this factor are given by Eq. 6 as:

$$\ln\left(\frac{0.58}{1-0.58}\right) + -1.98 + 0.48 \times 0 + 0.66 \times 0 + 0.69 \times 1 + 0.95 \times 1 = -0.14,$$

which implies that, given these two negative results and two positive results at the other wells, the posterior probability that Well 5 succeeds on Charge is 0.47, somewhat lower than the original (prior) probability of 0.55.

Comparing the initial and final probabilities for the other factors in Table 6, we see that these well results lead us to revise the probability of success on Rock and Seal upwards from 0.57 (for both Rock and Seal) to 0.77 and 0.84, respectively. The overall probability of geologic success, given as the product of these factor probabilities, increases from 0.18 to 0.30. Note that in this example, the probability of success for Well 5 has increased despite having failed at three of the four earlier wells! In this case, the impact of these failures on Well 5's probability of success is mitigated because the wells failed for different reasons and we succeeded fully on the well (Well 4) with the strongest relationship with Well 5: for each factor, the Lagrange multiplier λ_{45} is larger than the other Lagrange multipliers λ_{i5} involving Well 5.

Comparing Tables 4 and 5, we see that pairs of locations that have larger correlation coefficients tend to have larger Lagrange multipliers. However, the relationship between correlations and Lagrange multipliers is not perfect. For example, with Rock, the correlation between Wells 2 and 4 is 0.66 and between Wells 4 and 5 is 0.54, yet the corresponding Lagrange multipliers are 2.49 and 3.61, respectively. These differences should not be surprising because the two measures of pairwise dependence are different. In fact, as shown in Bickel and Smith (2006), it is quite possible to have positive pairwise correlation coefficients between two wells and a negative Lagrange multiplier for this pair.

In some cases, one may wish to omit certain pairwise assessments because of lack of time or resources for providing assessments. For example, Murtha and Peterson (2001) suggest designating a "key prospect" and assessing only the pairwise correlation coefficients involving this prospect. In their approach, the other correlations are implicitly chosen by the Monte Carlo simulation software. We can similarly omit correlations or conditional probability assessments by dropping the constraints in minimization problem (Eq. 3). If we omit a pairwise probability constraint, the optimal distribution (Eq. 4) assigns $\lambda_{ij} = 0$ to these omitted assessments. This implies that the two outcomes are conditionally independent given the outcomes of the other wells, but need not (and in general will not) imply that the pairwise correlation coefficient ρ_{ij} is zero. Instead of assuming conditional independence as a "default," one might instead use the log-odds interpretation of Eq. 6 to specify default multipliers directly or develop a statistical model that relates these parameters to the distance between wells or some other measure of similarity between prospects. Such a model might be analogous to kriging models that are commonly used in geostatistics to relate covariances among geologic attributes at different locations to the distances between locations [see, e.g., Goovaerts (1997)].

Determining the Optimal Exploration Strategy

Now suppose we have specified a joint probability distribution for each geologic factor. How do we determine the optimal drilling strategy? As discussed in the introduction, a standard decision tree model of the problem would be quite complex, even with moderate numbers of wells. We can, however, simplify the model by recognizing that many different early paths lead to the same state of information and future cash flows. For example, if we have drilled Wells 1 and 3, and Well 3 was successful but Well 1 failed because it had no seal, the future probabilities and cash flows are the same regardless of whether we drilled Well 1 or Well 3 first. In this section, we will describe a dynamic-programming-based solution procedure that exploits this recombining structure.

TABLE 6—CONDITIONAL PROBABILITIES FOR WELL 5 IN ONE SCENARIO

	Prior		Well Results				Lagrange Calcs		Posterior	
Well Factor	Prob	ln(Odds)	1	2	3	4	λ_5	$\Sigma\lambda_{5j} \times w_j$	ln(Odds)	Prob
Charge	0.55	0.20	0	0	1	1	–1.98	1.64	–0.14	0.47
Rock	0.57	0.28	1	1	1	1	–7.92	8.82	1.19	0.77
Seal	0.57	0.28	1	1	0	1	–7.14	8.48	1.62	0.84
Overall	0.18									0.30

The logic of the dynamic programming model is exactly the same as that in the decision tree model of Fig. 1: we determine the optimal drilling strategy by first considering scenarios where all but one well has been drilled; these are at the right side of the tree. To figure out what to do in these scenarios, we need to determine the probabilities for the outcomes of that well conditioned on the outcomes of the previous wells. We then work backward through the tree to determine the optimal actions in earlier states. The conditional probabilities that describe the probability of moving from one state to another when drilling a well are referred to as transition probabilities.

