Optimal Sequential Exploration: A Binary Learning Model

Article *in* Decision Analysis · March 2006

DOI: 10.1287/deca.1050.0052 · Source: DBLP

2 authors, including:

J. Eric Bickel
University of Texas at Austin
72 PUBLICATIONS **581** CITATIONS

SEE PROFILE

DECISION ANALYSIS
Vol. 3, No. 1, March 2006, pp. 16–32
ISSN 1545-8490 | EISSN 1545-8504 | 06 | 0301 | 0016

informs

DOI 10.1287/deca.1050.0052

Optimal Sequential Exploration: A Binary Learning Model

J. Eric Bickel
Department of Industrial and Systems Engineering, Texas A&M University, 236B Zachry Engineering Center, 3131 TAMU, College Station, Texas 77843-3131, ebickel@tamu.edu

James E. Smith
Fuqua School of Business, Duke University, Box 90120, Durham, North Carolina 27708-0120, jes9@duke.edu

In this paper, we develop a practical and flexible framework for evaluating sequential exploration strategies in the case where the exploration prospects are dependent. Our interest in this problem was motivated by an oil exploration problem, and our approach begins with marginal assessments for each prospect (e.g., what is the probability that the well is wet?) and pairwise assessments of the dependence between prospects (e.g., what is the probability that both wells i and j are wet?). We then use information-theoretic methods to construct a full joint distribution for all outcomes from these marginal and pairwise assessments. This joint distribution is straightforward to calculate, has many nice properties, and appears to provide an accurate approximation for distributions likely to be encountered in practice. Given this joint probability distribution, we determine an optimal drilling strategy using an efficient dynamic programming model. We illustrate these techniques with an oil exploration example and study how dependence and risk aversion affect the optimal drilling strategies. The information-theory-based techniques for constructing joint distributions and dynamic programming model for determining optimal exploration strategies could be used together or separately in many other applications.

Key words: decision analysis; probabilistic dependence; dynamic programming; maximum entropy; exploration
History: Received on July 7, 2005. Accepted by Don N. Kleinmuntz on November 29, 2005, after 1 revision.

1. Introduction

The motivation for this paper began with a consulting engagement for a client that wanted to prioritize its deep-water oil and gas exploration program. The client had grouped the oil and gas prospects into clusters it believed to be geologically dependent and wanted to understand how the results for one well change the chance of success for the remaining undrilled prospects and how this information should affect the drilling strategy. Given that the wells cost tens of millions of dollars to drill, an optimal drilling program could generate significant savings and make an otherwise unattractive exploration opportunity economically viable. Researchers exploring multiple related R&D projects face similar problems—given dependent projects, which projects should they pursue first?

In theory, this is a straightforward problem. We simply need to specify the joint probability distribution for all possible combinations of well outcomes and then build a decision tree that considers all possible drilling sequences and possible contingencies: Which well, if any, should we drill first? If that well is wet (or dry), which well do we drill next? And so on.

However, in practice this is a very difficult problem. First, the assessment of the joint distribution for n wells requires $2^n - 1$ probability assessments. If we decompose the assessment of the joint distribution into a series of conditional assessments, many of the assessments will be heavily conditioned. For example, with nine wells, the assessments for the ninth well will be conditioned on the outcomes of eight other wells and there will be 2^8 or 256 different conditioning scenarios to consider. Most decision makers and experts do not have the time or energy to provide this many assessments, and the complex conditioning makes it difficult to be consistent. Second, even if we could generate the necessary joint probability distribution, a straightforward decision tree model for determining the optimal drilling sequence

would be unmanageably large. For example, with $n = 9$ wells, this straightforward decision tree would include approximately 460 million scenarios.

In this paper, we develop a practical and flexible approach for analyzing sequential exploration problems. After reviewing related literature in the remainder of this section, we begin in §2 by describing an example involving six wells that we will use to demonstrate the methods proposed in the paper. This is a simplified version of the client's original problem. In §3, we describe the use of information-theory-based methods for constructing a joint probability distribution given marginal probabilities for each well and pairwise joint probability assessments for all pairs of wells; this basic information could be assessed in a variety of ways. The resulting joint probability distributions are straightforward to calculate and have an intuitive interpretation. In §4, we describe a dynamic programming model for determining the optimal drilling sequence in both risk-neutral and risk-averse settings and discuss the form of the optimal policies in our example. This dynamic programming model exploits the recombining structure of the problem and allows one to solve problems efficiently. A spreadsheet that implements the procedures of §§3–4 is available from the journal's website (http://da.pubs.informs.org/online-supp.html). In §5, we examine the accuracy of our approach for constructing joint probability distributions using a simulation study. In §6, we discuss some possible extensions and other applications of these models and conclude the paper.

There are several streams of decision analysis research that address related problems. First, the problem of modeling dependence in multiple well-drilling programs has long been of interest in the oil and gas decision analysis literature. For example, Newendorp (1975) argues that wells drilled in a single basin are typically dependent and discusses some of the challenges of modeling dependence in such a setting. More recently, Wang et al. (2000) discuss modeling dependence and propose a simple model for modeling dependence in the case where wells are assumed to be 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 specific wells succeeded or failed. The Wang et al. model does not allow the possibility that some wells are more likely to hit than others and/or more strongly related than others. Keefer (2004) develops an "underlying event" model for approximately capturing positive dependence among binary events; we will discuss this approximation in §5 and compare it to our proposed approach. Although the underlying event model has been successfully applied in some settings (e.g., Keefer et al. 1991), the underlying event and the exchangeable well assumptions both seem too restrictive for general use in sequential exploration problems.

Second, our probabilistic model is related to recent work on modeling dependence in decision analysis. In §3, we focus on constructing a joint distribution given marginal probability assessments and pairwise assessments of dependence; as we will see, the resulting distribution can be related to "copulas." Clemen and Reilly (1999) and Yi and Bier (1998) discuss the use of copulas in decision and risk analysis. The use of information-theoretic methods to construct probability distributions for decision analysis dates back to the seminal paper by Jaynes (1968). MacKenzie (1994), Lowell (1994), Smith (1995), and Abbas (2003, 2006) discuss the use of information-theoretic methods for constructing joint distributions in decision analysis. More specifically, MacKenzie (1994) uses entropy methods to construct a joint distribution given marginal distributions and pairwise correlation coefficients. In a study of sensitivity to dependence, Lowell (1994) uses entropy methods to construct joint distributions with constraints on the marginal and conditional probabilities. Smith (1995) studies bounds on values and policies given constraints on probabilities and the entropy of the distribution. Abbas (2003, 2006) illustrates the use of entropy methods with marginal, pairwise, and/or three-way probability assessments to specify a joint distribution and uses simulation methods to study the accuracy of this approach. Similar problems arise in the information-theoretic analyses of contingency tables (see, e.g., Gokhale and Kullback 1978), though there the focus is on developing a model of the response of one dichotomous "effect" variable as a function of the treatment categories, rather than constructing a full joint probability distribution.

Finally, this paper is related to recent work on dynamic decision modeling or "real options." There are a number of papers that have applied real-options methods to evaluating oil and gas investments (see, e.g., Paddock et al. 1988; Smith and McCardle 1998, 1999) that focus on information gathering and its impact on decision making. The optimal sequential drilling problem we consider here can be viewed as a specific example of a dynamic decision problem that illustrates general features of such problems. Smith and Thompson (2004) also consider sequential exploration as a real-options problem and analytically characterize optimal drilling orders in some special cases. They assume all the wells are "in the money," in that they are independently economically viable in isolation and describe ordering rules for the case of $n = 2$ wells. They also develop an ordering rule for larger n in the case where, in addition to assuming that all wells are "in the money," they assume a particular "common risk" model and have additional assumptions. Though we will discuss these rules in §4, our goal is to develop methods for solving the problem in general rather than identify simple rules that are optimal in specialized circumstances.

We view the primary contribution of this paper to be the development of a practical and flexible approach for analyzing a class of problems—sequential exploration problems—that occur frequently in practice. The information theory and dynamic programming methods that we use are quite general and well established; the contribution here is in identifying an appropriate and efficient formulation of this class of problems using these techniques. Although we focus on sequential exploration problems, we believe that our paper also makes some contributions that may have implications beyond this specific context. First, as discussed in §3, the information-theoretic approach with pairwise probability constraints leads to a nice probability distribution (see Equation (7)) whose parameters could be assessed directly or related to distances between prospects; this distribution could find uses outside the sequential exploration context. Second, the dynamic programming model illustrates the complex effects of dependence in sequential decision problems; our sequential exploration problem is a relatively simple example of a dynamic decision problem with dependence and, as discussed in §4, we already have policies that confound most intuitive "rules" for solving such problems. Third, our accuracy analysis in §5 provides a demonstration of the effectiveness of information-theoretic methods in a sequential decision-making context that adds to the body of work demonstrating its effectiveness in other contexts (see, e.g., Jaynes 1982 or Abbas 2006). Finally, by illustrating how information-theoretic and dynamic programming techniques can be used in decision analysis, our paper has some tutorial value that may help inspire the use of these methods in other contexts.

