

THE TWO-FACTOR PRICE PROCESS IN OPTIMAL SEQUENTIAL EXPLORATION

In a group of exploration prospects with common geological features, drilling a well reveals information about chances of success in others. With such dependencies, where do we drill first and what comes next given success or failure in previous wells? In addition, oil prices vary during the exploration campaign and with them so do the economics of wells and the optimal decision to drill. The grand problem then becomes a valuation of a group of prospects given geological dependencies and uncertain oil prices. Using a joint distribution for geological outcomes (a result of using information-theoretic methods) and a two-factor stochastic price process in a two-dimensional binomial sequence, we propose a Markov decision process that solves the optimal exploration problem. An Excel® VBA software implementation of this algorithm also accompanies this paper.

1. INTRODUCTION

Motivated by a petroleum exploration campaign in the Barents Sea, off the northern coasts of Norway, we came to revisit a solution to a prominent problem. When prospects are geologically dependent, what is the optimal sequence of drilling? A discovery or a dry hole in one well will affect the chances of success in the neighbouring prospects, so the exploration decisions should consider informational synergies between prospects. Furthermore, drilling in arctic waters and then evaluating the findings takes a long time, perhaps up to a year. By then the economic valuations for the upcoming wells expire and a new round of analysis completely changes the drilling policy. The management faces a new problem; variations in oil prices clearly change the optimal drilling policy, then what is the value of this exploration campaign and where should we drill first?

An optimal exploration policy should take into account the geological dependencies; yet considering the recent downturn in the markets that deeply affected the exploration business, a policy cannot ignore the possible future variations in prices. The unpredictable changes in prices make this problem like a restless bandit with multiple correlated arms. The state of the system changes not only because of the decision maker's actions, but also according to external, possibly random, factors. A solution to these problems is by solving the underlying Markov decision process (Puterman, 2014). In this paper, we apply a stochastic model that describes the dynamics of prices, and devise an aggregate algorithm for solving a moderate-sized problem of ten exploration prospects.

Similar problems, from developing pharmaceutical products to selecting R&D projects, benefit from a solution to the sequential exploration problem. When developing correlated compounds for products that arrive later in the market, or when selecting dependent projects with delayed outcomes, the decision makers deal with restless bandits with correlated arms. They face a trade-off between earning (drilling high-value wells) and learning (drilling prospects with most valuable information), but then values are not stationary; they change by the next decision epoch. We believe our valuation algorithm could also be useful for these other applications.

We rely on explicit definition of forward curves as the way of reflecting variability on project value. Contrary to the common belief, the value of discovering oil and gas is not a lump sum; it is the net present value of a stream of cash flows materializing perhaps years in the future. To show how project value varies with prices, we first need to show how the outlook of prices varies. In our model, we employ cash flow models for each well to estimate the effect of changing forward prices on the value of a discovery. Assuming oil prices follow the two-factor process of Schwartz and Smith (2000), we

apply the discrete binomial formulation of Hahn and Dyer (2012) to represent prices in our Markov decision process. For a group of exploration prospects, we do the following:

- We apply an information-theoretic approach, previously used in Bickel and Smith (2006), to generate a joint probability distribution incorporating marginal chances of success and geological dependencies.
- We model prices as a two-factor price process (Schwartz and Smith, 2000), and use the approach in Hahn and Dyer (2012) to construct dual-binomial lattices. The cash flow model takes the forward curves originating from each node of the lattice and estimates the value of discovery with respect to varying prices.
- We construct a Markov decision process representing the sequential exploration problem given the joint probabilities and binomial lattices. A recursive algorithm, incorporating transition probabilities and rewards from previous steps, returns the value of optimal sequential exploration.

In this paper, we build on the strand of literature describing sequential exploration. Bickel and Smith (2006) discussed optimal exploration of six prospects, with outcomes “dry” or “wet”, using a dynamic programming model. This led to $3^6 = 729$ states and the authors developed a spreadsheet model to handle the valuation. Bickel, Smith and Meyers (2008) extended the previous model to more intricate geological uncertainties, now three “layers” of uncertainty each could take “fail” or “success” states. Solving for five wells, their dynamic programming model had to handle around 59,000 states. Brown and Smith (2013) and Martinelli et al (2013) considered even larger problems, clusters of exploration prospects each containing many targets. They suggested approximate methods to address the curse of dimensionality. While these models are prevalent and insightful at early phases of screening, we believe the effect of well economics gains importance as decision makers proceed towards commitment to investments. Our valuation algorithm expands the model in Bickel and Smith (2006) and further includes the effect of stochastic prices on optimal decisions. We also provide a modular, open source software application that performs the valuation algorithm.

Our work also contributes to valuation of real options and applications of the two-factor price model in Schwartz and Smith (2000); a model realistic enough to reflect the dynamics of prices in the markets and simple enough to provide decision insights. In the oil and gas industry, Jafarizadeh and Bratvold (2013 and 2015) simulated the two-factor price process to evaluate real options. Yet in our discrete Markov decision process, simulation will be prohibitive; we need a finite number of discrete states for prices. Originally developed to approximate Geometric Brownian motions by Cox et al (1979) and recently adapted by Hahn and Dyer (2008) for mean-reverting processes, binomial lattices are promising provisions to our sequential exploration model. We use the dual-binomial lattice developed by Hahn and Dyer (2012) to approximate the two-factor price process.