To describe this procedure precisely, we introduce some notation to denote the possible scenarios or states of information that may be encountered. We focus on the case with five wells and three geologic factors, labeling the factors as the *x*, *y*, and *z* factors. We let the vectors $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{z}$ describe the state of each of these three factors across the 5 wells. Focusing on factor *x*, $\mathbf{x}=(x_1, \ldots, x_i, \ldots, x_5)$ where $x_i=0$, 1 or, "−" depending on whether the first factor fails (0), succeeds (1) or "−" if the well has not been drilled and the result has not yet been observed. For example, the initial state is $\mathbf{x}=(\text{-}, \text{-}, \text{-}, \text{-}, \text{-})$ when no wells have been drilled. The vector $\mathbf{x}=(0, \text{-}, 1, \text{-}, \text{-})$ represents the state in which Wells 1 and 3 have been drilled and the *x*-factor was present at Well 3 and absent at Well 1. The state vectors for the other factors are defined in the same way.

To calculate these transition probabilities, it is helpful to first define a total probability function. Let π_x, π_y, and π_z denote the joint probability distribution for the three geologic factors. (These joint probability distributions may be constructed using the method from the previous section or another method.) The total probability $\mu_x(\mathbf{x})$ associated with the vector $\mathbf{x}$ is constructed by summing $\pi_x(\mathbf{x})$ over the possible scenarios for the unknown outcomes. For example in the case of five wells, for $\mathbf{x}=(0, \text{-}, 1, \text{-}, \text{-})$,

$$\mu_x(0, \text{-}, 1, \text{-}, \text{-}) = \sum_{x_2, x_4, x_5} \pi_x(0, x_2, 1, x_4, x_5), \quad \ldots\ldots\ldots\ldots\ldots\ldots (7)$$

where x_2, x_4, and x_5 range over {0,1}; this summation would involve 2^3 probabilities. This is the probability that that the first well would fail and third well succeed on factor *x* (if both wells are drilled). The total probability functions μ_y and μ_z for the other factors are defined in the same way.

The transition probabilities required for the dynamic programming model can then be computed from this total probability function. Suppose that we start in a state where well *i* has not been drilled (thus x_i, y_i, and z_i are all equal to "−"). If we drill well *i*, the probability that well *i* succeeds on the *x*-factor is equal to $\mu_x(\mathbf{x}_i^1)/\mu_x(\mathbf{x})$ where $\mathbf{x}_i^1$ is identical to $\mathbf{x}$ except $\mathbf{x}_i=1$. The probability that Well *i* fails on factor *x* is $\mu_x(\mathbf{x}_i^0)/\mu_x(\mathbf{x})$ where $\mathbf{x}_i^0$ is identical to $\mathbf{x}$ except $\mathbf{x}_i=0$. This calculation is exactly analogous to the calculations in the two-well example, where the conditional probabilities were given as a joint probability divided by a marginal probability.

The dynamic programming model can now be formalized as follows: Let $v(\mathbf{x}, \mathbf{y}, \mathbf{z})$ denote the expected NPV of future cash flows (the continuation value) given that we start in state $(\mathbf{x}, \mathbf{y}, \mathbf{z})$. In this value calculation, we include the expected value for a successful or failed well when the well is drilled, and discount future cash flows using a discount factor δ that corresponds to the time required to drill a well.