2. An Illustrative Example

Suppose we are contemplating drilling wells at six sites. The marginal probability of success (i.e., a "wet" well) (p_i), the expected value given success (s_i), and the expected value given failure (f_i) for each well are shown in Table 1; these expected values are in millions of dollars and represent net present values (NPVs) in the period the well is drilled. The expected value given failure is the expected drilling cost. The expected value given success is the expected net present value of the hydrocarbon production stream less the drilling costs, costs of completion, production platforms, etc. These expectations take into account uncertainty in gas and/or oil prices, reserves, production, drilling costs, and all other uncertainties. The intrinsic values shown in Table 1 are the unconditional expected values: $p_i s_i + (1 - p_i) f_i$. In this example, the intrinsic values are all negative, meaning the wells are not "in the money," and the company would not choose to drill them if they were considered in isolation.

Table 1 Example Well Data

		Expected values		
Well	Probability of success (p_i)	Given success (s_i)	Given failure (f_i)	Intrinsic value
1	0.35	60	−35	−1.75
2	0.49	15	−20	−2.85
3	0.53	30	−35	−0.55
4	0.83	5	−40	−2.65
5	0.33	40	−20	−0.20
6	0.18	80	−20	−2.00

Though the assumptions shown in Table 1 do not reflect dependence among the different prospects, the sites were in the same general area and shared some common features. The company felt that by drilling wells sequentially, it could use the information from early wells to guide later drilling decisions and perhaps make this exploration opportunity attractive. As discussed in the introduction, to fully model this, we need to specify the joint distribution for the well results and then evaluate all the potential sequential drilling strategies.

The full specification of the joint distribution in this case would require 63 (2^6-1) probability assessments, including assessments conditioned on the outcome of five other wells. Although this full assessment was too difficult to manage, the company's experts were comfortable making pairwise conditional assessments $p(\text{well } j \text{ wet} \mid \text{well } i \text{ wet})$ for each pair of prospects. These are shown in Table 2 along with the pairwise correlation coefficients implied by these assessments. Note that well 4 was assumed to be pairwise independent of the other five wells: the conditional probabilities for well 4 being wet, given that well 1, 2, or 3 is wet are equal to the marginal probability (0.83) and the conditional probabilities for wells 5 and 6 given that well 4 is wet are equal to their marginal probabilities (0.33 and 0.18, respectively). The implied pairwise correlation coefficients involving well 4 are thus all zero. Wells 2 and 5 and wells 3 and 6 have the substantially larger pairwise correlation coefficients ($\rho_{25} = 0.459$ and $\rho_{36} = 0.359$) than the others because these sites share more common features.

There are a variety of ways to assess these pairwise dependencies. One could assess various combinations of marginal, joint, or conditional probabilities or assess pairwise correlation coefficients. The procedure for constructing a joint distribution in the next section starts with n marginal probabilities of success for the n individual wells and $n(n-1)/2$ joint probabilities for success at both well i and j. These joint probabilities were calculated from the marginal and pairwise conditional probabilities shown in Tables 1 and 2. As discussed by Moskowitz and Sarin (1983), the use of joint probability tables may help ensure the consistency of these assessments. Clemen et al. (2000) compare different methods for assessing dependence relationships and suggest that it may be easiest and most reliable to assess pairwise correlation coefficients. In the original consulting project, the conditional probabilities were assessed one way, but the implied joint and reverse conditional probabilities and correlation coefficients were also displayed to ensure that the assessments were consistent and that the expert understood the implications of the assessments.

The challenge is to use these pairwise assessments to construct an appropriate joint distribution and then determine the optimal drilling strategy.

3. Constructing the Joint Probability Distribution

We first consider the question of constructing a reasonable joint probability distribution from a set of marginal and pairwise conditional probability assessments. We chose to use an information-theoretic method to construct a joint distribution because the method performs well in a variety of contexts and has attractive theoretical properties (see, e.g., Jaynes 1982). We first describe the method for constructing the distribution and then discuss the results in the context of our example from §2. We discuss the accuracy of the approach in §5.

Table 2 Pairwise Assessments for the Example

	Direct conditional assessments $p(j \text{ wet} \mid i \text{ wet})$						Marginal	Implied correlation matrix (ρ_{ij})					
$i \setminus j$	1	2	3	4	5	6	(p_i)	1	2	3	4	5	6
1		0.59	0.63	0.83	0.39	0.31	0.35		0.147	0.147	0	0.094	0.248
2			0.65	0.83	0.55	0.24	0.49			0.236	0	0.459	0.153
3				0.83	0.42	0.31	0.53				0	0.203	0.359
4					0.33	0.18	0.83					0	0
5						0.26	0.33						0.146
6							0.18						

Method

Let $\mathbf{w} = (w_1, \ldots, w_n)$ be a vector of n binary random variables where $w_i = 1$ if event i occurs (e.g., the well is wet) and 0 otherwise. Our goal will be to construct a joint probability distribution $\pi(\mathbf{w})$ given information about the marginal probabilities for each event and pairwise joint probabilities. Specifically, we will assume that we have assessed the n marginal probabilities $p_i \equiv p(w_i = 1)$ and the $n(n-1)/2$ pairwise joint probabilities $p_{ij} \equiv p(w_i = 1, w_j = 1)$. We will assume these assessments are consistent in that $0 < p_i < 1$ and $0 < p_{ij} < p_i$. Beyond consistency, we make no specific assumptions about the assessed probabilities.[1, 2]

We will construct the joint distribution $\pi(\mathbf{w})$ by selecting a probability distribution to minimize the relative entropy or Kullback-Liebler (KL) distance relative to a reference distribution $\pi_0(\mathbf{w})$ that assumes independence, subject to the constraint of matching the specified marginal and pairwise joint assessments. This KL distance is defined as

$$D(\pi, \pi_0) \equiv \sum_{\mathbf{w}} \pi(\mathbf{w}) \ln\left(\frac{\pi(\mathbf{w})}{\pi_0(\mathbf{w})}\right) \tag{1}$$

and the independent reference distribution as

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

With $\pi_0(\mathbf{w})$ taken to be the independent reference distribution, the KL distance $D(\pi, \pi_0)$ is also called the *mutual information* among the binary random variables $\mathbf{w}$ with distribution π. Noting that $D(\pi, \pi_0) \geq 0$ and $D(\pi, \pi_0) = 0$ if and only if $\pi = \pi_0$, we can interpret the objective function in (1) as a measure of the strength of dependence in the joint distribution π. Our joint distributions will thus minimize the amount of dependence (according to this measure) or be "maximally uncertain," subject to the constraint of matching the specified assessments. In this sense, by minimizing the KL distance from the independent reference distribution, we are being conservative in estimating the amount of learning that can take place in the sequential exploration problem. (For more on KL distance and mutual information, see, e.g., Cover and Thomas 1991.)[3]

To represent the constraints on the joint distribution, we define the following functions:

$$\Omega_0(\mathbf{w}) = 1$$
$$\Omega_i(\mathbf{w}) = w_i \quad \text{for all } i$$
$$\Omega_{ij}(\mathbf{w}) = w_i w_j \quad \text{for all } i, j,$$

and let $\mathrm{E}_\pi[f(\mathbf{w})]$ denote the expectation of some function $f(\mathbf{w})$ when $\mathbf{w}$ has distribution π. Using this notation, we can represent the constraint on marginal probabilities as requiring π to be such that $\mathrm{E}_\pi[\Omega_i(\mathbf{w})] = p_i$ and the pairwise joint probabilities as requiring $\mathrm{E}_\pi[\Omega_{ij}(\mathbf{w})] = p_{ij}$. The constraint $\mathrm{E}_\pi[\Omega_0(\mathbf{w})] = 1$ requires the joint probabilities to sum to 1. We can then state the primal optimization problem for determining joint probabilities as

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

subject to

$$\mathrm{E}_\pi[\Omega_0(\mathbf{w})] = 1$$
$$\mathrm{E}_\pi[\Omega_i(\mathbf{w})] = p_i \quad \text{for all } i$$
$$\mathrm{E}_\pi[\Omega_{ij}(\mathbf{w})] = p_{ij} \quad \text{for all } i, j.$$

The primal objective function is convex in π, and the constraints are linear in π. Thus, any locally optimal solution to (3) is a global optimum.

We can simplify this optimization problem and learn more about the form of the optimal solution by considering the Lagrangian dual of this optimization problem (see, e.g., Luenberger 1989). The use of duality in information-theoretic analyses is standard (see, e.g., Jaynes 1968). To formalize this dual approach, let

[1] To streamline our notation, we will not explicitly list the ranges for $\mathbf{w}$, the events, or pairwise events in our summations and lists of constraints. The vector $\mathbf{w}$ will range over the 2^n possible combinations of outcomes of the n events. The marginal probabilities p_i and corresponding Lagrange multipliers will range from $i = 1, \ldots, n$. Similarly, the pairwise joint probabilities p_{ij} and corresponding Lagrange multipliers will range over the $n(n-1)/2$ unique pairs of events, which may be indexed as $i = 1, \ldots, n-1$ and $j = i+1, \ldots, n$.