In the next section, we discuss the details of constructing dual binomial lattices for prices. Then in section 3, we use these lattices along with the joint geological probability distribution in a recursive algorithm that solves the Markov decision process. This section also has a brief description of the information-theoretic approach to calculate the joint probability distribution. In section 4, we describe how we implemented the valuation algorithm in Excel VBA. Using the software, we solve a problem with ten exploration targets, and perform sensitivity analyses that support decisions (with more details about the software and example in appendices).

2. TWO-DIMNESIONAL LATTICE OF PRICES

In sequential exploration, we constantly compare the value from drilling with the value of learning; both depend on the price-driven expected future cash flows. Our models of price behaviour, whether a random walk process as in earlier studies, or a mean reverting model as in common industry's opinion and favoured by more recent studies, will have direct impact on the optimal sequential decisions. In this paper, we use the realistic, yet simple two-factor price process in Schwartz and Smith (2000) to describe the dynamics of prices.

The two-factor model assumes the prices are mean reverting and converge to a varying equilibrium price¹. The equilibrium, a random walk process itself, alters as a result of depleting resources or technological and political changes, while short-term deviations from this equilibrium, perhaps a result of temporary disruptions in supply, tend to disappear by time and follow an Ornstein-Uhlenbeck process. It is easier to think of this model as two nested processes for χ_t , the short-term factor, and ξ_t , the long-term factor, where the spot price, S_t is defined as $\ln S_t = \chi_t + \xi_t$ and

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dz_\chi \quad (1)$$

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \quad (2)$$

In the above equations, κ is the rate of mean-reversion, μ_ξ is the trend, σ_χ and σ_ξ are the volatilities for short- and long-term factors, and dz_χ and dz_ξ are the correlated increments of the standard Brownian motion processes with $dz_\chi dz_\xi = \rho_{\chi\xi} dt$.

For our Markov decision process, we need a representation of the above continuous diffusions in discrete sequences. The method in Hahn and Dyer (2012), based on a general approach in Nelson and Ramaswamy (1990), generates a two-dimensional binomial sequence that is simple and general for our purpose, and is not limited by grid size or the range of parameter values². In this formulation, four branches originate from each node at each discrete period, two states for each price factor, resulting in a dual binomial recombining lattice.

Assuming log of price at a specific node is $\xi_t + \chi_t$, the four states (and transition probabilities to each state) for the next epoch will be

State for $\ln S_t$	Probability	
$\xi_t + \Delta_\xi + \chi_t + \Delta_\chi$	p_{uu}	(3)
$\xi_t + \Delta_\xi + \chi_t - \Delta_\chi$	p_{ud}	(4)
$\xi_t - \Delta_\xi + \chi_t + \Delta_\chi$	p_{du}	(5)
$\xi_t - \Delta_\xi + \chi_t - \Delta_\chi$	p_{dd}	(6)

Where the increment for each factor is

$$\Delta_\xi = \sigma_\xi \sqrt{\Delta t} \quad (7)$$

$$\Delta_\chi = \sigma_\chi \sqrt{\Delta t} \quad (8)$$

¹ This model is in fact exactly equivalent to the stochastic convenience yield model of Gibson and Schwartz (1990). Mean-reversion as an appropriate assumption for commodities is discussed in e.g. Laughton and Jacoby (1993), Cortazar and Schwartz (1994), and Dixit and Pindyck (1994). Schwartz (1997) discusses mean-reversion in stochastic price models and their ability to price existing future contracts, as well as financial and real assets.

² Previous attempts to discretize two-factor price diffusions, notably the dual trinomial lattice approach in Hull and White (1994) or the improved version of Tseng and Lin (2007), worked only under a specific range of correlation values and had computational limitations.

The probabilities of moving to each state, p_{uu} to p_{dd} , are easier to calculate if we consider them as joint probabilities, the product of marginal probability of a move in ξ_t and a conditional probability of a move in χ_t . For example, $p_{uu} = p_u \times p_{u|u}$ where p_u is the marginal probability of “up” move in the long-term factor and $p_{u|u}$ is the conditional probability of “up” move in the short-term factor. The marginal and conditional probabilities for four transitions are

$$p_u = \frac{1}{2} + \frac{\mu_\xi \Delta t}{2\Delta_\xi} \quad (9)$$

$$p_d = 1 - p_u \quad (10)$$

$$p_{u|u} = \frac{\Delta_\xi(\Delta_\chi + \Delta t v_\chi) + \Delta t(\Delta_\chi \mu_\xi + \rho \sigma_\xi \sigma_\chi)}{2\Delta_\chi(\Delta_\xi + \Delta t \mu_\xi)} \quad (11)$$

$$p_{d|u} = \frac{\Delta_\xi(\Delta_\chi - \Delta t v_\chi) + \Delta t(\Delta_\chi \mu_\xi - \rho \sigma_\xi \sigma_\chi)}{2\Delta_\chi(\Delta_\xi + \Delta t \mu_\xi)} \quad (12)$$

$$p_{u|d} = \frac{\Delta_\xi(\Delta_\chi - \Delta t v_\chi) - \Delta t(\Delta_\chi \mu_\xi - \rho \sigma_\xi \sigma_\chi)}{2\Delta_\chi(\Delta_\xi + \Delta t \mu_\xi)} \quad (13)$$

$$p_{d|d} = \frac{\Delta_\xi(\Delta_\chi + \Delta t v_\chi) - \Delta t(\Delta_\chi \mu_\xi + \rho \sigma_\xi \sigma_\chi)}{2\Delta_\chi(\Delta_\xi + \Delta t \mu_\xi)} \quad (14)$$

We assumed $v_\chi = -\kappa \chi_t$ to simplify the equations. Furthermore, because equations (11) to (14) sometimes generate unacceptable results, we bound the probabilities for short-term factor between zero and one using the equation $p_{bounded} = \max(0, \min(1, p_{un-bounded}))$.