The expected NPVs and optimal strategies are calculated recursively, in the same way as in the decision tree of Fig. 1. If all of the wells have been drilled (i.e., $\mathbf{x}$, $\mathbf{y}$, and $\mathbf{z}$ are all vectors of zeros and ones), then there are no future cash flows and $v(\mathbf{x}, \mathbf{y}, \mathbf{z})=0$. For earlier states, the expected NPV associated with drilling Well *i* is given by the expected sum of the reward associated with Well *i* and the (discounted) continuation value associated with starting the next stage in the randomly determined next period state:

$$v_i(\mathbf{x}, \mathbf{y}, \mathbf{z}) = \sum_{j=0}^{1}\sum_{k=0}^{1}\sum_{l=0}^{1}\left(\frac{\mu_x(\mathbf{x}_i^j)\mu_y(\mathbf{y}_i^k)\mu_z(\mathbf{z}_i^l)}{\mu_x(\mathbf{x})\mu_y(\mathbf{y})\mu_z(\mathbf{z})}\right)\big(r_i(j, k, l) + \delta v(\mathbf{x}_i^j, \mathbf{y}_i^k, \mathbf{z}_i^l)\big), \quad \ldots\ldots\ldots\ldots\ldots\ldots (8)$$

where *j*, *k*, and *l* range over {0,1} and denote the outcomes of the three geologic factors. The $\mu_x(\mathbf{x}_i^j)/\mu_x(\mathbf{x})$ term represents the probability of transitioning from $\mathbf{x}$ to state $\mathbf{x}_i^j$ (similarly for the other factors) and $r_i(j, k, l)$ describes the reward associated with a particular outcome of that well. In our example, we will take $r_i(1, 1, 1)$ to denote the value of a successful well shown in Table 2. For other (*j*, *k*, *l*), we will take $r_i(j, k, l)$ to be the value of a failed well shown in Table 2. The optimal action in state $(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is to drill the well with the largest $v_i(\mathbf{x}, \mathbf{y}, \mathbf{z})$ or, if no well has a positive value, to not drill at all. The optimal value $v(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is thus $\max\{v_i(\mathbf{x}, \mathbf{y}, \mathbf{z}), 0\}$ where the maximum is taken over all available wells and not drilling (0).

To summarize, given joint probability distributions for each factor (π_x, π_y, and π_z), the workflow for determining an optimal drilling strategy requires two steps:

(1) Construct the total probability functions for each factor (μ_x, μ_y, and μ_z) using formulas of the form of Eq. 7.

(2) Determine the optimal values ($v(\mathbf{x}, \mathbf{y}, \mathbf{z})$) and optimal strategy using the recursive procedure outlined previously.

In our spreadsheet implementation, the total probabilities are calculated using matrix multiplication and the dynamic programming recursion is implemented using Visual Basic for Applications (VBA) in Microsoft Excel. Our five-well example takes approximately 5 seconds to evaluate on a PC.

Example Results. In our example, we assume a discount rate of 1% per stage, which implies $\delta=1/(1+0.01)\approx 0.99$. This discount rate would correspond to a rate of approximately 9% per year if the wells each take 6 weeks to drill.

Fig. 2 shows the initial stages of the optimal drilling strategy for this example. The expected value of this drilling strategy is $21.17 MM and we should begin by drilling Well 2. If we drill Well 2, there is a 48.9% chance that this well will be successful. If Well 2 is successful, we should then drill Well 3; the value starting in this state (with Well 2 drilled and successful) would then be $46.83MM. With the success of Well 2, Well 3 then has a 66.7% chance of success, compared to the prior probability of 53% (shown in Table 2). If Well 3 succeeds, we should continue on to Well 4.

The drilling strategies in the unsuccessful cases depend on why the well failed. If the first well drilled (Well 2) failed because it