[2] Though it would be valid to have $p_{ij} = 0$ or $p_{ij} = p_i$, this would imply zero probability joint events that cause numerical difficulties (infinite Lagrange multipliers) for the solution procedures we use below.

[3] If we had instead taken $\pi_0(\mathbf{w}) \equiv 1$, minimizing $D(\pi, \pi_0)$ would be equivalent to maximizing entropy $H(\pi) = -\sum_{\mathbf{w}} \pi(\mathbf{w}) \ln \pi(\mathbf{w})$.

λ_0, λ_i, and λ_{ij} denote the Lagrange multipliers associated with the constraints in (3) and $\boldsymbol{\lambda}$ the vector of these Lagrange multipliers; the Lagrangian is

$$L(\pi, \boldsymbol{\lambda}) \equiv \sum_{\mathbf{w}} \pi(\mathbf{w}) \ln\left(\frac{\pi(\mathbf{w})}{\pi_0(\mathbf{w})}\right) - \lambda_0(\mathrm{E}_\pi[\Omega_0(\mathbf{w})] - 1) - \sum_i \lambda_i(\mathrm{E}_\pi[\Omega_i(\mathbf{w})] - p_i) - \sum_{i,j} \lambda_{ij}(\mathrm{E}_\pi[\Omega_{ij}(\mathbf{w})] - p_{ij}).$$

We can solve the primal KL minimization problem (3) by solving the dual problem

$$\max_{\boldsymbol{\lambda}} \min_{\pi} L(\pi, \boldsymbol{\lambda}). \tag{4}$$

This dual problem will have a solution whenever the primal is feasible in that the probabilities specified as constraints are consistent. Moreover, the optimal value of the dual objective will equal that of the primal, and the optimal $\boldsymbol{\lambda}^*$ will correspond to the distribution π^* that solves the primal. If the assessments are not consistent (i.e., if there is no probability distribution satisfying the constraints in (3)), the dual problem (4) will be unbounded.

We can identify the form of the optimal solution π^* and simplify the dual problem with some additional analysis. Differentiating $L(\pi, \boldsymbol{\lambda})$ with respect to π and setting the result equal to zero, we obtain the form of the optimal joint distribution as a function of $\boldsymbol{\lambda}$:

$$\pi^*(\mathbf{w}, \boldsymbol{\lambda}) = \pi_0(\mathbf{w}) \exp\left(-1 + \lambda_0 + \sum_i \lambda_i \Omega_i(\mathbf{w}) + \sum_{i,j} \lambda_{ij} \Omega_{ij}(\mathbf{w})\right). \tag{5}$$

Substituting (5) into (4), the dual optimization problem (4) becomes

$$\max_{\boldsymbol{\lambda}} \left(-\sum_{\mathbf{w}} \pi^*(\mathbf{w}, \boldsymbol{\lambda}) + \lambda_0 + \sum_i \lambda_i p_i + \sum_{i,j} \lambda_{ij} p_{ij}\right). \tag{6}$$

One can check that this dual objective function is concave in $\boldsymbol{\lambda}$ and that the first-order conditions for optimality of (6) imply that the constraints of (3) are satisfied. Thus, rather than solving the original constrained optimization problem (3) involving 2^n unknown joint probabilities and $1+n+n(n-1)/2$ constraints, we can instead solve the dual problem (6), which is concave and involves $1 + n + n(n-1)/2$ Lagrange multipliers and has no constraints. In our example with $n = 6$, the primal problem (3) has 64 decision variables and 22 constraints; the dual problem (6) has 22 decision variables and no constraints. This dual problem (and larger ones) can be solved easily using spreadsheet-based optimization packages. In our example with $n = 6$ wells, this dual problem can be solved in one or two seconds using Solver in Excel. In contrast, Solver failed to reliably solve the primal formulation of this example.

The optimal distribution $\pi^*(\mathbf{w}, \boldsymbol{\lambda})$ in (5) has several nice properties and interpretations. First, it can be related to copulas: Because we take the reference prior $\pi_0(\mathbf{w})$ to be the independent distribution (2), $\pi^*(\mathbf{w}, \boldsymbol{\lambda})$ can be written as

$$\pi^*(\mathbf{w}, \boldsymbol{\lambda}) = (\text{Product of marginal distributions}) \times (\text{density weighting function}),$$

where the density weighting function (or copula density) is the exponential term in (5) and encodes the dependence among the n random variables (see, e.g., Clemen and Reilly 1999, Miller and Liu 2002). Our KL minimization problem can thus be interpreted as finding the copula that minimizes dependence subject to the given marginal and pairwise probability constraints.

Second, the conditional probability distributions implied by these joint distributions have a nice structure. Letting $\mathbf{w}_{-i}$ denote the vector of outcomes omitting event i, from (5) we find that the conditional log-odds for event i can be written as

$$\ln\left(\frac{p(w_i = 1 \mid \mathbf{w}_{-i})}{1 - p(w_i = 1 \mid \mathbf{w}_{-i})}\right) = \ln\left(\frac{p_i}{1 - p_i}\right) + \lambda_i + \sum_{j \neq i} \lambda_{ij} w_j. \tag{7}$$

That is, the conditional log-odds for one event (e.g., finding that well i is wet) is a linear function of the outcome of the other events. The marginal Lagrange multipliers λ_i (which are typically negative) describe the adjustment to the log-odds in the event that all other wells fail. The joint Lagrange multipliers λ_{ij} describe the increase in the log-odds of event i due to the occurrence of event j. Equation (7) has the form of a logistic regression model (see, e.g., Neter et al. 1996) and, thus, the joint distribution for the wells is

described by a system of a logistic regression equations in which the occurrence of event j increases the log-odds of event i by λ_{ij}, and the occurrence of the event i has a symmetric effect on the log-odds for event j.

Finally, the Lagrange multipliers can also be interpreted in the usual way for constrained optimization problems as describing how the objective function changes with changes in the constraint values. Here the objective function is a measure of the mutual information among the binary events, and the λ_i and λ_{ij} describe how changes in the probabilities p_i and p_{ij} (holding all marginal and other joint probabilities constant) will change this measure of dependence.[4]

Example Results

Table 3 shows the Lagrange multipliers associated with the example described in §2. These results illustrate several general features of these KL minimizing joint distributions. First, recall that in our initial assessments, the results for well 4 were assumed to be pairwise independent of the other wells. In Table 3, we see that the pairwise Lagrange multipliers involving well 4 are all zero, implying that well 4 is independent of all other wells. In general, if some event i is pairwise conditionally independent of all other events—that is, if $p(w_i \mid w_j) = p(w_i)$ for all $j \neq i$—then the resulting distribution will have $p(w_i \mid \mathbf{w}_{-i}) = p(w_i)$. This follows from the fact that the optimization procedure minimizes mutual information among the events; if there is nothing in the constraints that implies that w_i and $\mathbf{w}_{-i}$ are dependent, then the mutual information minimizing distributions will preserve this independence (see Cover and Thomas 1991, p. 27). However, as will be demonstrated below, pairwise independence between events i and j in the assessments—that is, if $p(w_i \mid w_j) = p(w_i)$ for some, but not all, j—need not imply that $\lambda_{ij} = 0$ in the optimal joint distribution.

[4] If we had maximized entropy instead of minimizing KL distance or mutual information (i.e., taking $\pi_0(\mathbf{w}) = 1$ instead of the independent prior defined by Equation (2)), we would arrive at the same joint distribution but with the marginal Lagrange multipliers adjusted to incorporate the marginal probabilities. Specifically, marginal Lagrange multipliers μ_i for the entropy formulation would be $\mu_i = \ln(p_i/(1-p_i)) + \lambda_i$. The pairwise Lagrange multipliers (μ_{ij} and λ_{ij}) are the same in the two formulations, and Equation (7) holds with the marginal log-odds incorporated into μ_i.