This will result in a dual binomial lattice for evolution of spot prices. At each node, the short- and long-term factors each can have an “up” or “down” tick, resulting in two connected binomial sequences. Although more comprehensible in three dimensional plots, we can still show the results in the lattice of figure 1 assuming four branches originate from each node and using parameter values in table 1. Here for example $S_1^{++} = e^{\xi_0 + \Delta_\xi + \chi_0 + \Delta_\chi}$ represents a move in spot price from $t = 0$ to $t = 1$ where both the short- and long-term factors have up ticks. The quadrinomial lattice shown in black solid lines generates four price states at $t = 1$ and nine at $t = 2$.

Table 1 Parameter values for the two-factor price process

Parameter of the process	Value
χ_0	0.15
κ	0.8
σ_χ	15%
ξ_0	4.1
μ_ξ	0
σ_ξ	20%
$\rho_{\xi\chi}$	0.3

Spot prices are often irrelevant to the economics of exploration decisions as expected cash flows of these projects appear years into the future. Instead, we are interested in the information that spot prices provide about the future trends. Forward curves provide such information by showing a riskless expectation of future price trends, and as in Jafarizadeh and Bratvold (2013 and 2015), can be theoretically reconstructed at each node of the lattice using the assumptions from the two-factor process. The relationship between forward prices and parameters of the process is

$$\ln F_{t,T} = r^{-\kappa(T-t)} \chi_t + \xi_t + \mu_\xi(T-t) + \frac{1}{2} \left((1 - e^{-2\kappa(T-t)}) \frac{\sigma_\chi^2}{2\kappa} + \sigma_\xi^2(T-t) + 2(1 - e^{-\kappa(T-t)}) \frac{\rho_{\xi\chi} \sigma_\xi \sigma_\chi}{\kappa} \right) \quad (15)$$




Figure 1 Dual binomial lattice showing the dynamics of spot prices (dark solid lines) and resulting forward curves (dashed lines). The line in solid red is the forward curve fitted to the observed forward prices at $t = 0$.

Originating from each node of the lattice in figure 1, the dashed lines in red represent the theoretical forward curves. When $t = 0$, the theoretical forward curve (shown in solid red) fits the observed forward prices in the market. Later on, as spot prices vary in the lattice, so do the corresponding forward curves. By $t = 2$ we will have a variety of curves from contango, e.g. the curve originating from the node S_2^{+++} , to normal backwardation as in the curve originating from S_2^{---} .

3. VALUATION ALGORITHM

The sequential exploration problem resembles a restless multi-armed bandit with dependent arms (Puterman, 2014). Drilling each well provides information about chance of geological success in other locations. Yet, economic success is a matter of success in geology and a good price outlook. Wells that seem economically viable now may become uneconomical by the end of period because the forward curve moved to an unfavourable position. The solution is to make drilling decisions by considering both the geological learning and the stochastic property of prices.

The next section discusses a method of integrating geological dependencies in the decision model. We then incorporate these dependencies along with dual binomial sequence of prices into a Markov decision process for the grand problem.

3.1. A Joint Probability Distribution for Geological Success

We use the Kullback-Leibler method of relative entropy, introduced by Jaynes (1968) and applied in Bickel and Smith (2006), to generate a joint probability distribution of geological outcomes. If $\mathbf{w} = (w_1, \dots, w_{10})$ is a vector of ten binary random variables for our wells (e.g. $w_5 = 1$ represents discovery in well 5), then we would like to construct a joint probability distribution $\pi(\mathbf{w})$ that reflects all the information at hand. The information-theoretic method utilizes the available information (chances of success for wells and geological correlations) and maximizes the relative entropy between $\pi(\mathbf{w})$ and a reference distribution $\pi_0(\mathbf{w})$ that assumes independence. The result of this optimization $\pi^*(\mathbf{w})$ is a joint probability distribution that manifests the individual chances of success along with geological dependencies.

To calculate $\pi^*(\mathbf{w})$ we need to solve a large optimization problem with a large set of constraints³ that may not be manageable in spreadsheet solvers. Instead, Bickel and Smith (2006) offer a workaround; a simplification based on the Lagrangian dual of the problem. If the marginal chance of success for each well is $p_i \equiv p(w_i = 1)$ and the joint pairwise probabilities are $p_{ij} \equiv p(w_i = 1, w_j = 1)$ then the Lagrangian dual of the optimization problem is

$$\max_{\lambda} \left(- \sum_{\mathbf{w}} \pi^*(\mathbf{w}, \lambda) + \lambda_0 + \sum_i \lambda_i p_i + \sum_{i,j} \lambda_{ij} p_{ij} \right) \quad (16)$$

Where λ_0 is the unit multiplier, λ_i and λ_{ij} are the Lagrangian multipliers associated with p_i and p_{ij} , and the vector λ represents all these elements. This problem has only $1 + n + n(n - 1)/2$ variables and no constraints. For our ten-well problem, the automated Excel® Solver in VBA reaches a solution for this optimization within a few seconds.

3.2. A Markov Decision Process

Assuming we have specified the geological probability distribution and dual-binomial price ticks, determining optimal drilling strategy may seem straightforward. As in figure 2, the decision to drill well 5 at the outset of the exploration program depends on its probability of success, the probability of price moves in the next period, and expected conditional value of drilling the remaining wells over the next nine decision periods. This decision tree, however, turns out to become unmanageably large in its complete form, with almost four quadrillion end-nodes⁴.