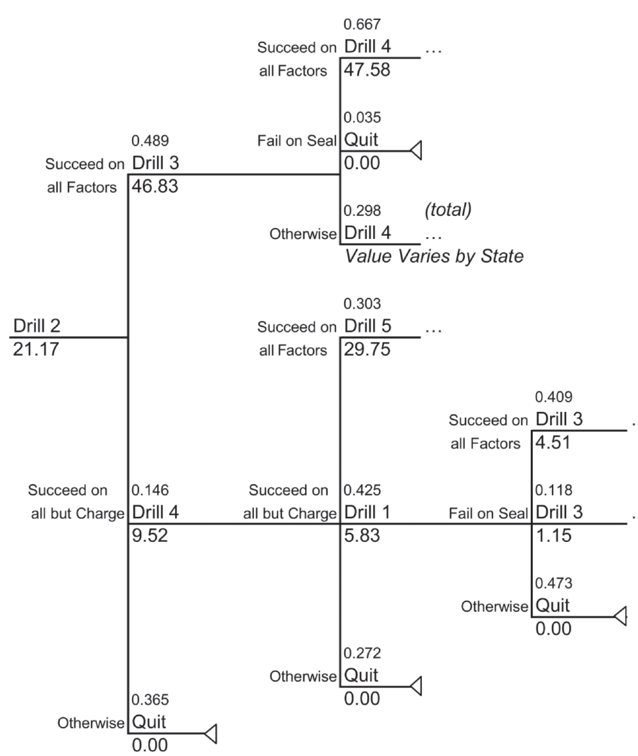

Fig. 2—The optimal drilling strategy.

had no charge, but the drilling results show that the location had good rock and a good seal, then we should continue by drilling Well 4 next. If, however, the rock or seal failed at Well 2, we should stop drilling after this first failed well. The later drilling decisions similarly depend on the results for the individual geologic factors. As shown in Fig. 2, there are some scenarios where we should quit after failing on one well, and others where we should continue to drill a fourth well despite having failed on our first three wells. In total, this optimal strategy leads to a 46% chance of drilling four or more wells and a 34% chance of drilling all five wells.

As discussed before in Eq. 8, the continuation values shown in Fig. 2 include current and future payoffs, but not past ones. For example, the continuation value of $46.83MM if Well 2 succeeds (the top branch of Fig. 2) includes the expected value of Well 3 and any future wells, but does not include the value of the successful Well 2. The initial value $21.17MM is given by Eq. 8 in terms of these continuation values as $0.489\times(\$15+\delta\$46.83)+0.146\times(-\$20+\delta\$9.52)+0.365\times(-\$20+\delta\$0)$, where $15 and −$20 are the rewards (r_i) associated with success and failure, respectively, at Well 2 and $\delta = 1/(1+0.01)$ is the discount factor. To determine that it is optimal to drill Well 2 first, we must compare this value to those given by similar calculations starting with the other wells.

As is evident in Fig. 2, the drilling strategies that exploit dependence and the information about geologic factors can be quite complicated. What is the benefit of pursuing such a complicated strategy? First, recall that all of the wells in this example have negative intrinsic values and therefore would not be attractive if considered in isolation. Here, the optimal dynamic strategy that exploits dependence leads to an expected value of over $21MM. Thus, by taking advantage of this information, we can make five prospects that are individually unattractive into an attractive multiprospect play!

Could we do as well by following some simpler strategy? Although we have not explored all "simple" strategies, the optimal policy in this example confounds many plausible rules for sequencing prospects. For example, it is not optimal to drill the well with the largest intrinsic value first; that would be Well 4 in this example. Nor it is optimal to drill the well with the highest probability of geologic success first (Well 3) or the cheapest first (Wells 2, 4, and 5 all have the same cost). Nor is it optimal to follow a rule that calls for quitting after some fixed number of failed wells: here we sometimes quit after one failure and other times continue despite having failed three times in three attempts.

Such simple heuristic strategies can capture some of the benefit of dependence, but perform worse than the optimal strategies. For example, we might consider a heuristic strategy of drilling the wells in order of their probability of success (well order: 3-2-1-4-5) and stopping whenever there have been two failures. If we can sample from a joint probability distribution that captures the dependence among prospects, we can evaluate such a heuristic strategy using Monte Carlo simulation with "if-then-else" logic for stopping. With this particular strategy, we would have an expected NPV of $11.71MM, compared $21.17MM for the optimal strategy. If we used the same drilling order, but instead always stop after the first or third failure we would have expected NPVs of $11.35MM or $4.11MM, respectively. Thus, these simple heuristic strategies can capture some of the potential benefit associated with dependence among the prospects, but are much less effective than the optimal strategy given by our dynamic programming approach.