Table 3 Lagrange Multipliers for the Example

	Optimal lambdas (λ_{ij})						
$i \setminus j$	1	2	3	4	5	6	Marginal (λ_i)
1		0.45	0.20	0.00	0.03	1.12	−0.57
2			0.64	0.00	2.09	0.18	−1.17
3				0.00	0.46	2.40	−0.84
4					0.00	0.00	0.00
5						0.29	−1.56
6							−2.44
λ_0							1.856

Second, notice that the pairwise Lagrange multipliers are all positive, except for those involving well 4. Following Equation (7), this implies that the outcomes of the wells (except well 4) are all positively related, meaning a wet (dry) result at one well increases the probability of a wet (dry) result at each of the other wells. Comparing the pairwise Lagrange multipliers in Table 3 with the pairwise correlation coefficients in Table 2, we see that larger multipliers tend to appear with larger correlation coefficients: The two largest Lagrange multipliers, $\lambda_{25} = 2.09$ and $\lambda_{36} = 2.40$, correspond to the two wells with the largest correlation coefficients, $\rho_{25} = 0.459$ and $\rho_{36} = 0.359$. However, the relationship between these two measures of dependence is not perfect: The rank of the top two are reversed for the two measures and, considering the relationships involving well 1, we see that $\rho_{12} = 0.147 \approx \rho_{13} = 0.147$, yet $\lambda_{12} = 0.45 > \lambda_{13} = 0.20$. These differences should not be surprising, because the two measures of dependence are fundamentally different: The Lagrange multipliers λ_{ij} describe the joint distribution (through Equation (5)) and the probabilities p_{ij} and correlation coefficients ρ_{ij} consider the pairwise relationships.

It is quite possible to have positive correlation coefficient ρ_{ij} and positive association ($p_{ij} > p_i \times p_j$), yet have a negative λ_{ij}; Table 4 shows the data for such an example with three events. In this case, all three events are positively correlated, but the negative Lagrange multiplier $\lambda_{13} = -0.37$ implies that occurrence of event 3 decreases the odds of event 1 (and vice versa). This negative Lagrange multiplier also implies that increasing the pairwise probability p_{13} (or equivalently increasing the correlation ρ_{13}), holding all other constraints constant, would reduce the

Table 4 An Example with Positive Correlation and Negative Lagrange Multiplier

	Conditional assessments $p(j \text{ wet} \mid i \text{ dry})$			Marginal	Correlation matrix (ρ_{ij})			Optimal lambdas (λ_{ij})			Marginal
$i \backslash j$	1	2	3	(p_i)	1	2	3	1	2	3	(λ_i)
1		0.75	0.60	0.50		0.50	0.20		2.40	−0.37	−1.012
2			0.75	0.50			0.50			2.40	−2.398
3				0.50							−1.012
λ_0											1.788

overall dependence in the problem as measured by the KL objective function. If we reduce the conditional assessment $p(3 \text{ occurs} \mid 1 \text{ occurs})$ from 0.60 to 0.50 (the marginal probability for event 3) so these two events are pairwise independent, the minimal KL objective would increase and λ_{13} would be even more negative ($\lambda_{13} = -13.59$). Thus, pairwise independence between a pair of events in the assessments need not (and in general will not) imply conditional independence in the optimal joint distribution, and decreasing the correlation between two events may lead to more dependence according to the KL measure. Conversely, $\lambda_{ij} = 0$ need not imply $\rho_{ij} = 0$.

The calculated Lagrange multipliers also provide a nice way to check the original assessments. For example, if one expected there to be a positive dependence among the wells and found some negative λ_{ij}s, then this would suggest the need to revisit the original assessments. Similarly, one might expect two wells that are near each other to have stronger relationships and larger λ_{ij} than two wells that are farther apart. If such a relationship does not hold in the optimal distribution, it again suggests a need to revisit the original assessments. Alternatively, after building some intuition about the relationship between the Lagrange multipliers and conditional probabilities and/or correlation coefficients, one might assess these Lagrange multipliers directly or develop a model that relates them to the distance between the wells or some other measure of similarity between the wells.

In some cases, one may wish to omit certain pairwise assessments, perhaps for lack of time or resources to complete the assessment. This causes no problems for the KL optimization procedure; the optimal distribution simply assigns a pairwise Lagrange multiplier $\lambda_{ij} = 0$ for these omitted assessments. This approach is conservative in that it minimizes the amount of mutual information among the variables, but one would have to consider whether this is desirable in a given application. If, for example, one expects a positive relationship among these variables, one might instead want to set λ_{ij} to some default positive value, perhaps based on a comparison to calculated λ_{ij} for nearby wells.

4. Determining the Optimal Exploration Strategy

Now suppose we have specified a joint probability distribution for the well outcomes; how do we determine the optimal drilling sequence? In this section, we describe a dynamic programming formulation of this problem, discuss the results in the context of our example problem, and then consider a risk-averse version of the basic formulation and an extension that allows synergies between prospects.

Formulation

The structure of the sequential exploration problem is straightforward: We need to decide which well, if any, to drill first; if that well is wet (or dry), which well do we drill next—and so on, through the n stages. A partial decision tree for this problem is shown in Figure 1. Although conceptually straightforward, these trees become quite complex even with moderate numbers of wells. For example, with $n = 6$ wells, this leads to a total of 113,959 scenarios.[5] Such a tree is large

[5] We define a scenario as an alternative that is considered or a possibility that must be contemplated when making a decision; this is the total number of branches in the tree. Equivalently, this is one less than the total number of nodes in the tree (counting terminal nodes), because every node lies at the end of a branch except the initial node. This measure is a reasonable indicator for the amount of computational effort required to solve the tree and can be fairly

Figure 1 A Partial Decision Tree for the Sequential Drilling Problem

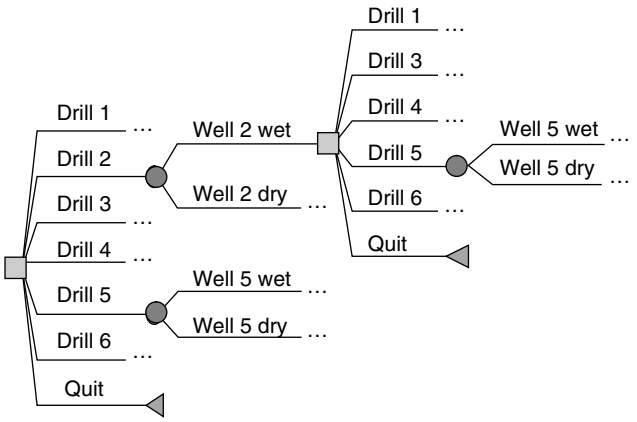

but could be handled using a professional decision tree program like DPL. As indicated in the introduction, with $n = 9$ wells the corresponding straightforward tree would include approximately 460 million scenarios.

We can simplify this decision tree model if we recognize 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 observed well 1 to be dry and well 3 to be wet, the future (conditional) probabilities and cash flows are the same regardless of whether we drilled well 3 first or well 1 first. This recombining feature of the problem is a consequence of Bayesian updating—posterior distributions do not depend on the order in which information is received—and does not rely on the specific properties of the optimal distributions discussed in §3. For the example with six wells, there are a total of 3^6 or 729 different possible well states (each well may be wet, dry, or undrilled) and a total of 2,187 possible scenarios to be considered, compared to 113,959 scenarios in the nonrecombining tree.[6] With $n = 9$ wells, the recombining tree will have approximately 79,000 scenarios to evaluate, compared to approximately 460 million in the nonrecombining tree. Though the complexity of the recombining tree grows rapidly with the number of wells, the recombining trees grow much more slowly than their nonrecombining counterparts.

We solve these recombining trees by working backward: We first figure out whether you should drill the "last" well conditional on the outcome the first $n-1$ wells. We then decide which well to drill if we had two wells remaining. And so on to the initial decision. To describe the solution procedure more precisely, let $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_i, \ldots, \omega_n)$ denote the state where $\omega_i = 0$ or 1 if the well is dry or wet and equal to $\omega_i =$ "–" if the well has not been drilled. For example, with $n =$ 6 wells, the initial state is $\boldsymbol{\omega} = (-, -, -, -, -, -)$, which means no wells have been drilled. The vector $\boldsymbol{\omega} =$ $(0, -, 1, -, -, -)$ represents the state in which well 1 was dry and well 3 was wet and the other wells have not yet been drilled.

Given the joint probability distribution π over well outcomes (constructed using the information-theoretic methods from §3 or any other method), it is straightforward to calculate the transition probabilities required for the dynamic programming model. First, let $\mu(\boldsymbol{\omega})$ be the total probability associated with the vector $\boldsymbol{\omega}$, constructed by summing $\pi(\mathbf{w})$ over the possible scenarios for these unknown events. For example in the case of $n = 6$ wells, for $\boldsymbol{\omega} =$ $(0, -, 1, -, -, -)$,

$$\mu(\boldsymbol{\omega}) = \sum_{w_2, w_4, w_5, w_6} \pi(0, w_2, 1, w_4, w_5, w_6),$$

where w_2, w_4, w_5, and w_6 range over $\{0, 1\}$; this summation would involve 2^4 probabilities. These probabilities can be easily computed from the joint

compared to recombining trees, which are discussed below. If there are n wells, we must consider stages $i = 0, 1, \ldots, n$, where the stage indicates how many wells have already been drilled. In stage i, there are ${}_nC_i$ (n "choose" i) possible combinations of previously drilled sites; $i!$ different possible orderings of the previously drilled wells; 2^i different outcomes of the i drilled wells; and $n - i + 1$ possible actions (undrilled sites + the option to quit) in that scenario. The total number of scenarios that must be considered is thus $\sum_{i=0}^{n} ({}_nC_i) i! 2^i (n - i + 1)$.