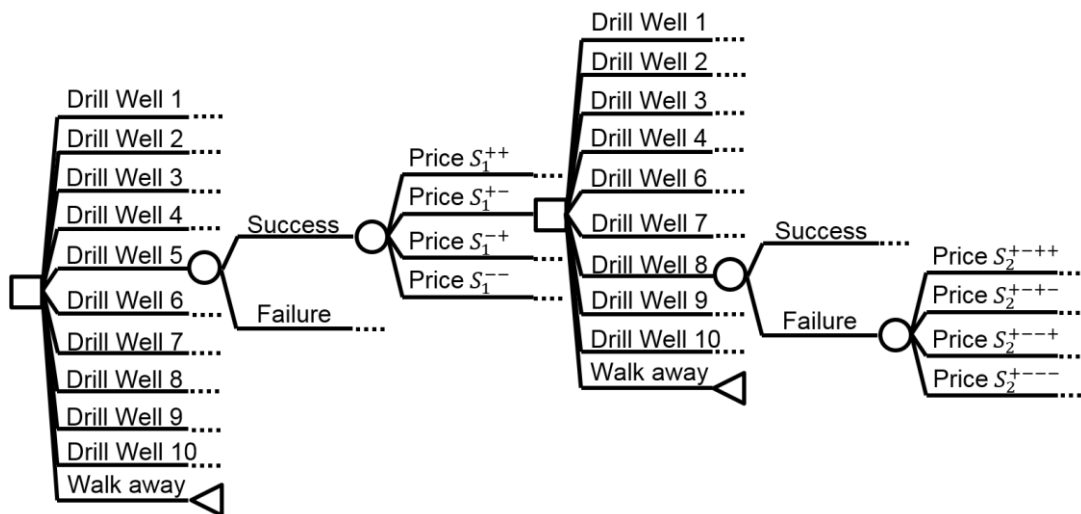


Figure 2 A partial decision tree showing the decisions and uncertainties in a sequential exploration problem

We can simplify this decision model by considering the fact that multiple end-nodes associate with identical information and future cash flows. For example, if we were successful in well 5 and then failed when we drilled well 8, the future geological probabilities are going to be the same regardless of the order in which we drilled the wells. The price at that point in time is also independent of its historical path; the forward curve we use to evaluate cash flows will only be a function of spot price. We describe the Markov decision model that draws on this recombining feature.

³ For our ten well application, there will be 1024 unknown joint probabilities ($2^n, n = 10$) and 56 constraints ($1 + n + n(n - 1)/2, n = 10$). Bickel and Smith (2006) explain the details of Kullback-Leibler procedure and its Lagrangian dual.

⁴ In the first decision epoch, we have 10 alternative wells each with 2 geologic outcomes and 4 price moves. The next epoch has 9 wells, each with 8 outcomes. Continuing this trend we will have $10! \times 8^{10}$ total outcomes.

We use a recursive algorithm, similar to the logic of solving a decision tree, to infer the optimal drilling decisions. In the final decision epoch, all wells except for one have been drilled. The decision to drill this last well depends on its conditional probability of success given the previous outcomes as well as its expected cash flows given the four prevailing forward curves at that point. After we determined the optimal decision for the scenarios of last epoch, we move backwards and calculate the optimal action in previous epochs. The transition probabilities (the probability of moving from one state to another at each decision epoch) will be composed of conditional geological probabilities and probabilities for price ticks.

To describe the state of wells at each decision epoch, we define $\boldsymbol{\omega} = (\omega_1, \dots, \omega_i, \dots, \omega_{10})$, where $\omega_i = "0", "1" \text{ or } "-"$. Here, "0" means "failure", "1" means "success", and "-" represents the case where the well has not been drilled yet. With this notation, $\boldsymbol{\omega} = (-, 0, 1, -, \dots, -)$ for example represents the state where we have drilled well 2 and 3, well 2 was a dry hole and well 3 was a success. Also, as we can only drill one well per epoch, the available eight alternatives at this state are all the wells except well 2 and 3.

The recursive algorithm selects the well that yields the highest expected value given the conditional chance of success and price levels. In fact, the algorithm looks beyond the immediate drilling results and considers the expected payoff that follows consequent to this drilling decision—we refer to this as the continuation value for price S and denote it by $v^S(\boldsymbol{\omega})$.

If we are in state $\boldsymbol{\omega}$ and well i is not drilled yet (therefore $\omega_i = "-"$), the expected value for well i , the NPV of future cash flows given that we observe price S and drill well i , is denoted by $v_i^S(\boldsymbol{\omega})$. This value depends on the expected future value of discovery at the current price level (d_i^S), failure (f_i), and the continuation value at the next period's four price levels. In other words

$$v_i^S(\boldsymbol{\omega}) = P(\boldsymbol{\omega}_i^1 | \boldsymbol{\omega}) \left(d_i^S + \delta \left(p_{uu} v^{S++}(\boldsymbol{\omega}_i^1) + p_{ud} v^{S+-}(\boldsymbol{\omega}_i^1) + p_{du} v^{S-+}(\boldsymbol{\omega}_i^1) + p_{dd} v^{S--}(\boldsymbol{\omega}_i^1) \right) \right) \\ + P(\boldsymbol{\omega}_i^0 | \boldsymbol{\omega}) \left(f_i + \delta \left(p_{uu} v^{S++}(\boldsymbol{\omega}_i^1) + p_{ud} v^{S+-}(\boldsymbol{\omega}_i^1) + p_{du} v^{S-+}(\boldsymbol{\omega}_i^1) + p_{dd} v^{S--}(\boldsymbol{\omega}_i^1) \right) \right) \quad (17)$$