A simpler strategy that is interesting to consider is an optimal dynamic strategy that considers whether a well succeeds or fails, but does not consider which geologic factors succeeded or failed. The recursive structure of the solution in this case is similar to that of Eq. 8 but simpler: rather than having to consider 9^5 different possible well states (each of the five wells could succeed or fail on three factors or not have been drilled yet, for a total of nine states per well), in this simpler model, we need only consider 3^5 states (each well succeeds or fails or has not been drilled). The difference in results between these two models highlights the value of postmortem analysis learning about the geologic factors and identifying why each well failed.

Applying this simpler dynamic programming model in our example, we find the optimal strategy shown in **Fig. 3.*** Comparing this strategy to that in Fig. 2, we see that, as before, we begin by drilling Well 2. However, the value of the optimal strategy is now $18.32MM, approximately $3MM less than before. Thus, the value of learning about the underlying geologic factors is worth approximately $3MM in this example.

The increase in value associated with the more detailed information about the geologic factors comes from following a different drilling strategy. Without the detailed information about the geologic factors, if we succeed on Well 2, we would drill Well 4 next. If we have more detailed information in this case, it is slightly better to drill Well 3 next instead (the value is $46.83MM for drilling well 3 vs. $46.62MM for drilling Well 4). The more significant difference is what happens if we fail on Well 2. Without detailed information about the individual geologic factors, it is optimal to quit in this scenario. With more detailed information, if Well 2 fails only because of the lack of a hydrocarbon charge, then it is optimal to drill Well 4 rather than quit. Intuitively, because the presence of the hydrocarbon charge was assumed to be less correlated between prospects (see Table 4), failing for this reason is less damning than other causes of failure. There is a 15% chance of failing for this reason and the difference in expected values in this scenario ($9.52MM for continuing and $0 for quitting) is quite significant.

Fig. 4 compares the cumulative probability distributions for the NPV given by following the optimal strategies with and without learning about the geologic factors. The probabilities associated with the good outcomes (on the far right of the cumulative probability distribution) are quite similar. In these scenarios, the initial wells succeed and many (or all) of the subsequent wells also succeed and the detailed information about which factors failed has relatively little impact. The main differences in the strategies are in middle and negative ranges of outcomes. Without learning about the factors, the probability of having a negative NPV is approximately 70% and the worst possible outcome is a loss of $40MM which occurs if we succeed on Well 2 and then fail on both Wells 4 and 3. With learning about factors, the probability of a negative NPV is reduced to approximately 60%, but we could lose a total of $110MM in the very unlikely (probability≈0.3%) scenario in which we drill Wells 2, 4, 1, and 3 and fail on all of them. The larger potential losses associated with following the "factor learning" strategy are more than offset in expected value terms by higher probabilities of positive outcomes. In total, without factor learning, we have an expected NPV of $18.32MM with a standard deviation of approximately $73MM. With factor learning, we have

* The joint distribution used in this simplified dynamic programming model was calculated from the joint distribution used in the more complex model that considered individual factors. The difference in results is thus entirely due to having more or less detailed information about why the wells failed.

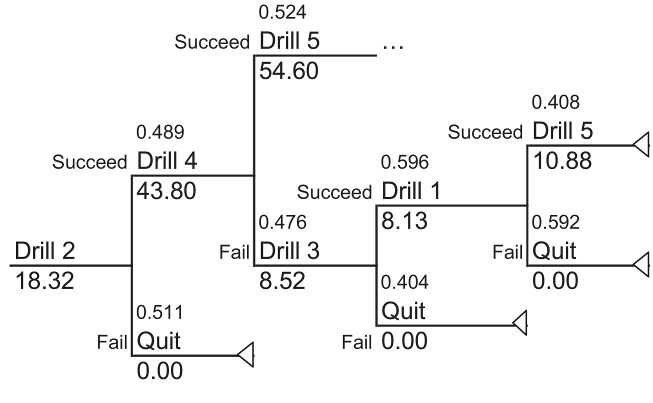

Fig. 3—Optimal strategy without factor learning (partial).