[6] These scenarios are defined as in the nonrecombining tree. If there are n wells, we must consider stages $i = 0, 1, \ldots, n$, where the stage indicates how many wells have already been drilled. In stage i, there are ${}_nC_i$ possible combinations of drilled sites; 2^i different outcomes of the i drilled wells; and $n - i + 1$ possible actions (undrilled sites + the option to quit) in each scenario. The total number of scenarios considered is thus $\sum_{i=0}^{n} ({}_nC_i) 2^i (n - i + 1)$. Comparing this to the scenarios in the nonrecombining tree, the difference is that in the recombining tree we recognize that the $i!$ different orders for previously drilled wells all lead to the same future prospects.

distribution π. Though $\mu(\boldsymbol{\omega})$ is a probability for each $\boldsymbol{\omega}$ (it lies between 0 and 1), μ is not a probability distribution, as $\sum_{\boldsymbol{\omega}} \mu(\boldsymbol{\omega})$ exceeds one.

We can use this total probability function to calculate the transition probabilities required for the dynamic programming model. Suppose that you start in a state $\boldsymbol{\omega}$, where well i hasn't been drilled (thus $\omega_i =$ "-"). If you drill well i, the probability that it is wet is $\mu(\boldsymbol{\omega}_i^1)/\mu(\boldsymbol{\omega})$, where $\boldsymbol{\omega}_i^1$ is identical to $\boldsymbol{\omega}$ except $\omega_i = 1$ and the probability that it is dry is $\mu(\boldsymbol{\omega}_i^0)/\mu(\boldsymbol{\omega})$, where $\boldsymbol{\omega}_i^0$ is identical to $\boldsymbol{\omega}$ except $\omega_i = 0$.

The dynamic programming model can now be formalized as follows. Let $v(\boldsymbol{\omega})$ denote the continuation value for state $\boldsymbol{\omega}$, that is, the expected NPV of future cash flows given that you start in state $\boldsymbol{\omega}$. In this value calculation, we include the expected future value for a successful well (s_i) or a failed well (f_i) in the period the well is drilled and discount cash flows using a discount factor δ that corresponds to the time required to drill the well. As discussed in the example in §2, s_i is the discounted expected net present value of production, in the period the well is drilled, from a successful well i (including drilling costs), and f_i is the expected cost of drilling an unsuccessful well i. If all the wells have been drilled (i.e., $\boldsymbol{\omega}$ is a vector of zeros and ones), then $v(\boldsymbol{\omega}) = 0$. For earlier states $\boldsymbol{\omega}$, the expected NPV associated with drilling well i is

$$v_i(\boldsymbol{\omega}) = \frac{\mu(\boldsymbol{\omega}_i^1)}{\mu(\boldsymbol{\omega})}(s_i + \delta v(\boldsymbol{\omega}_i^1)) + \frac{\mu(\boldsymbol{\omega}_i^0)}{\mu(\boldsymbol{\omega})}(f_i + \delta v(\boldsymbol{\omega}_i^0)), \quad (8)$$

where, as before, $\boldsymbol{\omega}_i^1$ is identical to $\boldsymbol{\omega}$ except $\omega_i = 1$ and $\boldsymbol{\omega}_i^0$ is identical to $\boldsymbol{\omega}$ except $\omega_i = 0$. The optimal action in state $\boldsymbol{\omega}$ is to drill the well with the largest $v_i(\boldsymbol{\omega})$ or, if no well has a positive value, to not drill at all. The optimal continuation value $v(\boldsymbol{\omega})$ is $\max\{v_i(\boldsymbol{\omega}), 0\}$, where the maximum is taken over all available wells and not drilling (0). There is no circularity in this definition of the value function because one never visits the same state twice; each time you drill a well, its state changes to either wet or dry.

Example Results

We can illustrate this procedure by considering the results in our example involving six wells. We assume a discount rate of 1% per period corresponding to a discount factor of $\delta = 1/(1 + 0.01) = 0.99$. Given that these deep-water wells take about six weeks to drill, this corresponds to an annual discount rate of approximately 9%. Our spreadsheet version of the dynamic programming model begins with the joint probability distribution $\pi(\mathbf{w})$ and calculates the total probabilities $\mu(\boldsymbol{\omega})$ by a matrix multiplication, where the matrix indicates which of the 2^6 states $\mathbf{w}$ are included in each of 3^6 or 729 total probability states $\boldsymbol{\omega}$. Given these total probabilities, the solution of the dynamic program requires calculating 1,458 equations of the form of Equation (8) (there are six actions and 3^6 or 729 possible well states, but two-thirds of these state-action pairs correspond to scenarios where the well has already been drilled) and then calculating 729 maximums of seven values. Despite the relatively large numbers of calculations involved, the calculations are all simple: Given a joint probability distribution, the total probabilities and solution of the dynamic program are determined almost instantaneously in Excel.

The resulting optimal strategy is summarized in Figure 2 and calls for drilling well 3 first, yielding an expected value of \$14.40 million. If well 3 is dry, you quit without drilling any more wells. If well 3 is wet, you drill well 6 next and the optimal continuation value is \$28.50 million. If well 6 is wet, you

Figure 2 The Optimal Drilling Strategy for the Example

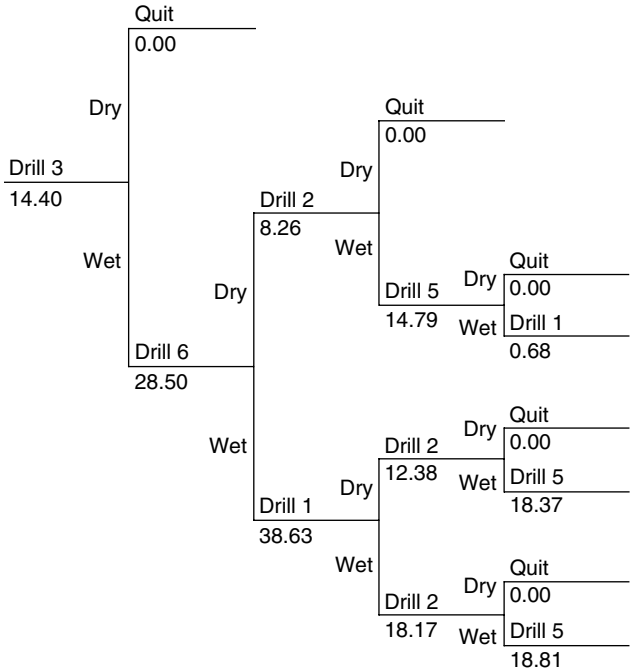

drill well 1 and then well 2, regardless of how well 1 turns out. If well 2 is wet, you then drill well 5. If well 6 is dry, you then drill well 2. If well 2 is dry, you quit, and if well 2 is wet, you continue to well 5. If 6 is dry but wells 2 and 5 are both wet, you then come back to drill well 1. You never drill well 4. As discussed in §2, if you had to drill using a nonflexible strategy (either drilling simultaneously or in a fixed schedule), you would prefer not to drill any of the wells. The optimal flexible drilling program thus adds \$14.40 million to the expected value. Following the optimal strategy, you are sure to drill one well, and there is a 23% chance that you would drill 5 wells, despite the fact that none of these wells is economically viable in isolation.

The optimal strategy for this example confounds most simple rules for recommending a drilling order. For example, we do not drill the well with the highest probability of success first; that rule would suggest drilling well 4 first, which we do not drill at all. Similarly, we do not drill the cheapest well first; that rule would suggest drilling wells 2, 5, or 6 first, all of which have the same minimal drilling costs. We do not drill the well with the smallest expected loss $((1-p_i)\times f_i)$; that rule would suggest well 4 again. Nor do we drill the well with the largest intrinsic value; that rule would suggest drilling well 5 first. Finally, we do not drill the potentially largest well, 6, first. Smith and Thompson (2004) show that, given a "common-risk structure" and "in-the-money" wells, it is optimal to first drill the well with the largest intrinsic value if it also has the largest probability of success. These assumptions do not hold in this example, because the well with the largest probability of success, well 4, does not have the largest or a positive intrinsic value. We could, however, increase the value given success (s_4) from \$5 million to \$10 million, thereby increasing its intrinsic value to \$1.5 million, giving it the highest intrinsic value and the largest probability of success; though we would now drill well 4, we still would not drill it first.

Why is the strategy of Figure 2 optimal? First, you don't drill well 4 because it has a negative intrinsic value and its result is probabilistically independent of the other wells; the information provided by the other well results does not affect its values or probabilities and hence cannot change this marginal decision. In addition, well 4 provides no information regarding other wells. If well 4 had a positive intrinsic value, you would drill the well regardless of how the other wells turn out; however, its place in the optimal drilling sequence will vary depending on how the other wells turn out. For example, if we increase well 4's value given success (s_4) from \$5 to \$10 million, the intrinsic value of well 4 is positive and it is optimal to drill well 4 somewhere from period 2 to 6, with the sequence and timing depending on how the others turn out. If we increase s_4 beyond \$15.25 million (corresponding to an intrinsic value of \$5.86 million), then it becomes optimal to drill well 4 first. The preference for the timing of well 4 is driven by discounting: If we didn't discount the cash flows, it wouldn't matter when you drill well 4.