Moreover, the continuation value for state $\boldsymbol{\omega}$ is the maximum expected value for all drilling alternatives

$$v^S(\boldsymbol{\omega}) = \max_i \left(0, v_i^S(\boldsymbol{\omega}) \right) \quad \text{for all un-drilled } i \quad (18)$$

In the above equations $\delta = \frac{1}{(1+r)^{\Delta t}}$ is the discount factor for rate r and the length of time between decision epochs Δt . Also, $\boldsymbol{\omega}_i^1$ is identical to $\boldsymbol{\omega}$ except that $w_i = 1$, and $\boldsymbol{\omega}_i^0$ is identical to $\boldsymbol{\omega}$ except that $w_i = 0$. Here $v^{S++}(\boldsymbol{\omega}_i^1)$ for example refers to the optimal value of $\boldsymbol{\omega}_i^1$ if both short- and long-term factors of prices go up in the next period.

We determine the geological transition probabilities $P(\boldsymbol{\omega}_i^1 | \boldsymbol{\omega})$ and $P(\boldsymbol{\omega}_i^0 | \boldsymbol{\omega})$ using our knowledge of the total probability distribution $\pi^*(\boldsymbol{w})$ calculated in the previous section. If $P(\boldsymbol{\omega})$ is the probability of state $\boldsymbol{\omega}$, it is calculated by summing up $\pi^*(\boldsymbol{w})$ over all possible scenarios for the unknown events. For example if $\boldsymbol{\omega} = (-, 0, 1, -, -, 0, 0, 1, 1, 1)$, then

$$P(\boldsymbol{\omega}) = \sum_{w_1, w_4, w_5} \pi(w_1, 0, 1, w_4, w_5, 0, 0, 1, 1, 1) \quad (19)$$

Where w , w_4 and w_5 range over $\{0, 1\}$. The probability of success and failure for well i conditional on the state ω (The transition probabilities, where well i has not yet been drilled and $\omega_i = \text{“--”}$) would be $\text{Prob}(\omega_i^1|\omega) = \text{Prob}(\omega_i^1)/\text{Prob}(\omega)$ and $\text{Prob}(\omega_i^0|\omega) = \text{Prob}(\omega_i^0)/\text{Prob}(\omega)$, respectively.

3.3. Calculating Immediate Cost and Reward

The recursive algorithm of the previous section assumes we receive the immediate reward d_i^S if drilling well i leads to discovery, or would have to pay the immediate cost $-f_i$ if our action leads to a dry hole. Assuming immediate realization of these values is in fact inaccurate in the context of sequential petroleum exploration as any loss or benefit (specifically benefit) will appear years into the future. For example, following discovery of oil or gas, appraisal and field development activities may take considerable time before production revenue starts. For this reason, in our formulation we use the then-expected net present value of discovery or dry hole.

Using the prevailing forward curve at the time of discovery to estimate production cash flows, and discounting the net cash flows (after deductions) to the time of discovery using a risk-free discount rate, we perform valuations in the risk-neutral paradigm. This approach to valuation, as discussed in Smith and Nau (1995) and implemented for exploration projects in Jafarizadeh (2017), conforms to the assumptions of the price process and the Markov decision model, and generates consistent valuations. Assuming $F_{t,\tau}^{\xi_t,\chi_t}$, $0 < \tau$ is the forward price with maturity τ originating from a node of the lattice with price elements ξ_t and χ_t , q_i^τ and c_i^τ are respectively the production and cost for well i during period τ , and Δt the time granularity of the problem

$$d_i^S = \sum_{\tau=0}^n \frac{F_{t,\tau}^{\xi_t,\chi_t} q_i^\tau - c_i^\tau}{(1+r)^\tau} \quad (20)$$

In the above equation, r is the risk-free interest rate and n is the length of project given discovery.

4. APPLICATION AND RESULTS

The algorithm in section 3 is beyond manual calculations and even laborious in spreadsheets, yet it may be efficiently implemented using array programming⁵. Our Excel VBA program uses large three-dimensional arrays to enumerate the states in the price lattice and Markov decision model. The program performs the optimization of the information-theoretic method using Excel Solver, and has built-in functions and subroutines to generate forward curves and calculate the associated net present values. Finally, the recursive algorithm processes on the elements of the arrays and generates the results. Appendix B explains the modules and their relationships in this program.

4.1. Implementation in Excel® VBA

The software follows the methodology we discussed in sections 2 and 3. It constructs three-dimensional arrays for χ_t and ξ_t , specifically, $\chi(10,10,10)$ and $\xi(10,10,10)$ where the first dimension shows the time steps, the second shows the changes in short-term factor, and the third

⁵ Although this application seems ideal for array programming languages, it is arguably not the ideal choice for users in academia and industry. For example, MATLAB® implementations of valuation algorithms (e.g. exploration waiting option in Jafarizadeh and Bratvold, 2015) are hampered by prohibitively high software license fees and scarce availability of programming skills. Yet Microsoft Excel is perhaps the platform of choice for small and medium scale analysis tasks in industry and dissemination of an open-source VBA application may be more beneficial.

shows the changes in long-term factor. The sum of these two arrays will reveal the log of prices for each state and each time-step. We then construct a forward curve for each element of this array and calculate d_i^S (NPV given discovery) in an array $\mathbf{d}_i(10,10,10)$ for each well i .