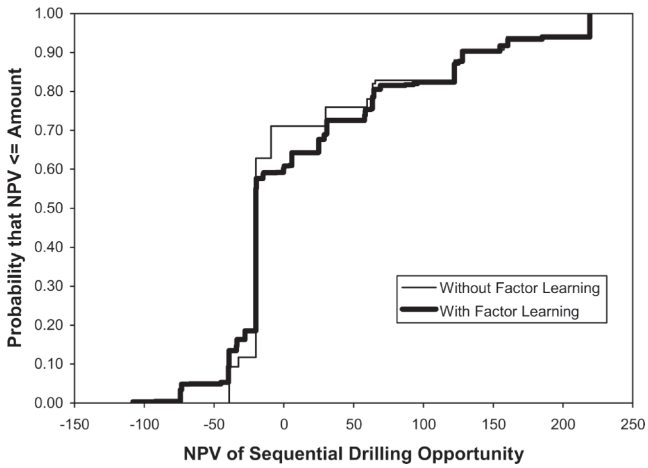

Fig. 4—Comparison of risks with and without factor learning.

a higher expected NPV of \$21.17MM and a slightly higher standard deviation of approximately \$76MM.*

Although we have focused here on choosing strategies to maximize the expected NPV, Bickel and Smith discuss an extension of the dynamic programming model that incorporates risk aversion. If we use a risk-averse criterion to choose strategies, we might be somewhat less aggressive and not take the risk of losing these large amounts.

In this particular example, the more detailed information about geologic factors leads to more aggressive drilling decisions. In other cases however, the more detailed information may obtain higher values by quitting in scenarios where we would have continued if we did not have more detailed information. For example, following the optimal strategy without factor learning (in Fig. 3), if we drill Well 2 first and succeed, we then drill Well 4. We then continue to drill regardless of how Well 4 turns out: if Well 4 succeeds, it is optimal to drill five, and if Well 4 fails, it is optimal to drill Well 3. In contrast, with more detailed factor learning, if Well 2 succeeds and we drill Well 4 next and it fails, the decision to continue drilling depends on why Well 4 failed. Specifically, if Well 4 failed on Rock or any two factors, then it is optimal to quit drilling. Thus, in this scenario, the more detailed information leads to less aggressive drilling strategies.

Conclusions

The approach we describe for modeling dependence among prospects represents the combination of two different techniques, each of which could be used independently in other applications. The models can also be extended or modified in a variety of ways.

The information-theoretic technique we used for constructing a joint probability from a limited number of assessments could obviously be applied to another set of geologic factors: there is nothing special about the Charge, Rock, and Seal decomposition considered here. It is also not difficult to extend the approach to consider more factors or factors that have more than two outcomes. Although we have focused on the use of this joint probability distribution to determine optimal drilling strategies by dynamic programming, one could use this form of distribution in other models. For example, it could be used in a simulation model that includes other uncertainties (e.g., production rates and product prices) and studies the overall financial risks of a multiprospect opportunity. To accurately capture the total risks in such a setting, it is important to accurately model the dependence among prospects. As noted earlier, the "black box" correlation routines used in the standard Monte Carlo packages do not generate appropriately correlated samples in this setting. Though we have not discussed the accuracy of our approach for approximating complex distributions here, Bickel and Smith (2006) provide some preliminary results along these lines that are quite encouraging.

Just as the probability model could be used independently from the dynamic programming model, the dynamic programming model could also be used with joint probability distributions determined using other approaches or with different sets of geologic factors. The dynamic programming model can also be extended in several useful directions. For example, though we have focused on a case where the "rewards" correspond to the success or failure of individual wells, it is not difficult to incorporate reward functions that reflect economies of scale or synergies in development or drilling costs among prospects. Although we have focused on choosing strategies to maximize the expected NPV, Bickel and Smith[14] describe an extension of the dynamic programming model that incorporates risk aversion.

Taken together, our information-theoretic approach to constructing a joint distribution and our dynamic-programming model for determining optimal drilling strategies allows us to address a problem that should be of interest to many explorationists: How should we incorporate our evolving knowledge of the underlying geologic factors into our drilling decisions? Moreover, how can we evaluate multiprospect plays in a way that takes this learning into account? Our illustrative example demonstrates the benefits of modeling dependence and considering dynamic strategies that exploit this dependence. In this example, we can exploit the informational synergies among prospects to make five individually unattractive prospects into an attractive exploration opportunity. In this and many other cases, modeling and exploiting dependence among prospects may help firms understand and realize the full value of a complex exploration program.