Why drill well 3 first? To understand the structure of the optimal policy, it helps to examine the expected NPVs for alternative initial well choices as shown in Table 5: These NPVs assume that an optimal drilling strategy is followed after the first well. (These are the $v_i(\boldsymbol{\omega})$ given by Equation (8) for the initial state $\boldsymbol{\omega}=(-,-,-,-,-,-)$ and are automatically calculated when solving the dynamic program.) In Table 5, we see that drilling well 3 first is slightly better (expected NPV = \$14.40 million) than drilling well 2 first (expected NPV = \$14.34 million). If we were to drill well 2 first, then it would be optimal to quit if well 2 is dry and to drill well 5 next if well 2 is wet. If we were instead to drill one of the other wells (1, 5, or 6) first, we would have a substantially lower expected NPV. Comparing wells 3 and 2 to the other wells, we see that these wells both have relatively high marginal probabilities of success (0.53 and 0.49, respectively) and strong links to other wells. In Table 3, we see that well 3 is strongly linked to well 6, as indicated by the Lagrange multiplier value of $\lambda_{36}=2.40$. Well 2 is strongly linked to well 5 with $\lambda_{25}=2.09$. These strong links help explain why there is considerable value associated with the information provided

Table 5 Expected NPVs for Alternative Initial Drilling Decisions

		Well					
	Quit	1	2	3	4	5	6
Expected NPV	0	10.88	14.34	14.40	11.61	11.44	10.64

by the early well results and why it is optimal to drill well 6 next if well 3 is successful and well 5 next if well 2 is successful.

The structure of this solution suggests that there are no simple general rules for determining the optimal drilling order. A simple rule would evaluate some index based on individual well data and prioritize wells based on that index. A classic example of such an index is the Gittins index (see, e.g., Ross 1983) that can be used to prioritize individual projects in the independent setting. In the dependent setting, one such simple index might be the intrinsic value of the well, which would change over time as the probabilities of success are updated. As discussed earlier, Smith and Thompson (2004) show that the ordering given by this index is optimal given particular assumptions. However, such rules cannot work with general dependence structures. For instance, in our example, if we increase the value given success for well 5 (s_5) from \$40 million to \$41 million, it becomes optimal to drill well 2 first rather than well 3. Thus changing one well's parameters may change the relative ordering of other wells! Intuitively this makes sense: Given the tight link between wells 2 and 5 and between wells 3 and 6, it is as if the pair of wells 2 and 5 is competing against the pair of wells 3 and 6 to go first. Improving well 5 improves the 2–5 pair and moves well 2 to the front of the drilling order, with well 5 following if successful. These pairs are not independent, however: If you drill well 3 first (or well 2 first) and fail, it is optimal to quit entirely rather than explore the other pair. Examples like this seem entirely reasonable and appear to preclude the possibility of using simple index rules for determining optimal policies; the optimal policies depend critically on the dependence among wells in addition to the parameters for the individual wells.

Risk-Averse Formulation

If we assume that the decision maker has an exponential utility function for the NPV of the cash flows generated by the exploration opportunity, we can capture risk aversion by recasting the dynamic programming recursion in terms of certainty equivalents rather than expected values. Specifically, suppose the decision maker's utility function is $u(x) = -\exp(-x/R)$, where x is the NPV and R is the decision maker's risk tolerance. With this utility function, the certainty equivalent of a gamble $\tilde{x}$ is given by $\mathrm{CE}[\tilde{x}] = -R\ln(\mathrm{E}[\exp(-\tilde{x}/R)])$. The recursive calculation of certainty equivalents exploits the "delta property" of the exponential utility function: If we add a constant Δ to all outcomes of a gamble $\tilde{x}$, its certainty equivalent increases by Δ; that is, $\mathrm{CE}[\tilde{x}+\Delta] = \mathrm{CE}[\tilde{x}] + \Delta$. This property allows us to add the rewards (s_i and f_i) to the continuation values as we did in risk-neutral recursion, given in Equation (8).

Let $v'(\boldsymbol{\omega})$ be the continuation certainty equivalent for state $\boldsymbol{\omega}$; that is, the decision maker's certainty equivalent for the exploration opportunity starting in state $\boldsymbol{\omega}$. Like the expected values $v(\boldsymbol{\omega})$ defined in Equation (8), these are stated in then-current values and are discounted through the recursion formula. The dynamic programming recursion can be defined in terms of certainty equivalents as follows. If all the wells have been drilled (i.e., $\boldsymbol{\omega}$ is a vector of zeros and ones), then $v'(\boldsymbol{\omega}) = 0$, as before. For earlier states $\boldsymbol{\omega}$, the certainty equivalent for drilling well i is

$$v_i'(\boldsymbol{\omega}) = -\rho(\boldsymbol{\omega})\ln\left(\frac{\mu(\boldsymbol{\omega}_i^1)}{\mu(\boldsymbol{\omega})}\exp\left(-\frac{s_i+\delta v'(\boldsymbol{\omega}_i^1)}{\rho(\boldsymbol{\omega})}\right) + \frac{\mu(\boldsymbol{\omega}_i^0)}{\mu(\boldsymbol{\omega})}\exp\left(-\frac{f_i+\delta v'(\boldsymbol{\omega}_i^0)}{\rho(\boldsymbol{\omega})}\right)\right). \quad (9)$$

Here, as before, $\boldsymbol{\omega}_i^1$ is identical to $\boldsymbol{\omega}$ except $\omega_i = 1$ and $\boldsymbol{\omega}_i^0$ is identical to $\boldsymbol{\omega}$ except $\omega_i = 0$. The risk tolerances $\rho(\boldsymbol{\omega})$ in these certainty equivalent calculations are a function of $\boldsymbol{\omega}$ and are inflated to reflect discounting: If the risk tolerance for current (period 0) dollars is R (as assumed above), then the corresponding risk tolerance for period t dollars must be inflated to $\rho(\boldsymbol{\omega}) = R \times \delta^{-t}$, where t is the period this well is drilled; in this model, the period is simply the number of previously drilled wells in state $\boldsymbol{\omega}$.[7] As before, the optimal action in state $\boldsymbol{\omega}$ is to drill the prospect with the largest $v_i'(\boldsymbol{\omega})$ or, if no well has a positive value, not to drill at all. The optimal continuation certainty equivalent $v_i'(\boldsymbol{\omega})$ is $\max\{v_i'(\boldsymbol{\omega}), 0\}$, where the maximum is taken over all available wells and not drilling (0).

Applying this risk-averse procedure in our example problem, we find that for risk tolerances below

[7] Rather than inflating risk tolerances, we could calculate the same present (period 0) certainty equivalents and optimal policies by discounting the cash flows before applying the dynamic programming procedure and taking $\delta = 1$ and $\rho(\boldsymbol{\omega}) = R$ in Equation (9).

$92 million, it is optimal to not drill at all. For risk tolerances between $92 million and $7.571 billion, it is optimal to drill well 2 first, followed by well 5 if well 2 is successful, with later elements of the policy changing slightly over this range. With risk tolerances greater than $7.571 billion, the optimal policy is identical to the risk-neutral case. Comparing the data for the wells in Table 1, we see that wells 2 and 5 have larger probabilities of success and lower costs given failure than wells 3 and 6, but lower values given success. Given that the expected values for these two strategies are close ($14.40 million vs. $14.34 million; see Table 5), a relatively modest degree of risk aversion—a risk tolerance less than $7.57 billion—leads to adopting a less risky strategy that begins with drilling well 2.

Synergies Among Prospects

Finally, we note that we can easily incorporate synergies between prospects, such as cost reductions resulting from transportation and production facilities that could be used to support multiple wells. In our description of the dynamic programming model, we have assumed that the rewards for a single well (s_i and f_i) are independent of the results for the other wells. Synergies can be easily accommodated in this framework by allowing the rewards to depend on the outcomes of other wells and considering nonconstant reward functions $s_i(\boldsymbol{\omega})$ and $f_i(\boldsymbol{\omega})$ in Equations (8) or (9). For example, if we would not build the production and transportation until all wells have been drilled, we can take the rewards associated with a successful well $s_i(\boldsymbol{\omega})$ to be zero, except for those scenarios when five wells have already been drilled. Alternatively, we could accomplish the same thing by taking all the rewards to be zero and taking the terminal value $v(\boldsymbol{\omega})$ to represent the value of the portfolio of wells in the case where all wells have been drilled (i.e., $\boldsymbol{\omega}$ is a vector of zeros and ones) rather than assuming $v(\boldsymbol{\omega}) = 0$ in this case. This extension does not complicate the dynamic programming solution procedure.