Similar arrays store information about joint and conditional probabilities to be used in the Kullback-Leibler optimization procedure. In the end, after enumerating states, the recursive algorithm culminates in processing an array of $\mathbf{V}(59049,10,10)$ to store $v^S(\omega)$ and similar sized arrays to store $v_i^S(\omega)$ for each state of each well. The first dimension of these arrays shows the total number of scenarios for ten wells when each can take three states “0”, “1” or “-”, a total of 3^{10} scenarios. The second and third dimensions show the variations in value due to variations in short- and long-term factors of prices. Applying the recursive algorithm on these arrays takes the bulk of processing time; a ten-well example takes approximately three to five minutes of CPU time to evaluate. Appendix B provides more details on the structure of this software.

A feature of this modularized program is its applicability and versatility. This open source code is composed of general modules; each accomplish a specific task and can effectively operate in other contexts. To show how this works, we organized the modules, with minimal programming effort, into sensitivity analysis subroutines. The next section shows how the program solves a complex exploration and how these subroutines can produce useful decision insights.

4.2 Example

We return to our arctic-circle exploration problem; a large, multi-prospect play with subsurface dependencies that requires a long time span to explore. The management believes that, with available resources, drilling a well and then analysis and interpretation of results will take at least a year. Assuming all wells require the same amount of resources, a complete exploration of this region would perhaps take a decade. In the meantime, dramatic variations in prices could sway the optimal policy. Hence, although at prevailing price projections the prospects are marginally uneconomical, the decision makers are interested in the expected value that an aggregate optimal exploration policy would bring about.

Applying our valuation algorithm along with the assumptions about price process and geological correlations (with more details in Appendix A) reveals that the optimal sequential exploration of this region would have a significant positive expected value. In other words, while each prospect was not economically viable in isolation, a sequential drilling that considers both prices and geological learning would become a sound investment. Considering the value of information and price option makes all the difference.

Figure 3 shows the valuation in three different versions; first, ignoring geological dependencies and variability of prices leads to the value of zero. Next, we assume prices vary. When prices follow a two-factor process, the expected value of the exploration play becomes positive. Finally, we include geological dependencies and notice another increase in valuation.

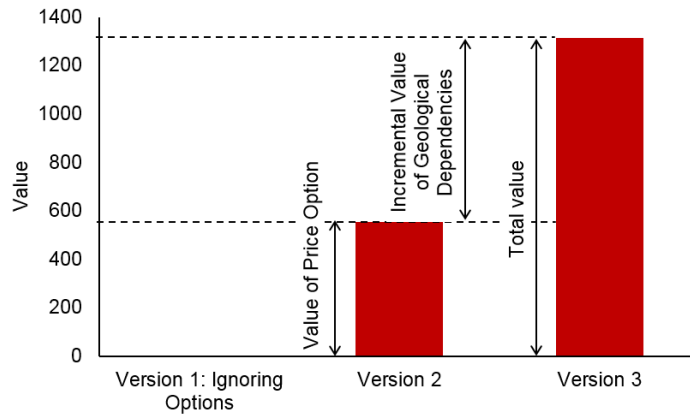


Figure 3 Expected value of sequential exploration. Although values are not additive, in this example geological dependencies and variability of oil prices almost equally drove the total value

Using the sensitivity analysis subroutine, we can take our valuation further by showing how key factors affect the value of sequential exploration. The univariate sensitivity analyses in figure 4 reveal that, for example, everything else unchanged, higher discount rates result in significantly reduced expected values.




Figure 4 One-way sensitivity analysis graphs

For correlated inputs, perhaps multi-way sensitivities are more insightful. In a two-way analysis of price volatility (figure 5), we notice that value is much more sensitive to σ_{ξ} , the volatility in the long-term factor. This is perhaps because of the long-term nature of investing in sequential exploration. With one-year intervals between drilling, the campaign takes almost ten years to conclude. Furthermore, production revenue of any discovery will take years to materialize.




Figure 5 Sensitivity analysis of value with respect to volatility in long- and short-term price factors

In a more comprehensive analysis, we can even gain insight on the interlinked nature of price options and geological learning. Effectively a large number of parameters influence the expected value: these

include individual project parameters, development solution given discovery, the price process and shape of the forward curves, and the configuration of prospect and their pairwise correlation. While evaluating the effect of such large variable set is prohibitive, we could still identify key elements and examine their effect.

In general, we expect higher price volatility and stronger geologic correlations to generate higher values. So what is the minimum volatility that makes a group of exploration targets (with a common correlation, ρ) valuable? In other words, we are looking for break-even volatility given various levels of common correlation. We can run the valuation algorithm for this group of targets and vary the price volatilities σ_χ and σ_ξ (while keeping all other parameters fixed) until we reach $v^{S_0}(\omega) = 0$. This would be the minimum σ_χ or σ_ξ to have a positive value for sequential exploration. A simple goal-seeking routine expedites the process. Figure 6 shows combinations of break-even volatility and common geologic correlation at specific spot price scenarios.




Figure 6 Sensitivity analysis of value with respect to price and average geological correlation; each contour line represents combinations of σ_ξ and ρ (solid lines) or σ_χ and ρ (dashed lines) that make $v^{S_0}(\omega) = 0$ at a specific spot price.