Nomenclature

C = hydrocarbon charge (0 = failure, 1 = success)
$E_\pi[f]$ = expectation of function f using distribution π
$KL(\pi, \pi_0)$ = Kullback-Leibler distance between probability distributions π and π_0
n = number of prospects
$p(A)$ = marginal probability of event A
$p(A|B)$ = conditional probability of event A given event B
$r_i(j, k, l)$ = immediate reward if well i is drilled and with outcomes j, k, l, (0 = failure 1 = success) on the three geologic factors.
R = appropriate reservoir rock (0 = failure, 1 = success)
S = appropriate seal (0 = failure, 1 = success)
$v(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ = expected NPV of future cash flows given current state of factors is $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{z}$, assuming an optimal drilling strategy is used.
$v_i(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z})$ = expected NPV of future cash flows, or the continuation value, if well i is drilled given current state of factors is $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{z}$
w_i = binary variable representing success(=1)/failure(=0) of geologic factor at location i
$\boldsymbol{w}$ = vector of n binary factor variables
$\boldsymbol{w}_{-i}$ = vector of factor outcomes omitting location i
x_i, y_i, z_i = state (0 = failure, 1 = success, "–" = unknown) of factor x, y, z at location i
$\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$ = n-vectors representing state of factors x_i, y_i, z_i at each location i

* The heuristic strategy discussed earlier, where we drill wells in order of their probability of success and stops after two failures, has an expected value of \$11.71MM and a standard deviation of approximately \$83MM; a lower mean and a higher standard deviation than both of these optimal strategies. The heuristic strategy where we use this same order but stop after one failure has a mean of \$11.35MM and a standard deviation of approximately \$67MM.

x_i^0, y_i^0, z_i^0 = identical to $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$ except $x_i=0$, $y_i=0$, $z_i=0$, respectively
x_i^1, y_i^1, z_i^1 = identical to $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}$ except $x_i=1$, $y_i=1$, $z_i=1$, respectively
δ = single period discount factor
λ_0 = Lagrange multiplier for unit probability constraint
λ_i = Lagrange multiplier for marginal probability constraint at location i
λ_{ij} = Lagrange multiplier for pairwise probability constraint for locations i and j
$\boldsymbol{\lambda}$ = vector of Lagrange multipliers
$\mu_x(\boldsymbol{x}), \mu_y(\boldsymbol{y}), \mu_z(\boldsymbol{z})$ = total probability functions for vectors $\boldsymbol{x}, \boldsymbol{y}$, and $\boldsymbol{z}$
π_0 = independent joint probability distribution
π = joint probability distribution
π_x, π_y, and π_z = joint probability distributions for factors x, y, and z
π^* = optimal joint probability distribution
ρ_{AB} = correlation between A and B
Ω_0 = unit probability constraint function
Ω_i = marginal probability constraint function for location i
Ω_{ij} = pairwise probability constraint function for locations i and j

Acknowledgments

The authors thank the five anonymous reviewers for their helpful suggestions.

References

Abbas, A.E. 2006. Entropy methods for joint distributions in decision analysis. *IEEE Transactions on Engineering Management* **53** (1): 146–159. DOI: 10.1109/TEM.2005.861803.

Bellman, R.E. 1957. *Dynamic Programming.* Princeton, New Jersey: Princeton University Press.

Bickel, J.E. and Smith, J.E. 2006. Optimal Sequential Exploration A Binary Learning Model. *Decision Analysis* **3** (1): 16–32. DOI: 10.1287/deca.1050.0052.

Cover, T.M. and Thomas, J.A. 1991. *Elements of Information Theory.* New York City: John Wiley & Sons.

Delfiner, P. 2003. Modeling Dependencies Between Geologic Risks in Multiple Targets. *SPEREE* **6** (1): 57–64. SPE-82659-PA. DOI: 10.2118/82659-PA.

Genrich, J.F. and Sommer, F.S. 1989. Novel Approach to Sensitivity Analysis. *JPT* **41** (9): 930–932, 983–985. SPE-17067-PA. DOI: 10.2118/17067-PA.

Goovaerts, P. 1997. *Geostatistics for Natural Resources Evaluation.* Oxford, UK: Oxford University Press.