5. Accuracy of the Distribution Approximation

It is natural to ask how well our KL-minimization approach approximates the "true" distributions that are likely to be encountered in practice. In thinking about the accuracy of this approach, it is important to remember that in practice, as in our example, the true distribution will typically never be known. Indeed, the purpose of the process is to construct a reasonable joint distribution from a limited number of assessments: If the true distribution were known, we could simply use it. Nevertheless, we would like to have some assurance that the joint distributions that we construct are reasonable and that the method provides good approximations across the range of joint distributions that might arise in practice. In this section, we will present results of some preliminary tests of accuracy using the simulation procedure proposed by Keefer (2004) to generate joint probability distributions that have positive pairwise correlations; this simulation algorithm is described in detail in the appendix. Abbas (2006) also uses simulation to evaluate the accuracy of maximum entropy approximations based on marginal and pairwise (also three-way) constraints but does not consider dynamic decision problems like the one considered here. We use Keefer's simulation procedure because it is designed to generate joint distributions that are realistic for positively related exploration prospects.

To provide a benchmark for evaluating the accuracy of the KL approximation, we will compare the results of the KL approximation to those given by an *independent approximation* that starts with marginal probabilities for each event and assumes that the events are independent. Although such an approach will not capture learning, in practice dependence is often overlooked or not captured, and it is important to understand the nature and magnitude of the errors introduced by neglecting it. In our six-well example, described in §2, the intrinsic value of each well is negative, and therefore, the independent approximation would suggest that we do not drill any wells and yields a value of $0, compared to the $14.40 million given by the KL approximation.

An alternative model that is simpler than the KL approximation but captures some degree of dependence is Keefer's (2004) *underlying event* (UE) approximation. The UE approximation posits the existence of a hypothetical UE that must occur in order for any wells to be wet and such that, given the occurrence of the underlying event, all the wells are independent. Keefer calibrates this model by assessing the

marginal probabilities of success for all events and assessing one conditional probability. Given the structural assumption of the UE model, this is sufficient to fully determine the joint distribution (see Keefer 2004 for details). This simplifying structural assumption has some strong consequences for learning: As Keefer notes, once one well is wet, you know that this UE has occurred and the remaining wells are probabilistically independent. In this scenario, learning stops and the optimal drilling strategies call for drilling the same set of wells regardless of how the remaining wells turn out. In general, the UE approximation may overestimate or underestimate the value in a sequential exploration problem. In our six-well example, the UE approximation suggests not drilling at all and therefore suggests a zero value.

In our simulations, we will focus on cases with six wells and will consider three measures of the accuracy of these approximations of the joint distribution: (i) the mean absolute difference between the approximate and true joint probabilities, (ii) the maximum absolute difference in these joint probabilities, and (iii) the KL distance from the true distribution to the approximation. The KL distance is given by Equation (1) but with the true distribution in place of the reference prior $\pi_0(\boldsymbol{\omega})$. These error measures are all such that a value of zero indicates no error and larger values indicate larger errors. Measures (i) and (ii) were used in Keefer's analysis; Abbas considers (i), (ii), and (iii).

In addition to these probability-based measures of accuracy, we will consider two economic measures in the sequential exploration setting: (iv) the absolute value of the difference between the optimal expected value given by the true probabilities and the optimal expected value given by the approximate probabilities and (v) the difference between the expected value given by following the true optimal policy and the expected value by following the approximate optimal policy, with both expected values calculated using the true joint probabilities: We refer to this as the expected value lost. The error in the value estimate (measure iv) is an appropriate error measure if the analysis is intended to estimate the value of the exploration opportunity, for example, if contemplating selling or acquiring the exploration opportunity. The expected value lost (measure v) indicates how close you come to identifying the optimal exploration strategy.

In each step of our simulation, we randomly generate a true joint probability distribution using Keefer's procedure. We then use this true joint distribution to calculate the marginal and pairwise conditional or joint probabilities required for each of the approximations and construct the corresponding approximate joint distributions. Next, we solve the dynamic programming model of §4 for the true joint probability distribution and for each of the approximations. Finally, we calculate and record the five different error measures described above. Although we randomly generate the joint probability distribution, in our study, we always use the expected values of success and failure (s_i and f_i) assumed in Table 1.

The results of a simulation involving 5,000 trials are shown in Table 6. The rows in the table correspond to the different measures of accuracy and the columns to the different approximations. For each error measure, we report the mean error, the standard deviation, and the fraction of times the simpler method outperformed the KL approximation. To place these numerical results in context, the average joint probability in these models must always be $1/2^6 = 0.0156$. The true

Table 6 Simulation Results

	Approximation		
	Independent	UE	KL
i. Mean absolute difference			
Mean	0.0148	0.0095	0.0017
Std dev	0.0037	0.0026	0.0007
Fraction < KL (%)	0.00	0.00	—
ii. Max absolute difference			
Mean	0.2471	0.1326	0.0115
Std dev	0.0720	0.0568	0.0060
Fraction < KL (%)	0.00	0.06	—
iii. KL distance from true distribution			
Mean	0.913	0.444	0.026
Std dev	0.443	0.208	0.016
Fraction < KL (%)	0.00	0.00	—
iv. Absolute error in value estimate ($ million)			
Mean	15.25	6.39	0.31
Std dev	8.89	4.77	0.29
Fraction < KL (%)	0.24	2.08	—
v. Lost value from approximate policy ($ million)			
Mean	15.25	5.57	0.13
Std dev	8.89	3.66	0.20
Fraction < KL (%)	0.06	0.50	—

values for the sequential exploration opportunities in the simulations averaged \$82.8 million, with a standard deviation of \$51.4 million; approximately 3% of these randomly generated distributions lead to a truly optimal value of zero.

Examining these results, we see that on all measures the mean errors for the KL approximation are much smaller than the corresponding errors for the other methods, typically by an order of magnitude. Moreover, examining the Fraction < KL results, we see there are very few cases where the simpler approximations outperform the KL approximation on any of these measures. The best Fraction < KL is the UE approximation's performance on the absolute value of the error in value estimates (measure iv), where the UE model approximation outperforms the KL approximation in 2% of the cases. However, even in these cases, the superior performance of the UE approximation is usually illusory: Though the UE value estimate may sometimes be closer to the true value (it is sometimes high and sometimes low), in only 0.50% of the cases considered does the policy suggested by the UE approximation actually perform better than that suggested by the KL approximation when evaluated by the true probabilities. In these scenarios, the difference in expected values averages only 0.15. Thus, in the very few cases where the UE approximation truly outperforms the KL approximation, the two are typically quite close.

Of course, it should not be too surprising that the KL approximation outperforms these other approximations; it requires substantially more information than these other methods. The results do show that the errors associated with the independent approximation—that is, assessing the n marginal probabilities (six probabilities in the example) and ignoring the dependence—can be substantial, with the average error in value estimate and value lost (these two measures are always equal for the independent approximation) being \$15.25 million, that is, 18.42% of the average true value of \$82.8 million. The UE model requires more information than the independent model (n marginals and one conditional probability; a total of seven probabilities in the example) and performance improves, but the errors are still fairly substantial, the average errors in value estimate and value lost being 7.71% and 6.73% of the average true value. The KL approximation requires n marginal and $(n-1)n/2$ pairwise conditional or joint probabilities (a total of 21 probabilities in the example) but puts this additional information to good use. Overall, the KL approximation performs quite well in these simulations: a mean error of 0.31 and mean value lost of 0.13, only 0.37% and 0.16% of the average true value. Although we find these results quite reassuring, this is a preliminary analysis of accuracy; a more complete analysis of the accuracy of these approximations would consider alternative value assumptions, different numbers of wells, alternative methods for generating true distributions, and the impact of errors in the assessments.

6. Discussion and Conclusion

The modeling tools developed in §§3 and 4 work well together, but each could be used independently of the other and in other contexts. For example, one could use this information-theoretic approach for generating joint probability distributions for a simulation model. Similarly, one could use the dynamic programming model with joint distributions generated some other way. For example, in the motivating consulting application, rather than assessing probabilities for "wet" and "dry" (hydrocarbons present or absent) for each well directly, the experts actually specified probabilities for the necessary conditions for hydrocarbons to be present: The hydrocarbons must have formed; the rocks must be porous enough to allow the hydrocarbons to flow into the well; the geologic formation must be of the appropriate shape, and the geologic formation must have a seal to trap the hydrocarbons. All four of these factors must be present for a well to be "wet." These factors were considered to be independent at each site, but the individual factors were dependent across sites. The experts assessed pairwise probabilities (as in §2) for each of these underlying factors, and the information-theory-based approach of §3 was used to construct a joint probability distribution for each factor. The joint distribution for the wells being wet or dry was then calculated from these joint distributions for the underlying factors.