In addition to a value estimate, we can also determine the starting well in this optimal sequence of decisions. Because in this valuation, the interactions between geological learning and stochastic prices are intricate, a complete strategy map would neither be feasible nor beneficial for decision makers. However, running the software at any point in time and learning about the next well in the optimal strategy would be enough to make value-maximizing decisions.

5. CONCLUSIONS

This paper provides an algorithmic solution to the Markov decision process of sequential exploration. We combine the simplicity of binomial lattices with the power of recursive solutions and implement our method in an effective computer application. Furthermore, we solve a problem of sequential exploration consisting of ten wells and show how sensitivity analyses can provide deeper insights into exploration decisions. We integrate all the modelling tools in a single package, but note that each can also work independently; for example, the subroutine for binomial lattice is also suitable for valuation of commodity options, the embedded functions for forward curve work elsewhere within the spreadsheet, and the recursion subroutine can be adopted in other restless bandit problems.

Our valuation model is not limited to applications in the oil and gas industry. Comparable problems, for example developing drugs from common compounds in the pharmaceutical industry or sequential R&D projects, have similar characteristics. In these problems, once the relationship between value of a project and market uncertainties are understood, the applications of the model is effortless.

Finally, we note an extension of the problem that can readily use this evaluation framework. In some contexts, individual discoveries may be too small to justify a development solution. Success in two or more wells could be bundled in a project that uses common production and export facilities to drain this cluster of discoveries. We can adjust the rewards for our Markov decision process and handle these functional synergies. The modifications of the code should also be straightforward.

The Excel® VBA software is available at the following link

https://www.dropbox.com/s/mrtr5tk17ow2khr/Sequential_Exploration%202.1.xlsm?dl=0

REFERENCES

- Bickel, J. E., & Smith, J. E. (2006). Optimal sequential exploration: A binary learning model. *Decision Analysis*, 3(1), 16-32.
- Bickel, J. E., Smith, J. E., & Meyer, J. L. (2008). Modeling dependence among geologic risks in sequential exploration decisions. *SPE Reservoir Evaluation & Engineering*, 11(02), 352-361.
- Brown, D. B., & Smith, J. E. (2013). Optimal sequential exploration: Bandits, clairvoyants, and wildcats. *Operations research*, 61(3), 644-665.
- Cortazar, G., & Schwartz, E. S. (1994). The valuation of commodity contingent claims. *Journal of Derivatives*, 1(4), 27-39.
- Cox, J. C., Ross, S. A., & Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of financial Economics*, 7(3), 229-263.
- Dixit, A. K., Dixit, R. K., Pindyck, R. S., & Pindyck, R. (1994). *Investment under uncertainty*. Princeton university press.
- Hahn, W. J., & Dyer, J. S. (2008). Discrete time modeling of mean-reverting stochastic processes for real option valuation. *European journal of operational research*, 184(2), 534-548.
- Hahn, W. J., & Dyer, J. S. (2011). A discrete time approach for modeling two-factor mean-reverting stochastic processes. *Decision Analysis*, 8(3), 220-232.
- Hull, J., & White A. (1994). Numerical procedures for implementing term structure models II: Two-factor models. *Journal of Derivatives*, 2(2), 37-48.
- Jafarizadeh, B., & Bratvold, R. (2012). Two-factor oil-price model and real option valuation: an example of oilfield abandonment. *SPE Economics & Management*, 4(03), 158-170.
- Jafarizadeh, B., & Bratvold, R. B. (2015). Oil and gas exploration valuation and the value of waiting. *The Engineering Economist*, 60(4), 245-262.
- Jaynes, E. T. (1968). Prior probabilities. *IEEE Transactions on systems science and cybernetics*, 4(3), 227-241.
- Laughton, D. G., & Jacoby, H. D. (1993). Reversion, timing options, and long-term decision-making. *Financial Management*, 225-240.

Martinelli, G., Eidsvik, J., & Hauge, R. (2013). Dynamic decision making for graphical models applied to oil exploration. *European Journal of Operational Research*, 230(3), 688-702.

Nelson, D. B., & Ramaswamy, K. (1990). Simple binomial processes as diffusion approximations in financial models. *The review of financial studies*, 3(3), 393-430.

Puterman, M. L. (2014). *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

Schwartz, E. S. (1997). The stochastic behavior of commodity prices: Implications for valuation and hedging. *The journal of finance*, 52(3), 923-973.

Schwartz, E., & Smith, J. E. (2000). Short-term variations and long-term dynamics in commodity prices. *Management Science*, 46(7), 893-911.

Tseng, C. L., & Lin, K. Y. (2007). A framework using two-factor price lattices for generation asset valuation. *Operations Research*, 55(2), 234-251.

APPENDIX A: PETROLEUM EXPLORATION VALUATION

This appendix provides more details on the example of sequential exploration. Assume we have identified ten exploration prospects in a frontier region. Each well could be a “discovery” or “dry hole” and its outcome will likely affect the probabilities in other prospects. The value of a well given “dry hole” is the present value of its drilling cost, while value given “discovery” is the net present value of all its costs and expected production revenue. Table 2 shows the information about each well at current price levels

Table 2 information about exploration wells in the example

	Chance of Success	Value given “Success”	Value given “dry”	Expected Value
Well 1	0.25	859	- 300	- 10
Well 2	0.2	1596	- 400	- 1
Well 3	0.25	752	- 300	- 37
Well 4	0.3	795	- 350	- 6
Well 5	0.2	1585	- 400	- 3
Well 6	0.15	2065	- 400	- 30
Well 7	0.3	882	- 400	- 15
Well 8	0.25	882	- 350	- 42
Well 9	0.3	834	- 400	- 30
Well 10	0.33	749	- 400	- 21