Iman, R.L. and Conover, W.J. 1982. A Distribution-Free Approach to Inducing Rank Correlation Among Input Variables. *Communications in Statistics—Simulation and Computation* **11** (3): 311–334. DOI: 10.1080/03610918208812265.

Jaynes, E.T. 1968. Prior Probabilities. *IEEE Transactions on Systems Science and Cybernetics* **4** (3): 227–241.

Jaynes, E.T. 1982. On the rationale of maximum-entropy methods. Proceedings of the IEEE **70** (9): 939–952.

Keefer, D.L. 2004. The Underlying Event Model for Approximating Probabilistic Dependence Among Binary Events. *IEEE Transactions on Engineering Management* **51** (2): 173–182. DOI: 10.1109/TEM.2004.826014.

Kokolis, G.P., Litvak, B.L., Rapp, W.J., and Wang, B. 1999. Scenario Selection for Valuation of Multiple Prospect Opportunities: A Monte Carlo Play Simulation Approach. Paper SPE 52977 presented at the SPE Hydrocarbon Economics and Evaluation Symposium, Dallas, 21–23 March. DOI: 10.2118/52977-MS.

Lerche, I. 1992. *Oil Exploration Basin Analysis and Economics.* San Diego, California: Academic Press.

Magoon, L.B. and Dow, W.G. 1994. The Petroleum System. In *The Petroleum System—From Source to Trap,* AAPG Memoir 60. ed. L.B. Magoon and W.G. Dow, 3–24. Tulsa: American Association of Petroleum Geologists.

Megill, R.E. 1984. *An Introduction to Risk Analysis.* Tulsa: PennWell Publishing.

Moore, L.R. and Mudford, B.S. 1999. Probabilistic Play Analysis From Geoscience to Economics: An Example From the Gulf of Mexico. Paper SPE 52955 presented at the SPE Hydrocarbon Economics and Evaluation Symposium, Dallas, 21–23 March. DOI: 10.2118/52955-MS.

Murtha, J.A. 1996. Estimating Reserves and Success for a Prospect With Geologically Dependent Layers. *SPERE* **11** (1): 37–42. SPE-30040-PA. DOI: 10.2118/30040-PA.

Murtha, J.A. and Peterson, S.K. 2001. Another Look at Layered Prospects. Paper SPE 71416 presented at the SPE Annual Technical Conference and Exhibition, New Orleans, 30 September 30–3 October. DOI: 10.2118/71416-MS.

Newendorp, P. 1975. *Decision Analysis for Petroleum Exploration.* Tulsa: PenWell Publishing.

Smith, J.E. 1993. Moment Methods for Decision Analysis. *Management Science* **39** (3): 340–358.

Stabell, C.B. 2000. Alternative Approaches to Modeling Risks in Prospects With Dependent Layers. Paper SPE 63204 presented at the SPE Annual Technical Conference and Exhibition, Dallas, 1–4 October. DOI: 10.2118/63204-MS.

Wang, B., Kokolis, G.P., Rapp, W.J., and Litvak, B.L. 2000. Dependent Risk Calculations in Multiple-Prospect Exploration Evaluations. Paper SPE 63198 presented at the SPE Annual Technical Conference and Exhibition, Dallas, 1–4 October. DOI: 10.2118/63198-MS.

J. Eric Bickel is an Assistant Professor in the Department of Industrial and Systems Engineering and Director of the Decision Analysis Systems Lab at Texas A&M University. His research interests focus on the theory and practice of decision analysis, with a special focus on the oil and gas arena. He holds a PhD degree from the Department of Engineering Economic Systems at Stanford University. **James E. Smith** is a Professor in the Decision Sciences area at the Fuqua School of Business at Duke University. His research interests also focus on the theory and practice of decision analysis. He holds a PhD degree from the Department of Engineering Economic Systems at Stanford University. **Jennifer L. Meyer** is a Senior Engagement Manager at Strategic Decisions Group in Palo Alto, California. She has worked in strategy development, economic evaluation, and business portfolio modeling for clients in the transportation, oil and gas, and technology industries and trains clients in strategic decision-making and risk management. She holds a PhD degree in Operations Research from Stanford University.