Although this paper was motivated by a particular oil and gas example, the probability model developed here could be applied in other oil and gas applications, as well as in other contexts. For example,

rather than consider multiple wells, one might consider dependence among results for different targets or formations reached from a single well. In pharmaceutical R&D, one might use the probability model of §3 to capture dependence in the safety or efficacy of similar compounds. Alternatively, one might be interested in modeling dependence of default risks in a set of bilateral contracts with different counterparties. Moreover, though we have focused on the case where the events are binary (e.g., the well is either wet or dry), the information-theory-based procedure for determining joint probabilities can be generalized to problems with additional outcomes or additional uncertainties (e.g., learning about volumes or oil quality).

The dynamic programming model can also be extended in many ways. For example, we might consider the possibility of exploring in parallel as well as in sequence. The sequential approach maximizes the information available when you examine prospects, but this comes at the expense of delaying the potential benefits of a successful well. In an R&D context where the individual projects or compounds take longer to evaluate, parallel exploration strategies may be particularly attractive. The model of §4 can easily be extended to handle this: The state space would still consist of the same 3^n possible combinations of wet, dry, or unknown wells, but, if we are going to evaluate the possibility of drilling up to two wells in each period, we would need to consider up to $n^2 + n$ actions,[8] rather than the maximum of $n + 1$ actions in the current formulation. It is also straightforward to extend the model to consider the case where there is a bound on the total number of wells that can be drilled. We could also expand the state space to consider additional outcomes or uncertainties beyond the individual well outcomes. For example, in the oil exploration context we might consider the oil price or the "day rate" for the drilling rigs as an additional uncertainty that may affect values and drilling decisions. A more complex extension would consider dependence between the values given success (the s_i) as well as the success events themselves. The size of this state space in this case would be such that we would probably have to use approximate dynamic programming methods (e.g., Bertsekas and Tsitsiklis 1996) to solve these problems rather than the exact methods used here.

Though several generalizations are possible, our model illustrates both the challenges and rewards of dynamic decision modeling. The key challenges are constructing an appropriate yet tractable model that describes the uncertain "learning" over time and building a decision model that can efficiently contemplate the many scenarios and options that arise. Here we used information-theoretic methods to construct a joint probability distribution that captures the learning from drilling wells sequentially. We then used dynamic programming techniques to take advantage of the recombining structure of the problem and to keep the decision model reasonably tractable. These general techniques can be applied in many dynamic decision problems. The benefit of the dynamic modeling is apparent in our example: By being smart in choosing which wells to drill when, we were able to turn a collection of six individually unattractive wells into a promising exploration play. Here and in general, the optimal strategies are quite sensitive to the dependence in the learning model and are difficult to evaluate without the aid of a dynamic decision model.

An online supplement to this paper is available on the *Decision Analysis* website (http://da.pubs.informs.org/online-supp.html).

Acknowledgments

The authors thank three anonymous referees and editor Don Kleinmuntz for helpful comments and suggestions. Eric Bickel thanks Jennifer Meyer of Strategic Decisions for encouraging this research; she was the project manager for the consulting engagement that motivated this work. Eric also thanks David Lowell and Ali Abbas for helpful conversations about using maximum entropy methods in the original consulting application.

Appendix: Keefer's Procedure for Generating Joint Distributions

In §5 we use Keefer's (2004) procedure for generating probability distributions wells that exhibit positive pairwise dependence among binary variables. The procedure generates conditional probabilities for wells 1–6 in order. In each simulation:

[8] This is $(n+1)$ choices for the first well (n wells plus not drilling) times n choices for a second well ($n-1$ other wells plus not drilling a second well).

1. We generate a probability of success for well 1, $p(w_1 = 1)$, by drawing a random number from a uniform distribution on $(0, 1)$.

2. We use a two-stage process to generate conditional probabilities for well 2: We first generate a random p from the uniform on $(0, 1)$ and then generate $p(w_2 = 1 \mid w_1 = 0)$ by uniform sampling from $(0, p)$ and generate $p(w_2 = 1 \mid w_1 = 1)$ by drawing a random number from $(p, 1)$.

3. We also use a two-stage process for well 3: We generate two uniform random numbers, labeling the smaller one p_1 and the larger one p_2. We then assign $p(w_3 = 1 \mid w_2 = 0, w_1 = 0)$ a random number drawn from $(0, p_1)$, $p(w_3 = 1 \mid w_2 = 0, w_1 = 1)$ and $p(w_3 = 1 \mid w_2 = 1, w_1 = 0)$ random numbers drawn from (p_1, p_2), and $p(w_3 = 1 \mid w_2 = 1, w_1 = 1)$ a random number drawn from $(p_2, 1)$.

4. For wells 4–6, we generate conditional probabilities by sampling from a beta distribution and then sorting the probabilities so that more successes at early wells leads to a higher probability of success for the present well. The beta distribution's parameters α and β were determined by randomly drawing its mode, given by $(\alpha - 1)/(\alpha + \beta - 2)$, from a uniform distribution. The values of $\alpha + \beta$ were taken to be 2.1, 2.5, and 2.9 for wells 4–6.

Keefer chose these particular distributions and parameters so that the pairwise correlations would average close to 0.5. Our simulation is implemented in Excel and replicates Keefer's results for the pairwise correlations as well as the error results for the independent and UE approximations for the two error measures (i and ii in Table 5) that we have in common.

References

Abbas, Ali E. 2003. Entropy methods in decision analysis. Ph.D. dissertation, Stanford University, Stanford, CA.

Abbas, Ali E. 2006. Entropy methods for joint distributions in decision analysis. *IEEE Trans. Engrg. Management* **53**(1) 146–159.

Bertsekas, Dimitri, John N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA.

Clemen, Robert T., Terence Reilly. 1999. Correlations and copulas for decision and risk analysis. *Management Sci.* **45**(2) 208–224.

Clemen, Robert T., Gregory W. Fischer, Robert L. Winkler. 2000. Assessing dependence: Some experimental results. *Management Sci.* **46**(8) 1100–1115.

Cover, Thomas M., Joy A. Thomas. 1991. *Elements of Information Theory*. John Wiley & Sons, New York.

Gokhale, D. V., Solomon Kullback. 1978. *Information in Contingency Tables*. Marcel Dekker, Inc., New York.

Jaynes, Edwin T. 1968. Prior probabilities. *IEEE Trans. Systems Sci. Cybernetics* **SSC-4**(3) 227–241.

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

Keefer, Donald L. 2004. The underlying event model for approximating probabilistic dependence among binary events. *IEEE Trans. Engrg. Management* **51**(2) 173–182.

Keefer, Donald L., F. Beckley Smith Jr., Harry B. Back. 1991. Development and use of a modeling system to aid a major oil company in allocating bidding capital. *Oper. Res.* **39**(1) 28–41.

Lowell, David G. 1994. Sensitivity to relevance in decision analysis. Ph.D. dissertation, Stanford University, Stanford, CA.

Luenberger, David G. 1989. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, MA.

MacKenzie, George R. 1994. Approximately maximum-entropy multivariate distributions with specified marginals and pairwise correlations. Ph.D. thesis, University of Oregon, Eugene, OR.

Miller, Douglas J., Wei-han Liu. 2002. On the recovery of joint distributions from limited information. *J. Econometrics* **107** 259–274.

Moskowitz, Hebert, Rakesh K. Sarin. 1983. Improving the consistency of conditional probability assessments for forecasting and decision making. *Management Sci.* **29**(6) 735–749.

Neter, John, Michael H. Kutner, Christopher J. Nachtsheim, William Wasserman. 1996. *Applied Linear Statistical Models*. Irwin, Chicago, IL.

Newendorp, Paul D. 1975. *Decision Analysis for Petroleum Exploration*. Penwell Publishing, Tulsa, OK.

Paddock, James L., Daniel R. Siegel, James L. Smith. 1988. Option valuation of claims on real assets: The case of offshore petroleum leases. *Quart. J. Econom.* **103**(3) 479–508.

Ross, Sheldon M. 1983. *An Introduction to Stochastic Dynamic Programming*. Academic Press, San Diego, CA.

Smith, James E. 1995. Generalized Chebychev inequalities: Theory and applications in decision analysis. *Oper. Res.* **43**(5) 807–825.

Smith, James E., Kevin F. McCardle. 1998. Valuing oil properties: Integrating option pricing and decision analysis approaches. *Oper. Res.* **46**(2) 198–217.

Smith, James E., Kevin F. McCardle. 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. *Oper. Res.* **47**(1) 1–15.

Smith, James L., Rex Thompson. 2004. Managing a portfolio of real options: Sequential exploration of dependent prospects. Working paper, Department of Finance, Edwin L. Cox School of Business, Southern Methodist University, Dallas, TX.

Wang, B., G. P. Kokolis, W. J. Rapp, B. L. Litvak. 2000. Dependent risk calculations in multiple-prospect exploration evaluations. Paper SPE 63198, presented at the 2000 Society of Petroleum Engineers, Technical Conference and Exhibition, Dallas, TX.

Yi, Woojune, Vicki M. Bier. 1998. An application of copulas to accident precursor analysis. *Management Sci.* **44**(12S) 257–270.