These wells are not attractive in isolation. Their negative expected value shows that with current level of information, they will not create value. However, geologists in the company believe the wells are geologically dependent according to table 3

Table 3 Geological correlations

Wells↓→	1	2	3	4	5	6	7	8	9	10
1	1	0.1	0.2	0.1	0.2	0.2	0.1	0.2	0.2	0.2
2		1	0.2	0.3	0.4	0.3	0.3	0.4	0.3	0.3
3			1	0.1	0.2	0.1	0.1	0.2	0.4	0.4
4				1	0.1	0.2	0.1	0.2	0.2	0.2

5	1	0.1	0.3	0.4	0.3	0.3
6		1	0.1	0.2	0.1	0.1
7			1	0.2	0.1	0.1
8				1	0.2	0.2
9					1	0.3
10						1

APPENDIX B: SOFTWARE DETAILS

This section describes the operation of the Excel® VBA functions and subroutine. In brief, the software collects data from the spreadsheet and performs the Kullback-Leibler procedure on probabilities and correlations. It then constructs the double-binomial lattices for prices and outcomes, and finally, carries out the recursive algorithm. The result of these tasks is the value of group of prospects under the optimal exploration strategy.

As discussed before, the generalized modules in the program each perform a specific task and then pass the necessary arguments to the next units. This open-source structure allows users to manipulate and construct other special-purpose programs such as the sensitivity analysis subroutine we discussed in section 4. This appendix expands on this modular structure and provides further details on the specific functions and subroutines.

The program is composed of two main subroutines, “DP Main” and “KL Main”, that act like control centres; they have module-level variables that pass to other subroutines to perform tasks. In the first stage, “KL Main” calls a subroutine that collects the input data from the spreadsheet. All the data is stored in arrays and then passed to the subroutine that performs the Kullback-Leibler optimization using Excel’s Solver. The result is the joint probability distribution that is given to “DP Main” subroutine that solves the Markov decision process using our recursive algorithm. In fact, this second stage in the program calls subroutines that generate price lattice and the scenario probabilities, and passes these results to a subroutine called “DP Recursion” that performs the recursive valuation algorithm. The result of this process is the value of the optimal sequential strategy.

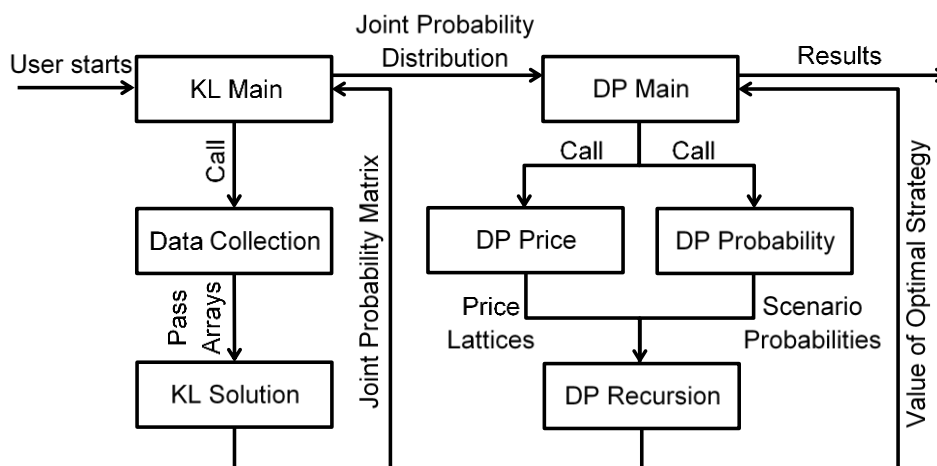


Figure 7 Software’s modules and their structure

The above flow chart shows the structure of the program and the series of tasks that the subroutines perform. Each subroutine may also utilize functions (not shown in the flowchart) that perform part of the processing. Table 2 shows a list of functions and subroutines in this software.

Table 4 list of functions and subroutines

Name	Type	Task
FCURVE()	Function	Returns the forward price $F_{0,T}$ based on parameters of the two-factor price process and time to maturity T .
Puthenu()	Function	Returns the conditional probability $P_{u u}$ in a dual lattice
Puthend()	Function	Returns the conditional probability $P_{d u}$ in a dual lattice
Pdthenu()	Function	Returns the conditional probability $P_{u d}$ in a dual lattice
Pdthend()	Function	Returns the conditional probability $P_{d d}$ in a dual lattice
KL Main	Subroutine	Main subroutine for Kullback-Leibler procedure
Data Collection	Subroutine	Collects input data from the spreadsheet and stores it in arrays.
KL Solution	Subroutine	Converts the correlations to pairwise joint probabilities, then arranges probabilities into arrays, transforms the arrays to the spreadsheet and runs the Excel Solver to complete the KL procedure.
DP Main	Subroutine	Main subroutine for solving the Markov decision process
DP Price	Subroutine	Constructs the three-dimensional arrays to represent the binomial lattice for χ_t and ξ_t . It then generates a forward curve for pairs of χ_t and ξ_t and calculates the then-NPV of discovery for each well.
DP Probability	Subroutine	Uses the joint probability matrix generated in previous stage and constructs the array for probability scenario.
DP Recursion	Subroutine	Uses arrays generated in previous stages and performs the backward recursion algorithm. Then returns the value of optimal sequential drilling and the first well in the sequence.